TU Delft

MSc thesis in circuit design

A Resonant Pulser Capable Of Generating High Voltage Pulse Train For Ultrasound Transducers

Yishi Liang (5657016)

November 2024

To obtain the degree of Master of Science in Microelectronics at the Delft University of Technology, to be defended publicly on Friday November 29th, 2024

TU Delft

Yishi Liang (5657016): *A Resonant Pulser Capable Of Generating High Voltage Pulse Train For Ultrasound Transducers* (2024)

The work in this thesis was carried out in the:

Electronic Instrumentation Laboratory
Delft University of Technology

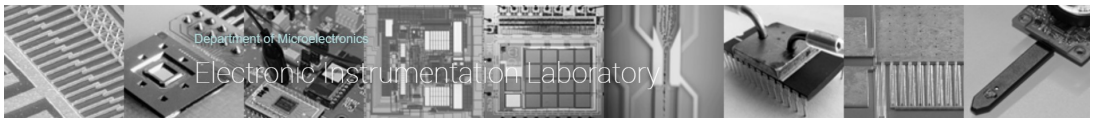

Supervisor:	dr. ir. M.A.P. Pertijs
Thesis committee:	dr. Q.Fan
	I.Bellouki

Abstract

High-voltage pulse train excitation for ultrasonic transducers offers exceptional signal-to-noise ratio (SNR) and enables the possibility of beamforming. Conventional high-voltage pulses cannot be directly generated from low-voltage batteries without a high-voltage source, requiring conversion to high-voltage DC before being shaped into pulses. Recent advancements in resonant pulser circuit architectures have made it feasible to convert low-voltage battery power directly into high-voltage pulses. However, these designs are limited to producing a single half-sine pulse. This work proposes a novel circuit architecture based on a 180nm BCD process, aiming to further enhance the performance of resonant pulsers and enable the generation of pulse trains.

Acknowledgements

First and foremost, I would like to express my heartfelt gratitude to Michiel Pertijs, particularly for his patience and professionalism as a mentor. I vividly remember how, 15 months ago, as I started searching for the thesis project, he dedicated four two-hour meetings to discuss potential directions within his research group that aligned with my study interests. Ultimately, he offered me the project I found to be very suitable. I will always keep in mind the opportunity he gave me and the substantial, invaluable time he invested in my work.

After nearly a decade of moving between universities, I've come to deeply appreciate that while excellent scholars are abundant in academia, they don't always translate into exceptional teachers. But over the past year, I've been fortunate enough to work with someone who embodies both.

My deepest gratitude also goes to Imad and Zehao. Imad, with his deep familiarity with the theme of my project, guided me through related past work during the year and inspired me to build on it. Beyond that, our coffee chats offered me insights into university life from an employee's perspective, and also free coffee. Thanks a lot, Imad. Zehao, from another perspective, brought extensive experience in analog circuit design. As my fellow student, he often provided sharp observations about potential transistor-level issues in my circuits and shared valuable opinions. I am so grateful for his insights and wish him all the best in China.

I also want to extend my thanks to Minghao, Junqiang, Xuanrui, the rest of my fellow students in the same office. We made a lot of what might be considered "boring" jokes, but I genuinely believe those are time of quality.

Although we didn't interact too much, I would also like to thank Nuriel, whose helpfulness stood out. He introduced me many tips for the circuit simulation environment at the beginning of the year and was always willing to spend 20 minutes elaborating on his knowledge for a question that might only take 20 seconds to answer.

Lastly, I want to thank my parents. Beyond financial support, some of my fondest memories from this time are of our hiking adventures in Bergen and Stavanger, and family dinners on the terrace, interrupted halfway by rain, only to be finished in a rush. These moments of shared laughter and resilience are invaluable to me and made our vacations extraordinarily joyful.

To my friends who stood by me during this journey, thank you for the moments of uncontrollable laughter and the times we cried together, only to laugh again. To the girls who crossed paths with me along the way, thank you for your companionship, for discussing love, and for genuinely supporting each other as we faced life's challenges.

Liang Yishi
19/11/2024

Acronyms

Contents

List of Figures

List of Tables

1 Introduction

Modern ultrasound systems are increasingly complex and powerful, prompting significant miniaturization efforts. Today, laptop-sized ultrasound devices can perform on par with traditional larger machines. The emergence of handheld devices, such as the GE Vscan, illustrates a trend toward highly integrated ultrasound imaging solutions[1]. The development of ultrasound medical applications aims to promote the use of portable and potentially wearable ultrasound devices, making them suitable for both hospital settings and home use. Advances in integrated circuits (ICs) have played a crucial role in this process, particularly with the rapid development of the CMOS industry, which has directed more research and development resources toward miniaturizing ultrasound systems. The goal is to integrate complete imaging functionality into a single Application-specific integrated circuit (ASIC). Achieving high integration of all necessary electronic components while maintaining excellent imaging quality has become the main challenge in this endeavor.

The principle of ultrasound imaging is based on pulse-echo measurements, where an electromechanical transducer emits and receives sound waves, and the propagation time of the wave is used to calculate the distance between the transducer and the imaging target. Ideally, the intensity of the received echo would depend solely on the impedance difference between the tissue and the reflecting interface. However, as ultrasound signals propagate through the human body, they exhibit exponential attenuation, causing the signal strength to diminish with increasing distance, which introduces certain challenges in imaging quality and signal processing. The attenuation characteristics in soft tissue follow the Beer-Lambert law[2], which is expressed as:

$$I(z) = I_0 \cdot e^{-\mu z} \tag{1.1}$$

where I(z) represents the signal intensity at the detector, I_0 is the initial intensity at the source, z is the distance between the source and the detector, and μ is the attenuation coefficient, which varies depending on the material.

To achieve the high signal-to-noise ratio (SNR) required for the receiving chain, a high excitation voltage of several tens of volts must be applied to the transducer. To maintain such high-voltage (HV) on the chip, the design of pulse generators typically employs specialized HV components, which occupy a significant portion of the chip area[3]. In each transmission cycle, usually only a single pulse or a short pulse sequence is used to excite the transducer elements, which generally have a center frequency in the range of several MHz.

In certain specific cases, continuous pulses may be executed for imaging purposes (at which point the system operates in full duplex mode). Additionally, the system often needs to

excite not only individual transducer elements but also multiple transducer elements in one-dimensional (1D) or two-dimensional (2D) arrays, allowing for 2D or 3D imaging. This presents new challenges for developing advanced technologies that can achieve a controlled plane of focused ultrasound waves through beamforming. Based on this, this thesis will specifically focus on a novel HV transmitting circuit designed to provide continuous HV pulse trains for future portable ultrasound ASICs to enable 2D or 3D imaging.

Compared to other medical imaging methods, miniature ultrasound offers an economical, non-invasive solution for diagnosing and monitoring patients, as well as for medical interventions. On-chip ultrasound imaging applications include bladder monitoring and intracardiac imaging through probes or catheters, such as intracardiac echocardiography (ICE). Previous miniature ultrasound imaging systems typically relied on a large external imaging system, which processed all ASIC data and supplied power to the ASIC. Today, integrated ultrasound systems are widely used in various non-invasive and point-of-care diagnostic scenarios; some applications require high-resolution, fast imaging technology like Doppler imaging, while others have more basic imaging requirements.

An important advancement in this field is the development of compact, portable ultrasound systems, with ultrasound patches being a common example (a example of ultrasound patches is shown in Figure1.1. These patches promise a cost-effective way to provide continuous patient monitoring or diagnosis without repeated intervention by medical personnel. However, developing these devices presents key challenges, one of which is the lack of available power sources capable of generating HV pulses or efficiently producing HV DC. To address this, a new circuit architecture has been proposed [4], which will be explained in detail in the next chapter.

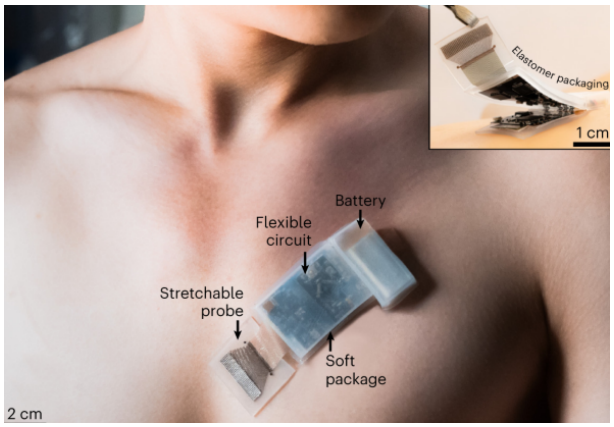

Figure 1.1: A photograph of the encapsulated ultrasound patch laminated on the chest for measuring cardiac activity via the parasternal window. The inset shows a folded circuit [5]

2 Prior art

2.1 Traditional ultrasound tranceivers

The block diagram of a traditional ultrasound transceiver is shown in Figure 2.1. First, the signal processor initiates the scanning process, generating a pulse signal or pulse train that is sent to a pulse generator. Then, the pulse generator drives the transducer, converting electrical energy into ultrasound waves. After emission, the ultrasound waves travel to the target object, which reflects the waves back to the transducer. The received echo signals are processed through the ultrasound receiver (RX) chain, which includes a Low Noise Amplifier (LNA), a Time Gain Compensation (TGC) amplifier, and an Analog-to-Digital Converter (ADC). The travel time of the ultrasound waves is directly related to the distance between the transducer and the target object, while the intensity of the echo signals depends on the acoustic impedance and distance of the target object. The signal processor constructs an image based on the information from the travel time and intensity of the echo signals (To create 2D or 3D images, a transducer array should be available to be excited)[6]. In the workflow of a traditional ultrasound transducer, the processes of ultrasound emission and reception, as well as the operations of the signal processor, work together to achieve imaging of the target object. It is worth noting that in the entire system, only the pulser requires HVDD power. This is because, as mentioned in the previous chapter, an excitation voltage of up to tens of volts is needed to drive the transducer, which LVDD cannot provide.

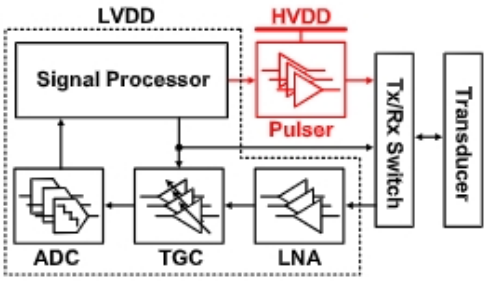

Figure 2.1: Block diagram of the ultrasound transceiver for conventional ultrasound imaging systems [7]

2.2 Model of the ultrasound transducer

Figure 2.2 shows the Butterworth-Van-Dyke (BVD) model of the ultrasound transducer, which consists of electrical components C_m, L_m, R_m, and C_p [8]. The BVD model includes ca-

pacitance C_m (motional mass), inductance L_m (compliance), resistance R_m (acoustic emission and motional loss), and parasitic capacitance C_p. When a pulse is applied to the transducer, electrical energy is converted into acoustic energy and flows into the motional branch composed of C_m, L_m, and R_m. In the frequency domain, the ultrasound transducer acts as a band-pass filter, typically operating at the fundamental frequency ω_f to minimize reactive power [1]. The dynamic power consumption of the transducer is primarily caused by C_p, if it dominates the resonant impedance, as the pulse generator needs to charge and discharge the transducer each cycle, leading to dynamic power loss. In most traditional probes, a matching network is often employed to minimize the impact of C_p. However, in the target application of this work, the matching network is omitted to reduce both size and cost[9]. Nevertheless, this power consumption does not affect image quality because only the power dissipated by R_m is related to acoustic emission. Due to the pulse generator charging and discharging the transducer every cycle, a dynamic power loss of $C_p \cdot HVDD^2 \cdot f$ occurs.

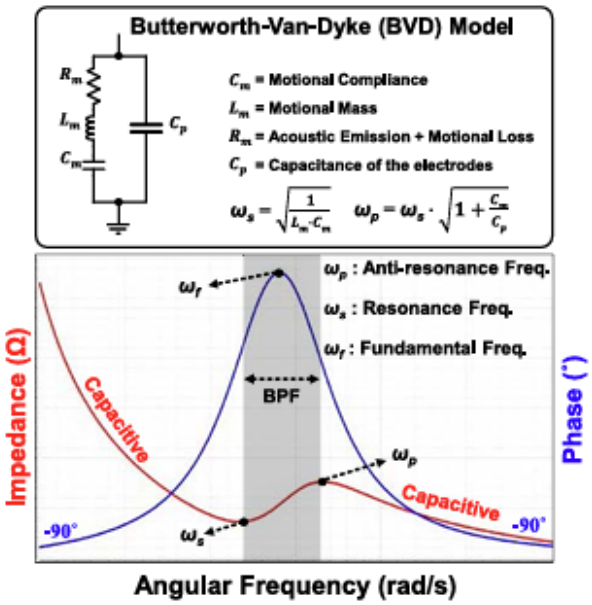

Figure 2.2: Equivalent circuit model of the ultrasound transducer and its impedance and phase characteristics in the frequency domain [7]

2.3 State of the art – HV pulser design

An important goal of the pulser design in this thesis is to achieve high energy efficiency. In the current state-of-the-art, several methods have already been proposed. In this following subsections 2.3.2 and 2.3.3, two methods will be introduced.

2.3.1 Conventional Class-D Pulser

Among the various pulse generator topologies, the switch-based Class D structure is the most widely used. Despite its very simple structure, pulse generators based on this topology can completely deplete the energy charged in C_p without causing heating issues [10]. However, this traditional design is still troubled by $C_p \cdot HVDD^2 \cdot f$ excessive losses caused by the large C_p driven by the high supply voltage.

As shown in the figure 2.3, the power switches in the Class D structure are implemented using large-size transistors to achieve HV isolation. Therefore, this structure not only occupies a large area but also results in high parasitic capacitance, leading to excessive dynamic power consumption [11]. This imposes a power burden on the gate driver. Even with specialized processes that provide HV transistors (such as DMOS, and LDMOS), which can withstand large drain-source voltages (V_{DS}), the allowable gate-source voltage (V_{GS}) of HV transistors is not large. Even when using HV CMOS that can withstand high (V_{GS}), higher gate drive losses can make the whole story fairly energy-inefficient[1]. Therefore, an additional level shifter is required to maintain V_{GS} = 5V, and this additional level shifter necessitates extra HV devices.

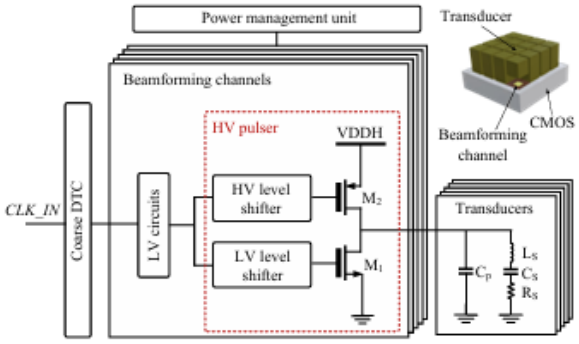

Figure 2.3: HV pulser based on the conventional class-D structure [12]

2.3.2 Multi-level pulsing

The first method to achieve high energy efficiency is known as multilevel pulsing, which improves power efficiency by using multiple voltage levels[1] to drive ultrasonic transducers. An N-level pulser uses $(N-1)$ regulated voltage sources to charge and discharge the capacitor in a stepwise manner, which reduces the wasted dynamic power to $\frac{CV^2f}{(N-1)}$ [13]. The energy savings stem from the charge recycling mechanism implemented during the discharge process, enabled by the regulated voltage supplies. In a conventional square wave scenario, all of the capacitor charge CV is discarded to ground. However, in the case of multi-level pulsing, when the capacitor switches from one voltage level to the next lower one, a portion of the charge $\frac{CV}{(N-1)}$ is recycled back to the power supply. This process can be repeated, allowing for up to $(N-2)$ charge packets of $\frac{CV}{(N-1)}$ to be recycled before the final packet is dumped to ground. As a result, the dynamic power is reduced by a factor of

$(N-1)$.

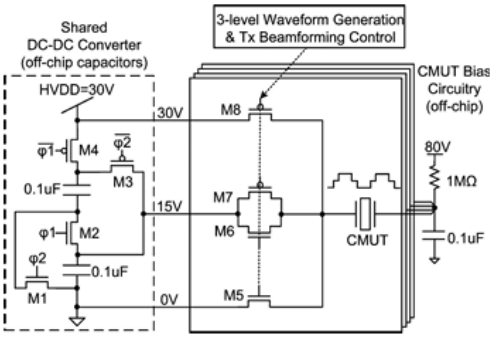

Figure 2.4: Circuit schematic of the four-channel three-level pulsers with the middle-voltage generation (all transistors are HV devices) [1]

As shown in Figure 2.4, pulse signals are divided into three voltage levels: 30V (HV), 15V, and 0V (ground), thus 50% of CV^2 is saved. The intermediate voltage of 15V is generated by a DC-DC converter that switches capacitors in a parallel-series configuration, which is shared among multiple channels. The only off-chip components in the entire system are two 0.1uF capacitors. Since the proposed three-level pulse generator has the characteristic of charge recycling and the Capacitive micromachined ultrasonic transducers (CMUT) load (approximately 40 pF per channel) is much smaller than 0.1uF, the converter can operate at a very low frequency (10-100 Hz), thereby saving power and consuming less than 1% of the total pulse power.

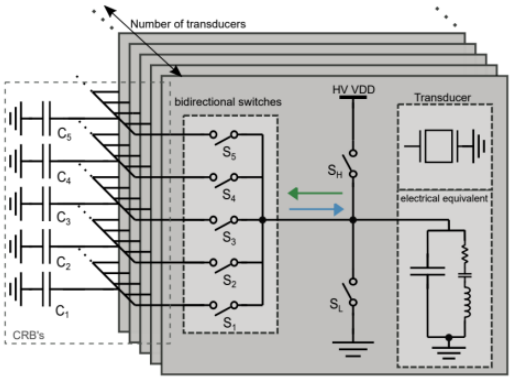

Figure 2.5: A 7-level pulser implementing large reservoir capacitors instead of generating intermediate voltage levels explicitly [14]

Furthermore, the Figure 2.5 illustrates a more advanced implementation of a 7-level pulser that achieves better efficiency, without explicitly generating 5 intermediate voltage levels[14].

2.3.3 Charge redistribution

Another method to improve efficiency is the charge redistribution technique. The general idea of the charge reuse unit is to reuse half of the charge in each driving cycle by separating the system ground from the array ground. In traditional TX driving, the ground electrode of the transducer is connected to the ground, while the driving electrode is mainly charged to VDDH. When the discharge mode starts, the driving and ground electrodes float and then short-circuit with each other. The charge quickly redistributes, causing both electrodes to approach VDDH/2 relative to the system ground.

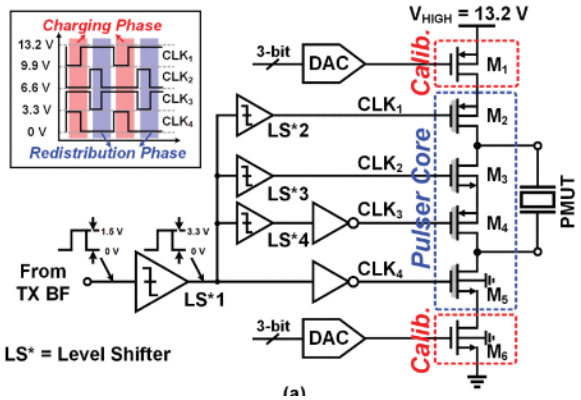

Figure 2.6: Charging redistribution scheme vs. conventional charging[15]

This charge reuse technique takes advantage of this phenomenon. During the next charging mode, each electrode charges/discharges by VDDH/2, thereby reusing half of the initial charge from the load capacitor. This energy-saving CR-TX technique achieves similar acoustic performance compared to traditional methods[16]. The works presented in [15] show implementations using this 'charge redistribution' on PMUT transducer's, this is also shown in Figure 2.6.

2.3.4 Buck-Boost-converter based waveform generation

In addition to achieving high energy efficiency, another goal is to enable the HV pulser to be powered by a low-voltage (LV) battery. This subsection briefly introduces a bidirectional buck-boost converter. As shown in Figure 2.7, it integrates a reverse buck function[17] with an inductor boost converter[18], providing advantages for forming an energy storage system and effectively managing energy storage and release to generate HV waveforms that can drive PZT transducers in an energy-efficient manner. However, its drawbacks are also evident. Apart from requiring multiple off-chip passive components, the major disadvantage of using this converter is that its operating frequency is much higher than the frequency of the generated waveform [17]. This approach is impractical for ultrasound imaging applications that require pulses in the MHz range.

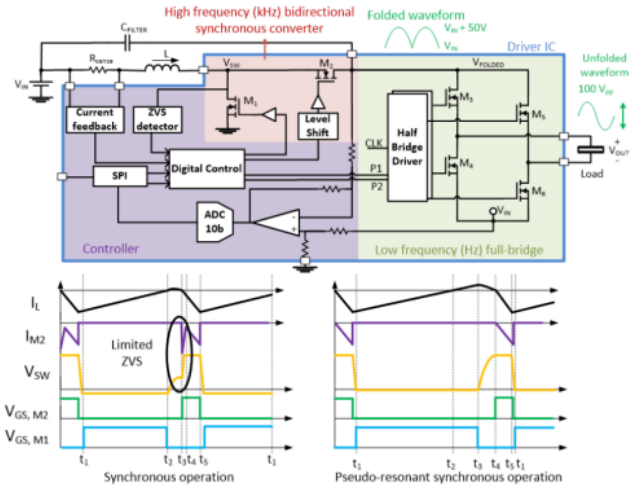

Figure 2.7: Architecture of a bidirectional synchronous buck-boost converter to drive a piezoelectric element at 100V peak-to-peak[18]

2.3.5 Energy efficient resonant HV pulser

Based on the pulser designs presented in the previous subsections, it could be readily asked whether it is possible to design a pulser that can directly convert LV battery supply into HV pulse with high energy efficiency, so that it could rely on fewer off-chip components and effectively reduce the losses associated with energy conversion. To answer this doubt, the idea of a resonant pulser was proposed and partially implemented recently[4]. It utilizes the resonance principle of an LC circuit to transfer electrical energy from the power supply to the transducer via an inductor and then recycle the transferred energy back again through the inductor after the pulse generation phase.

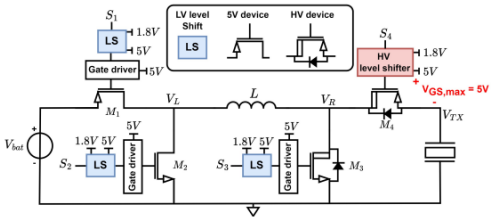

Figure 2.8: Simplified schematic of the proposed resonant pulser core, showing the minimum required 4 transistors needed for operation with an HV PMOS as high side switch[4]

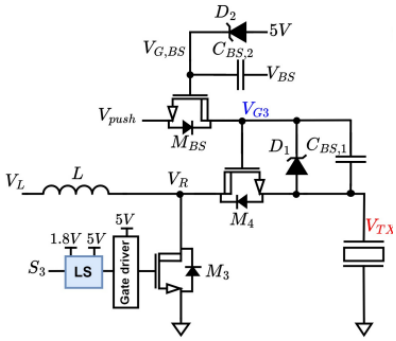

Figure 2.9: The detailed implementation for M4 in the left figure[4]

As shown in the figure 2.8 and the figure 2.9, the structure of this pulser is very similar to that of a traditional boost converter. The primary innovation here is the direct connection

of the power supply, off-chip inductor, and transducer, enabling direct conversion from LV DC to HV pulses. An off-chip inductor acts as a bridge between the power supply and the transducer. First, the transducer is disconnected from the circuit while the inductor is pre-charged by the power supply. Then, the power supply is disconnected, and the inductor and transducer resonate in an LC circuit to generate a HV sinusoidal pulse for the transducer. Vice versa, using a symmetric operation logic, once the resonant energy has fully transferred back to the inductor, precisely implemented switches enable complete energy recovery under ideal conditions. The specific operating principle will be explained in detail in the next chapter.

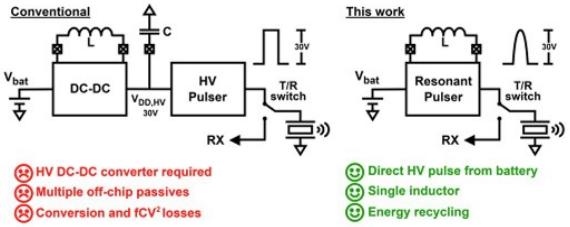

Figure 2.10: Envisioned battery-powered US patch and comparison of the resonant pulser to a conventional solution[4]

As shown in the figure, compared to a traditional class-D pulser, this idea significantly reduces the size of on-chip components and the number of off-chip components. Most importantly, it enables the recycling of transducer energy. However, due to limitations in switch design, this architecture cannot allow high resonant frequencies needed for generating pulse trains, therefore, the output is a single half-cycle sine wave, which makes it impossible to drive multiple transducer elements. This also means that TX beamforming, as required by the application context, cannot be achieved. To design a pulser capable of generating a pulse train, this work is carried out now.

3 Architecture-level study

3.1 Design objectives

The first step in circuit design is to clarify the design objectives. It is important to note that this is not a design from scratch; rather, it builds upon the circuit presented in sub-section 2.3.5 in the previous chapter. The goal is to address the shortcomings of that circuit and develop an improved version. As shown in the figure 3.1, if the objectives are broken down, the first version of the design successfully generated single half-sine HV pulse. Unfortunately, it was unable to generate pulse trains and could only drive one transducer element at a time, making the ultimate beam-forming goal impossible to achieve.

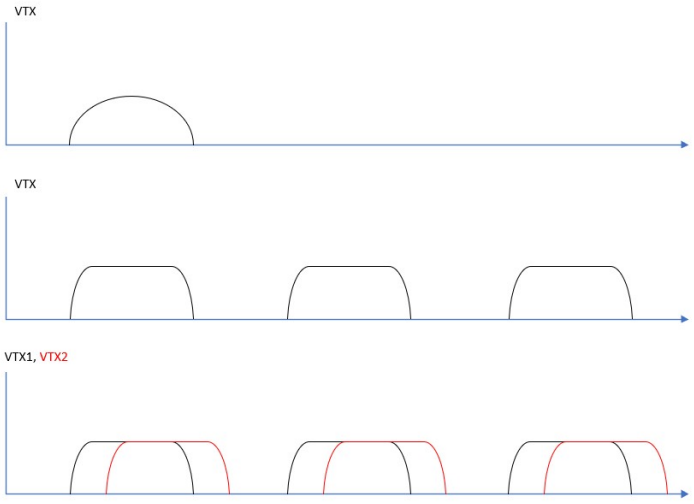

Figure 3.1: The objectives can be broken down into three stages, from top to bottom, First Stage (Previous Version), Second Stage (Current Goal), Third Stage (Ultimate Goal)

Using a half-sine wave to drive the transducer elements has two limitations. First, it makes achieving the beamforming goal impossible. Considering another design requirement—to adhere to using only one off-chip inductor—multiple transducer elements must be driven simultaneously (with time overlap) to enable beamforming. For this to happen, each transducer element must be able to ”float” at a high potential after being charged. This floating time allows other branches’ transducer elements to charge, ultimately allowing an array of elements to be driven with different phases. This state, where the transducer remains at a high potential, is called the ”hold phase”.

Second, due to the fact that the amplitude of the pressure wave generated by the transducer driven by a half-sine wave is much lower than that of a pulse train, the roundtrip SNR of the ultrasonic signal is also lower than that of the pulse train.

3.2 Operating principle

But what prevented the previous design from generating pulse trains? To answer this question, an ideal model of the resonant pulser and a simplified timing diagram will first be introduced.

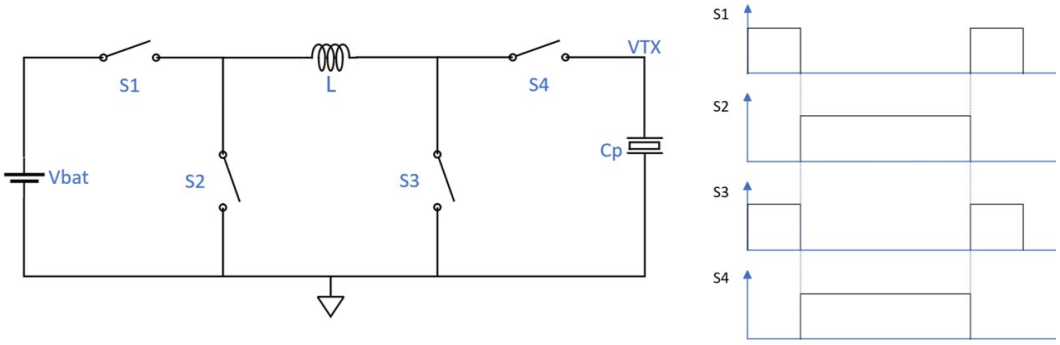

Figure 3.2: Simplified circuit model of the pulser core and its timing diagram in [4]

As shown in figures 3.2, the basic operating principle of the circuit will be detailed. There are four phases within one period. The first phase is called the inductive charging phase. In this phase, switches S1 and S3 are closed to charge the off-chip inductor. When the energy stored in the inductor is sufficient to charge the transducer and generate the required amplitude of the HV pulse, switches S1 and S3 are opened, and switches S2 and S4 are closed, beginning the second phase, known as the resonant energy transfer—rising edge. During this phase, the pre-charged inductor charges the transducer. When the energy has been fully transferred from the inductor to the transducer, the voltage pulse reaches its peak. At this point, the third phase begins, without changing the control of the switches. The energy originally transferred to the transducer returns to the inductor. When the voltage on the transducer drops to zero, this phase, known as the falling edge, ends. At this moment, switches S2 and S4 are opened, and switches S1 and S3 are closed again. The fourth phase, the inductive discharge phase, begins, where all the remaining energy in the inductor after the resonant energy transfer process is returned to the LV DC power source[4].

To perform an ideal quantitative analysis of the entire process, the equation:

$$V_{bat} = L \cdot \frac{di}{dt} \tag{3.1}$$

describes the charging rate during the inductive charging phase. The equation:

$$\frac{1}{2}LI_{pk}^2 = \frac{1}{2}C_pV_{TX}^2 \tag{3.2}$$

represents energy conservation during the resonant energy transfer phase. The equation:

$$f_0 = \frac{1}{2\pi\sqrt{LC_p}} \tag{3.3}$$

defines the resonant frequency. By combining these three equations, the equation:

$$V_{TX} = 2\pi T_{\phi_1} f_0 V_{bat} \tag{3.4}$$

is derived. It is evident that generating a HV pulse V_{TX} with a certain amplitude requires more inductor charging time T_{ϕ_1} as the energy resonant frequency f_0 decreases. In the previous design, the inductor charging time combined with the inductor discharging time needed for an resonant frequency of approximately 2.5 MHz far exceeded half of the transducer's resonant period. Returning to the question raised earlier, this means that after generating the first half-sine pulse, the off-chip inductor does not have enough time to fully discharge and recharge to produce a second HV sine pulse, making the generation of a pulse train impossible. Conversely, to generate a HV pulse train with the specified amplitude, the resonant frequency f_0 needs to be increased to reduce T_{ϕ_1}, however, this results in steeper rising and falling edges of the generated pulses, the half-period reserved for resonant energy transfer will have some surplus, and this can be effectively utilized for the pulser's "hold phase."

Returning to figure 3.1, the objectives can be divided into three stages. In the first stage, the previous design made the generation of HV pulses possible, but it only produced half-sine pulses. The goal of the second stage is to successfully design a pulser capable of generating a pulse train with the "hold phase." The ultimate goal of the third stage is to achieve a circuit design that can drive multiple transducer elements simultaneously (with phase difference that causes them to partially overlap), thereby enabling beamforming applications. The focus of this work is to accomplish the second stage.

To achieve the goal in the second phase, the basic circuit model of the main pulser doesn't require significant modifications; the main adjustments needed are in the operation timing of several core switches. The original circuit operation, which consisted of four phases, now needs to be expanded to five phases by adding a "hold phase"[4] between the rising and falling edges of the resonant energy transfer. To accomplish this, after the energy stored in the inductor has completely transferred to the transducer, S4 should be immediately opened, causing the transducer to remain at a high potential and be isolated from the left side of the circuit. At the same time, S3 should be closed, grounding both ends of the inductor through

S3 and S2, allowing any residual energy in the inductor to be dissipated. At the end of the "hold phase," S3 is opened again, S4 is closed, and the falling edge of the resonant energy transfer begins. This whole process can be reflected in figure 3.3.

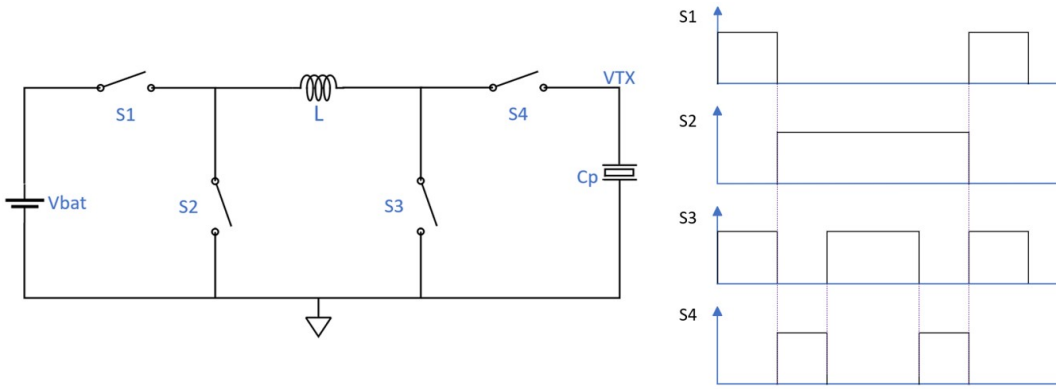

Figure 3.3: Simplified circuit model of the pulser core and its timing diagram in this work

3.3 Pulse width trimming circuit

This section and the next two sections will explain all the components of the control circuit and its architecture, as well as how its sub-blocks work together. Overall, the control circuit includes a pulse trimming circuit, an amplitude calibration circuit, and an energy recycling calibration circuit. The pulse width trimming circuit will be elaborated in this section, and the transistor-level implementation of some of its key blocks will be discussed in the next chapter.

The control of the core switches in the pulser is crucial. Accurately and timely controlling their on/off states determines whether pulses can be appropriately switched between different phases, which further affects the efficiency of the entire circuit and the feasibility of the architecture. In this high-speed pulser, various forms of delays, offsets, and process variations in the off-chip inductor and transducers significantly impact the timing control of the switches. Therefore, trimming and calibration in advance are essential.

3.3.1 Basic working principle

The HV pulse needs to be trimmed mainly due to the process variations of the off-chip inductor and transducers. Here, the tolerance of both the inductance and capacitance is assumed to be $\pm 20\%$. According to equation 3.3, this will have a significant impact on the resonant frequency f_0. In order to make the pulse timing and resonant frequency match each other, a trimming circuit needs to be implemented.

The trimming circuit only needs to be activated once at the initial stage of the first measurement. Once the resonant frequency is matched, the trimming circuit can be turned off and does not need to be activated again, as the inductance and capacitance do not change over time. To match the resonant frequency, the pulser first generates a test pulse based on a rough estimate of the timing. Simply put, this rough estimate can be the timing corresponding to the resonant frequency of an ideal inductor and transducers without any process variations.

By using a comparator to detect polarity of the actual voltage at the falling edge of the resonant energy conversion at the theoretical zero-crossing point of the voltage, it is determined whether the actual resonant frequency is higher or lower than the ideal resonant frequency. A Successive Approximation Register (SAR) then controls a Digital-to-Analog Converter (DAC) to adjust the behavior of the LC resonance, aligning the pulse timing with the resonant frequency. This process takes several cycles to complete. Theoretically, the falling edge and the rising edge are symmetrical, so trimming only the falling edge is sufficient to determine a unified DAC setting for both periods.

It is important to note that the transducer voltage needs to be attenuated before being sensed by the comparator because the transducer generates HV pulses which can be as high as 30V, while the comparator itself is powered by only 1.8V. Moreover, the comparator needs to be a continuous-time comparator, as clock-based sampling cannot accurately detect the zero-crossing in such a relatively high-speed pulse. The output of the comparator is connected to a SR latch to stabilize and sustain the output result from the comparator.

3.3.2 DAC selection

Due to the process variations mentioned earlier, the width of the rising and falling edges of the resonant energy transfer will change. To ensure that the actual switch operation timing adapts to the resonant frequency, there are essentially two strategies: the first is to adjust the timing to match the actual resonant frequency, and the second is to adjust the resonant frequency to match the given timing. The first strategy involves using some kind of time DAC (TDAC) based on delay cells, which can supply time for the rising and falling edges to match the time required by the actual resonant frequency (see figure 3.4).

The second strategy uses a CDAC, which is connected in parallel with the transducer to supply additional capacitance, thereby adjusting the resonant frequency to match the rising and falling edges of the given timing, similar to what was done in [4] (see figure 3.5).

When considering these two strategies, it quickly becomes apparent that the CDAC adds to the capacitance of the transducer, which means it can only decrease the resonant frequency to reach the lower limit of speed when both the inductor and transducer fluctuate by -20%. This implies that, to some extent, speed is sacrificed. In contrast, using the TDAC strategy, the compensation of the time cells are entirely based on the actual resonant frequency required, meaning the resonant frequency will maintain its actual value and fluctuate between the lower and upper limits. Thus, speed will not be affected by the DAC. However, it should be noted that when using the TDAC strategy, the initial test pulse can only be given the

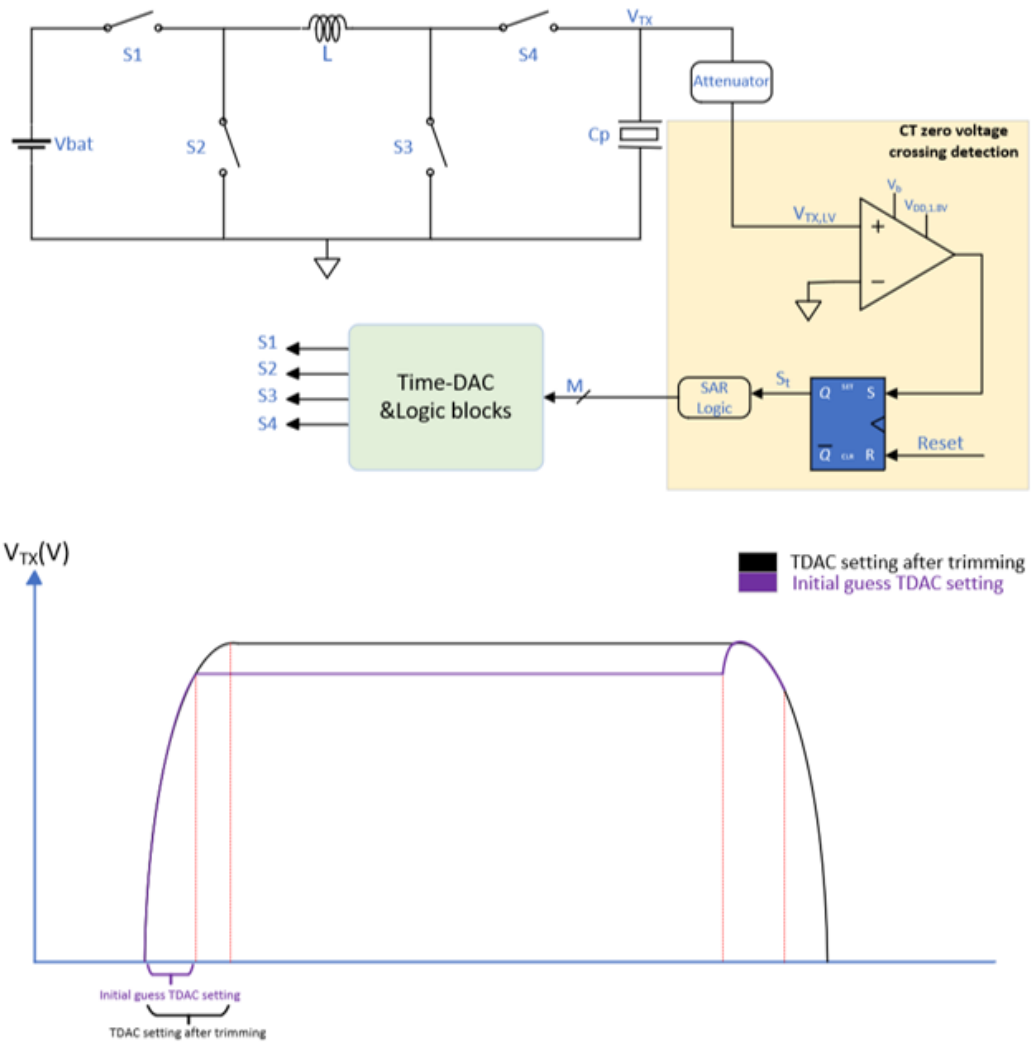

Figure 3.4: Simplified schematic and waveform of the transducer TDAC trimming circuit

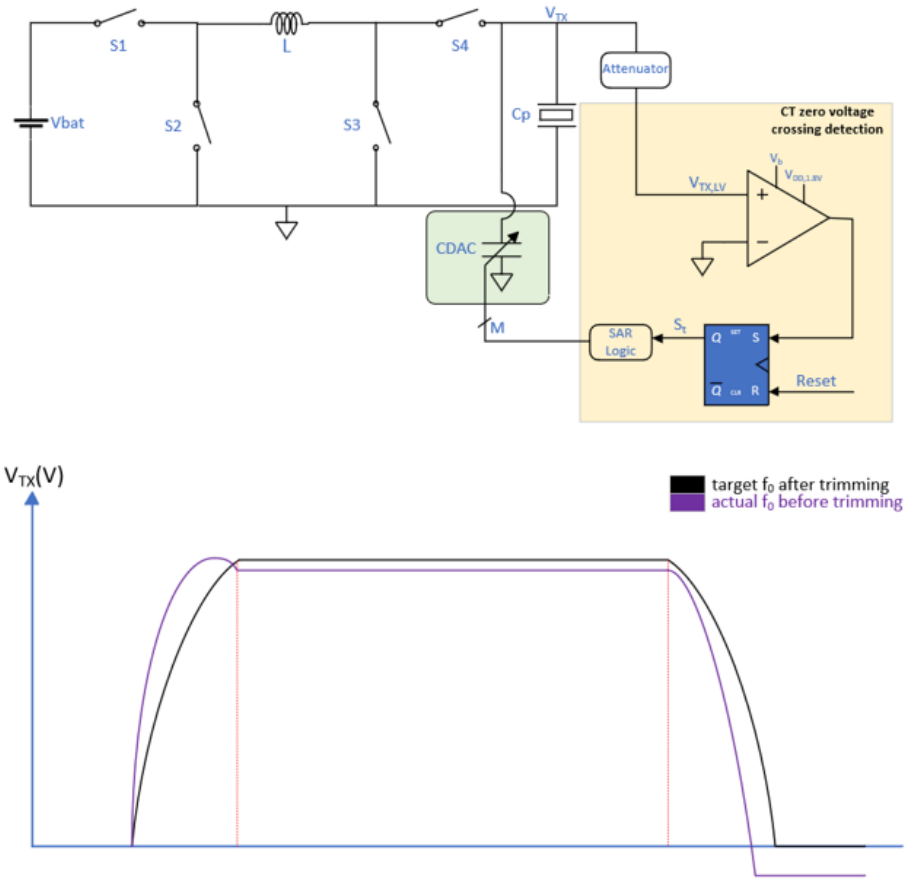

Figure 3.5: Simplified schematic and waveform of the transducer CDAC trimming circuit

shortest timing at the speed limit, because TDAC, by itself, can only provide additional time units and cannot reduce the time amount.

Another important metric when selecting the DAC is the resolution required for the DAC, which is crucial because higher resolution demands not only significantly increase the complexity of circuit implementation but also the number of cycles the trimming circuit needs to execute, leading to increased power consumption. To determine the resolution required for the two types of DACs, it is necessary to first establish the full scale and the minimum tolerable Least Significant Bit (LSB) for each DAC, which represents the maximum error. Determining the full scale is straightforward, it can be based on the maximum time units or the capacitance units that might need to be compensated in the fastest or slowest resonant frequency cases. The LSB, on the other hand, is constrained by the direction of the body diode of M4, which does not allow the falling edge voltage to drop below -0.7V. Ultimately, through estimation and simulation, both TDAC and CDAC require a resolution of 7 bits.

In terms of power consumption, the CDAC scheme will always consume more energy than

the TDAC scheme, because the TDAC does not alter the circuit's inherent resonant frequency, whereas the CDAC compensates the circuit with additional capacitance. This means the circuit has a larger capacitance load, leading to increased power consumption, while the TDAC only adds time cells to the circuit. Additionally, from the perspective of circuit implementation complexity and potential die area, the TDAC is also a more promising option. Therefore, trimming using a TDAC is selected. Furthermore, for similar considerations, only TDACs, rather than CDACs, are selected for all calibration and trimming circuits in this circuit design.

3.4 Amplitude calibration circuit

In addition to the timing control of the rising and falling edges, the inductive charging and discharging phases also need to be calibrated. The charging time is calibrated to ensure the required amplitude of the high-voltage pulse is accurately achieved, while the discharging time is calibrated to recycle as much of the LC resonance energy back into the battery as possible, with S1 immediately disconnected upon completion of the process to prevent any reverse recharging of the inductor. However, unlike the pulse width trimmed in section 3.2, these times are also affected by variation of V_{bat} (see equation 3.4), even if the resonance frequency is perfectly calibrated and the required high-voltage pulse amplitude is determined. Considering the inevitable energy losses, V_{bat} actually slowly decreases, so these two times need to be calibrated once before each of the measurement period.

The figure 3.6 presents a simplified schematic of the amplitude calibration circuit. As in the previous section, the same attenuator is employed here. After trimming the pulse width, the attenuated $V_{TX,LV}$ is fed into the peak amplitude detection comparator, where it is compared with V_{ref}(target pulse amplitude/attenuation factor). Based on the polarity of the comparator's output, it determines whether the inductor charging time should be increased or decreased. Similarly, the SAR logic controls the settings of the corresponding TDAC to stabilize the output pulse amplitude within the permissible error range. It should be noted that in this calibration circuit, a high-speed continuous-time comparator is no longer required; instead, only a single sampling and comparison at the end of each rising edge is necessary. Therefore, a dual-tail discrete-time comparator is used for this purpose, as implemented in [4].

3.5 Recycling calibration circuit

The energy recycling calibration is the final calibration step needed at the architectural level. In principle, it sets the corresponding TDAC by detecting the polarity of the voltage spike at the node connecting S1 and S2 at the end of the inductive discharging phase [19]. This process is iterated through SAR logic to ultimately achieve complete energy recycle within an acceptable error range, see figure 3.7. The detailed circuit implementation is elaborated in [4].

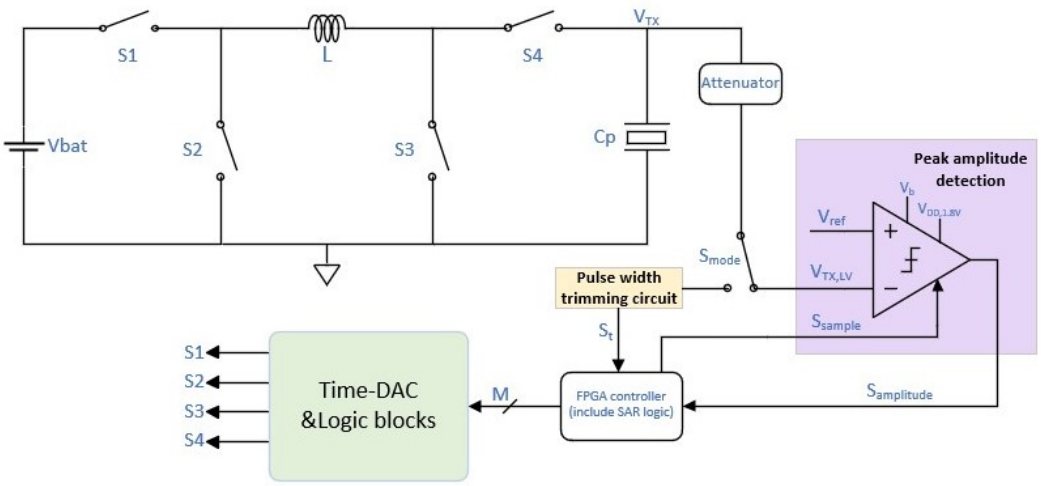

Figure 3.6: Simplified schematic of the amplitude calibration circuit

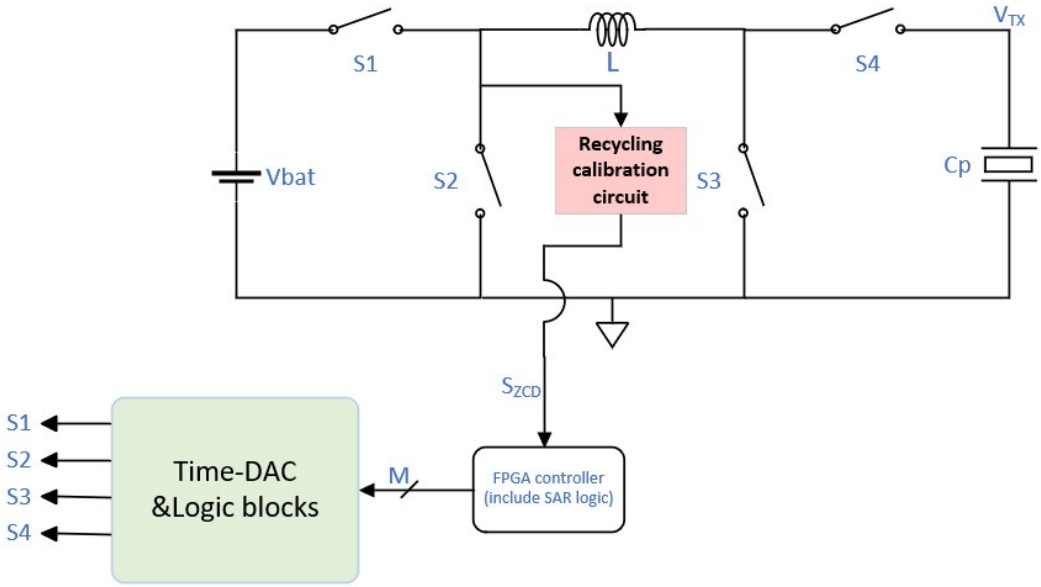

Figure 3.7: Simplified schematic of the recycling calibration circuit

4 Transistor-level circuit implementation

4.1 The core high-side HV switch

4.1.1 Implementation

As mentioned in the previous chapter, the circuit architecture of the main pulser doesn't require significant changes, but the operation principles of the switches need to be adjusted. As illustrated in figure 4.1, a simplified transistor-level circuit implementation is shown, where S1, S2, S3, and S4 are renamed as M1, M2, M3, and M4 at this level. M1, M2, M3, and all the level shifters and gate drivers that control them can be relatively easily implemented, simply the way they are implemented in [4], and the basic circuit architecture also remains consistent with previous version of design [4]. It is important to note that M3 is a 40V HV NMOS since its drain is connected to the HV node V_R.

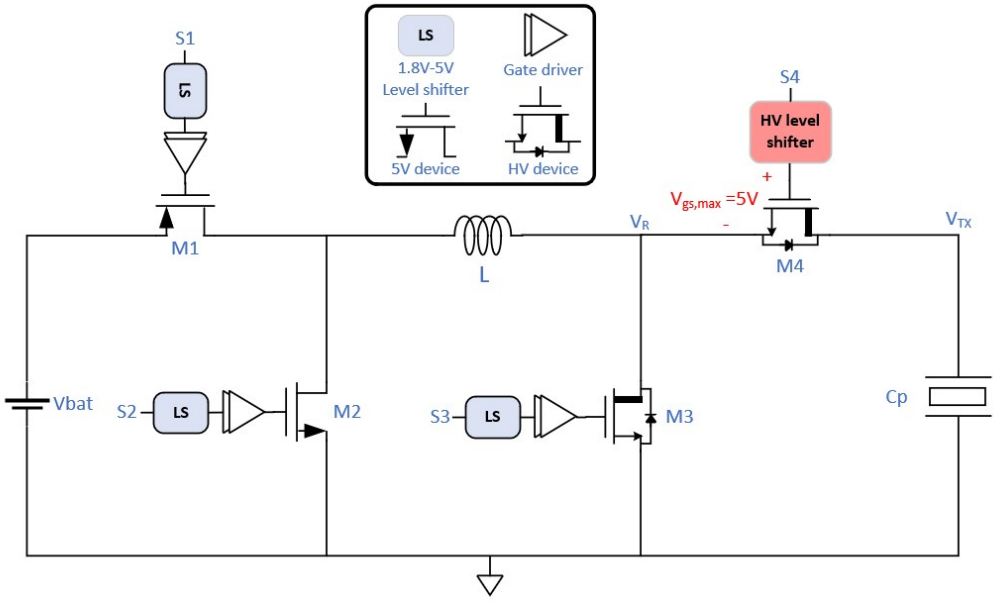

Figure 4.1: Simplified schematic of the proposed resonant pulser core, showing the minimum required 4 transistors needed for operation with an HV NMOS as high side switch

The real challenge in the pulser core circuit lies with the high-side HV switch M4 and its auxiliary driving circuit. Unlike the previous version of design, this switch needs to open at a high potential of 30V to enter the "hold phase" and then close again at the same 30V high

potential after the ”hold phase” ends to complete the falling edge of the resonant energy transfer. Considering the absence of a HV supply in the system, driving M4 is not a trivial task.

To address the driving issue of M4, the circuit architecture shown in the figure 4.2 is proposed. In order to allow M4 to remain on during both the rising and falling edges without relying on a HV supply, its gate V_g is bootstrapped to its source V_R using capacitor Cg. To turn the switch on, Cg needs to be charged to a voltage of about 5V. This is done by means of an additional HV NMOS M4-1. To turn it on near the end of the inductive charging phase, V_A is raised to 5V, and V_B to 10V, quickly charging node V_g. When the rising edge begins, V_B is lowered back to 5V, turning off M4-1 and leaving node V_g floating. At the same time, M3 also needs to be turned off, leaving node V_R floating. In this way, during the entire rising edge, V_R pulls V_g up, ideally maintaining a 5V voltage difference, ensuring that M4 remains in the required on-state during the rising edge phase [20].

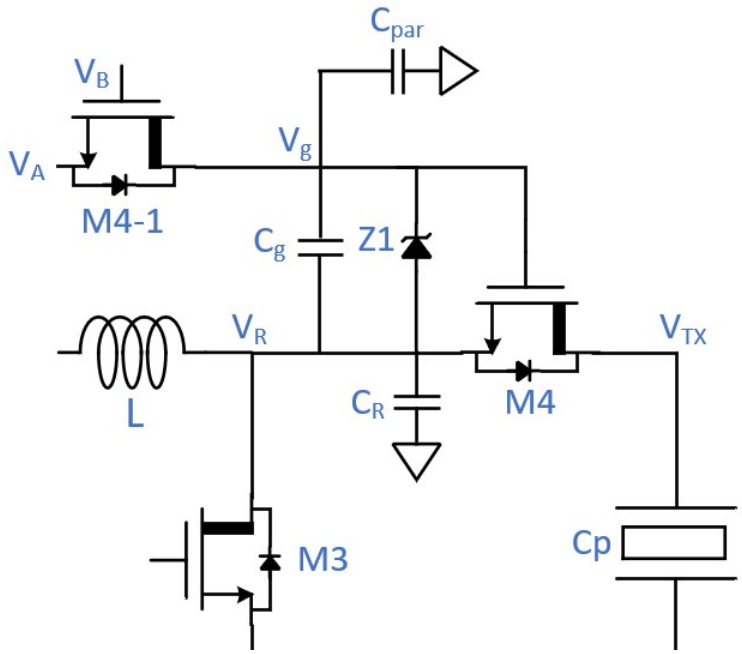

Figure 4.2: Proposed M4

When the energy in the inductor has been fully transferred to the transducer and the pulse reaches its peak amplitude, M4 should be turned off to allow the transducer to float at a high potential and enter the ”hold phase.” To achieve this, V_A is first lowered to 0V, which turns M4-1 back on and discharges node V_g. After discharging for a short time interval (typically, a few hundred picoseconds to one nanosecond), V_g drops below V_R, completely turning off M4. M3 is then turned back on, quickly dropping V_R back to 0V, C_g drags V_g down together with it and eventually also settles at around 0V. Subsequently, V_B is also lowered to 0V, and the circuit fully enters the ”hold phase”.

The circuit operation logic during the falling edge is essentially the same as that during the

rising edge. As shown in the timing diagram figure 4.3 for each control node within one period, V_g needs to be charged quickly. This is because, to fully turn on M4 and achieve a 5V gate-source voltage difference, M3 must remain on during this process to keep V_R at 0V. Otherwise, V_R would rise along with V_g, always preventing M4 from turning on. For the period before the falling edge, when V_g is being charged, the shorter the duration that M3 remains on while M4 is slowly turning on, the better. This is because the loop formed by them could leak energy, originally stored in the transducer and about to be returned to the inductor, to the ground. The longer this process takes, the less energy can be recycled. In this design, considering the high resonant energy transfer frequency, this duration should be limited to around 1 ns.

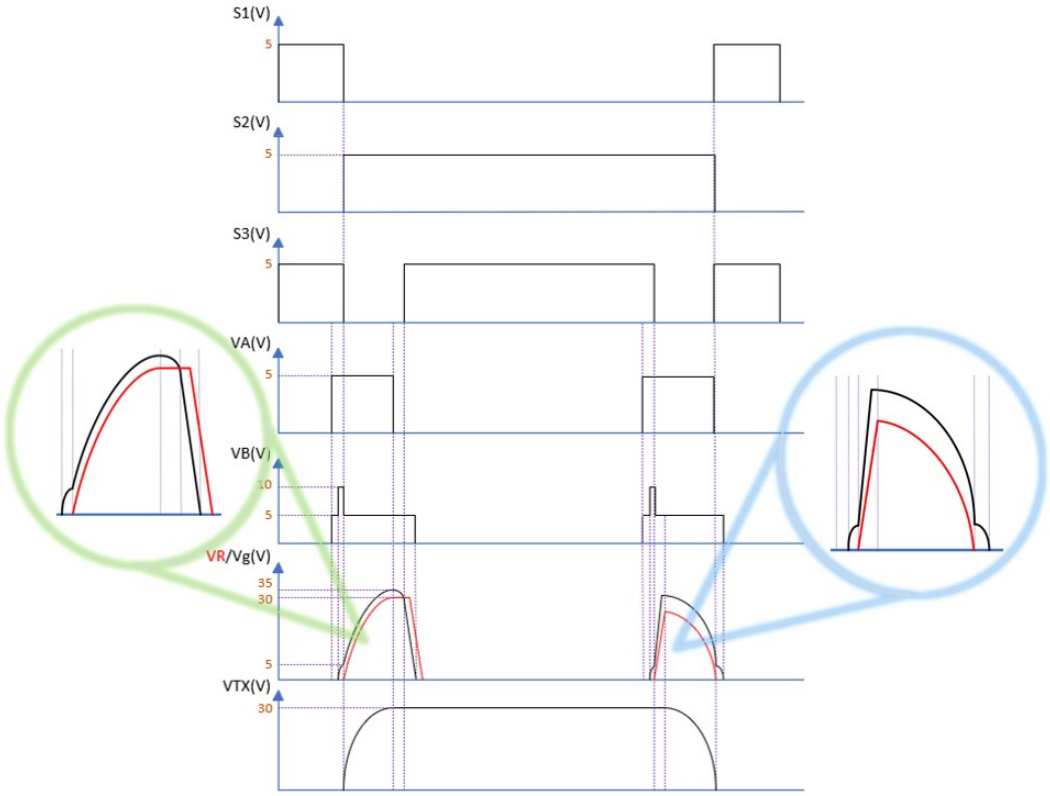

Figure 4.3: The timing of the main pulser with M4 implemented

It is evident that within one period, M4 needs to be turned on twice and turned off twice, while M4-1 needs to be turned on and off four times. Each time M4 is turned on or off, M4-1 must also be turned on and off once. Understanding this is crucial for estimating the energy losses of M4 and its auxiliary circuits and then finalizing their size.

It is worth noting that, unlike during the rising edge, V_A remains at 5V until the falling edge is completely over. During the initial tiny time period of the inductor discharging phase, M4-1 discharges V_g, with M3 already on. This helps to completely clear any residual energy stored in V_g and V_R. It might seem that M3 and M4 are again both on during a time, shorting

the transducer, but this is not an issue anymore, since, after the falling edge, there is fairly little residual energy left in the transducer.

Returning to figure 4.2, C_g and C_R are two additional capacitors that have been inserted, while the parasitic capacitance C_{par} at node Vg, primarily from M4-1, is shown grounded. To explain the purpose of C_g and C_R, one first need to focus on the equations 4.1,4.2 that describe the capacitive voltage division, which is a crucial phenomenon during the energy resonance process.

$$\Delta V_g = \frac{C_{gs,total}}{C_{gs,total} + C_{par}} \cdot \Delta V_R \tag{4.1}$$

During the rising edge, V_R drags V_g upwards. Due to the parasitic capacitance C_{par} at the V_g node, the change in V_g is actually smaller than the change in V_R, as proven in the equation 4.1. As a result, the initial 5V voltage difference between V_g and V_R does not remain constant during the rising edge—it will inevitably decrease. However, to ensure that M4 remains properly turned on throughout the rising edge, the total gate-source capacitance, $C_{gs,total}$, needs to be much larger than the parasitic capacitance, C_{par}. Since the parasitic capacitance between the gate and source of M4 alone is hard to be sufficient, an additional capacitor, Cg, is introduced.

$$\Delta V_r = \frac{C_{gs,total}}{C_{gs,total} + C_{r,total}} \cdot \Delta V_g \tag{4.2}$$

Equation 4.2 describes the capacitive voltage division that occurs right after the rising edge, when M4-1 discharges V_g. At this point, unlike the previous equation, V_g drags V_R down, so the change in V_R is always smaller than the change in V_g. The key difference between this process and the rising edge is that the capacitive voltage division during the rising edge should be avoided, as it effectively reduces the voltage stored between the gate and source of M4. However, the capacitive voltage division after the rising edge is beneficial and should to be utilized to some extent. Due to this non-ideal effect, V_g can potentially drop to the same or even a lower potential than V_R in a short time, which helps the pulser turn off M4 and enter the "hold phase." To achieve this, the total capacitance at V_R, $C_{r,total}$, must not be too small relative to $C_{gs,total}$. Therefore, an additional capacitor, C_R, is inserted to meet this requirement.

4.1.2 Auxiliary circuitry

Up to now, the bootstrap switch M4-1, used to drive the core switch M4, has been implemented. It is important to note that this needs to be a 45V HV NMOS, as the voltage at its drain, connected to V_g, potentially can be 5V higher than V_R and pulse amplitude in different corners can reach up to about 40V. The driving of V_A and V_B is another matter,

especially for V_B, which will switch between three voltage levels (0V, 5V, and 10V), implying improper driving sequences might cause M4-1 breakdown, making its implementation non-trivial. Based on this, the circuit shown in the figure 4.4 is proposed. Among them, M4-2 is a 12V HV NMOS.

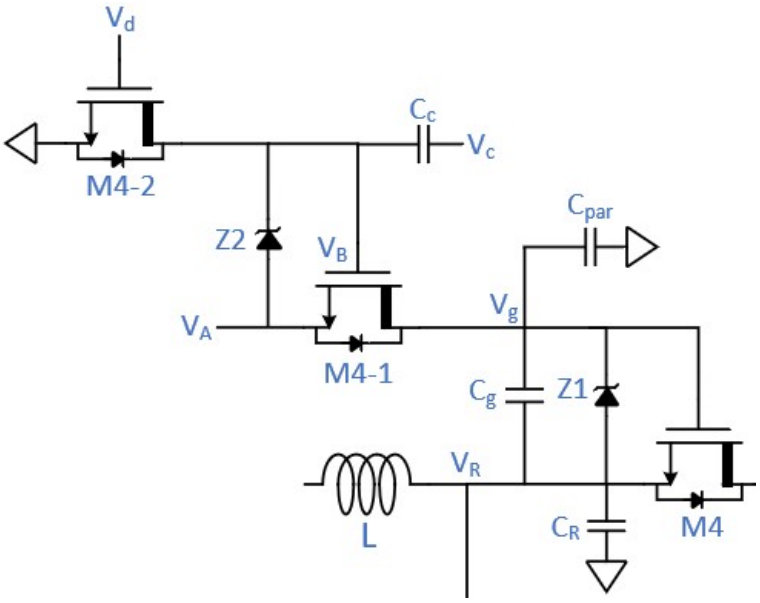

Figure 4.4: Proposed M4-1

As described in the previous subsection, just before the resonant energy transfer's rising edge begins, the LV logic signal V_A, controlled by TDACs and logic blocks, first starts its rising edge. At this point, due to the forward biasing of Zener diode Z2, V_B will also be raised to approximately 4.3V. Once V_B reaches 4.3V, V_c is immediately pulled from 0V to 5V to try to raise V_B to about 9.3V. This turns on M4-1, allowing the gate of the core switch M4 (V_g) to begin charging. After approximately 1ns of charging, V_c is again immediately dropped back to 0V, turning off M4-1, and the pulse enters its rising edge.

As shown in the figure 4.5, a 1ns short pulse needs to be generated for V_c just before the rising and falling edges of the resonant energy transfer. As mentioned in the previous subsection, this 1ns timing requirement is relatively strict, especially for the falling edge, as this duration will determine the energy efficiency of the pulser. Besides the strict limitation on the pulse width of V_c, the time delay between the rising edges of V_c and V_A is also strictly constrained. Theoretically, once V_B reaches around 4.3V, V_c needs to be immediately raised to 5V. Otherwise, the body diode of M4-1 will continue to be forward biased, causing V_g to be charged by the body diode, which is undesirable. Additionally, M4 might turn on even before the V_c pulse arrives, which would increase the time that M4 and M3 are both on to over 1ns, reducing the amount of recyclable energy during the falling edge and lowering the efficiency of the pulser, which is apparently a situation that should be avoided. It should also be noted that the time for V_B to be charged from 0V to 4.3V must be optimized to the shortest by adjusting the size of Z2, in order to minimize the impact of the inevitably forward biasing of M4-1 during this period.

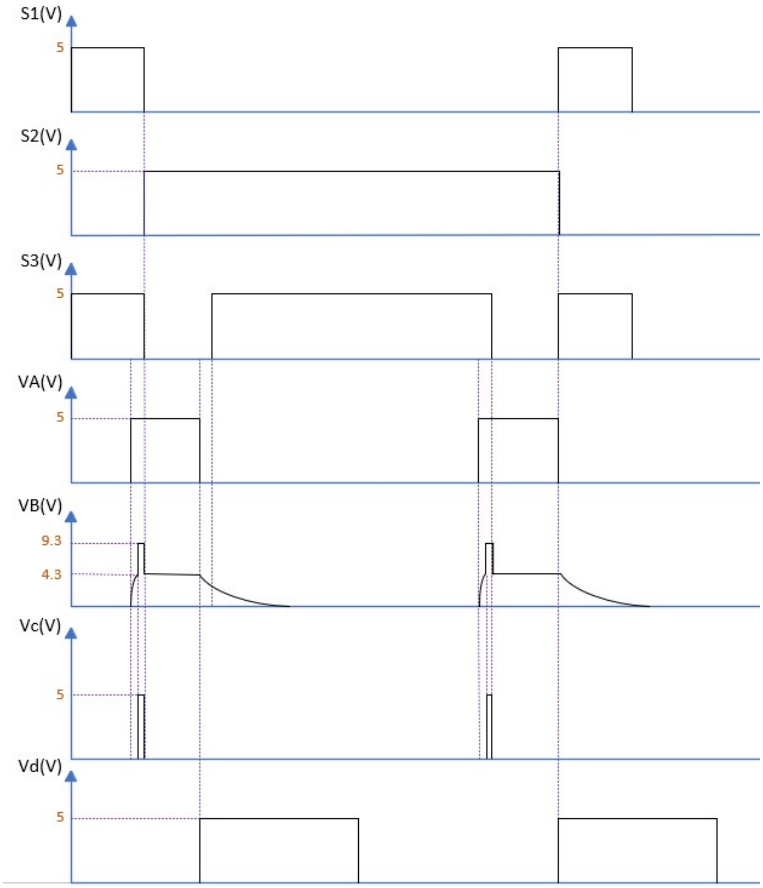

Figure 4.5: The timing of the main pulser with M4-1 implemented

On the other hand, an excessively short time delay between the rising edges of V_A and V_c can also lead to undesirable consequences. In this case, V_B may be prematurely forced to rise the other 5V before it reaches 4.3V, this would cause the maximum potential V_B can achieve to be lower, and M4-1 would start charging the gate of M4 without being fully turned on. This could potentially prevent M4 from fully turning on as well, increasing the on resistance during both the rising and falling edges of the pulse, thereby reducing the overall efficiency of the pulser as well. Thus, there is an optimal delay that maximizes energy efficiency.

4.1.3 Sizing determinations and trade-off considerations

Referring back to figure 4.2, as emphasized in the previous subsection, M4 needs to be turned on within 1 ns before the rising and falling edges. This is achieved by charging node V_g, and the charging speed of V_g is determined by the RC loop formed by the capacitance

at V_g (mainly C_g in pF level) and the on-resistance of M4-1. Here, to clarify the design specification, charging V_g from 0 V to above 4.65 V within 1 ns is defined as having "sufficient speed." To achieve this "sufficient speed", C_g needs to be relatively small, or M4-1 needs to be relatively large. If C_g is too small, based on equation 4.1, the difference between V_g and V_R (i.e. V_{gs} of M4) at the end of the rising edge might also be small, or even negative, making it insufficient for M4-1 to remain on during the entire rising edge. On the other hand, if M4-1 itself is too large, the parasitic capacitance C_p dominated by it will also be large, leading to a risk that M4-1 cannot stay on throughout the entire rising edge either. This implies that there is an optimal regarding the sizes of C_g and M4-1 that needs to be determined.

To determine the sizes of M4-1 and C_g, an initially arbitrary size is set for M4 (with a total width of around 1 mm). The smallest M4-1 that can provide "sufficient speed" to V_g is then identified across a range of C_g values from 4 pF to 12 pF. For each combination of M4-1 and C_g, the minimum V_{gs} that M4 can reach during the rising and falling edges is determined and compared to its threshold voltage (approximately 1.45 V). Combinations of M4-1 and C_g that cannot keep V_{gs} of M4 above 1.45 V throughout the whole period of both edges are discarded, and from the remaining combinations, the one with the lowest power consumption is chosen, as shown in the tables 4.1, 4.2. Finally, C_g is set to 10 pF, and M4 is set to a width of 300 µm. The V_{gs} of M4 can be charged to 4.65 V and 4.7 V within 1 ns before the rising and falling edges, respectively, and its minimum value at the end of each edge reaches 2.27 V and 1.89 V, both above 1.45 V. It is worth noting that C_R is simply set to $0.4 \times C_g$ to ensure a healthy ratio for the capacitive voltage division that occurs right after the rising edge, as described in equation 4.2.

Minimum M4-1 width for which V_g charging has "sufficient speed" (μm)	**C_g (pF)**
146	4
198	6
248	8
300	**10**
350	12

Table 4.1: Minimum M4-1 for a range of C_g (for which V_g charging has "sufficient speed")

C_g (pF)	**Minimum M4 V_{gs}/rising edge (V)**	**Minimum M4 V_{gs}/falling edge (V)**	**M4 V_{th} (V)**
4	0.85	1.52	1.45
6	1.57	1.62	1.45
8	1.99	1.71	1.45
10	**2.27**	**1.89**	**1.45**
12	2.46	2.11	1.45

Table 4.2: Minimum M4 V_{gs} during the rising/falling edges

Based on this, the size of M4 is further settled, mainly by considering the balance between switching losses and conduction losses to ensure minimal total energy loss. Finally, M4 is set to a total width of 1.4 mm. The sizes of M4-1 and C_g do not need to be redefined after determining the size of M4, as adjustments in the millimeter range for M4 will not significantly affect the ratios illustrated in equations 4.1, 4.2.

Another key design trade-off to consider lies in M4-2, C_c, and Z_2 (where C_c is an additional capacitor introduced to boost V_B when V_c rises). First, since the sizes of M4-1 and C_g have already been determined, the time required to charge V_B from 0 V to 4.3 V needs to be minimized, as described in the previous subsection, to minimize the risk of latch-up caused by M4-1 being forward biased during this period. This duration is mainly constrained by the RC loop formed by the in resistance of Z_2 and the capacitance of C_c, and it can be reduced by widening Z_2 or decreasing C_c. According to the equation,

$$\Delta V_B = \frac{C_c}{C_c + C_{gs,M4-1}} \cdot \Delta V_C \tag{4.3}$$

excessively sacrificing the size of C_c will significantly affect the voltage increase of V_B during V_c boost, preventing it from being raised to 9.3 V. However, an overly large Z_2 will also introduce more parasitic capacitance, increasing $C_{gs,M4-1}$, and this will also reduce the voltage increase of V_B during V_c boost. Therefore, an optimal for the sizes of C_c and Z_2 needs to be determined.

It is found that increasing the size of Z_2 beyond 150 μm no longer significantly enhances the speed at which V_B is charged to 4.3 V. Therefore, the total width of Z_2 is set at 150 μm. Meanwhile, the capacitance of C_c is increased to 10 pF to effectively dominate the larger parasitic capacitance between the gate and source of M4-1, ensuring that the ratio in the 4.3 remains healthy. It is worth noting that the width of M4-2 can be very small, as there is ample time for V_B to discharge after the rising and falling edges, as shown in figure 4.5, the total width of M4-2 is set to 10 μm.

4.2 Other switches on the main pulser

Similar to M4, the dimensions of M1 to M3 are also determined based on a balance between switching loss and conduction loss. Generally speaking, when the switch size is relatively small, the on-resistance is usually quite large, meaning the switch has substantial conduction losses that dominate the its energy consumption. However, as the switch size increases, the parasitic capacitance between the gate and source gradually increases as well, making it increasingly difficult for the switch to turn on, which in turn increases switching losses. When the switch size exceeds a certain value, the reduction in conduction loss brought by further increasing the switch size becomes insignificant, while the switching losses continue to rise, meaning total energy consumption also begins to increase and has already past the diminishing point, which indicates the optimal switch size.

Based on this, the sizes of M1 to M4 can all be determined. However, even so, a global estimate of total energy consumption is still necessary to execute, to observe how changes in the size of one switch affect the energy consumption of the others, as these switches more or less influence each other. In particular, when the total width of M3 increases from 2mm to 4mm, the average V_{gs} of M4 on the falling edge also increases, which in turn reduces the on resistance of M4, leading to a decrease in both the conduction energy of M4 and the total

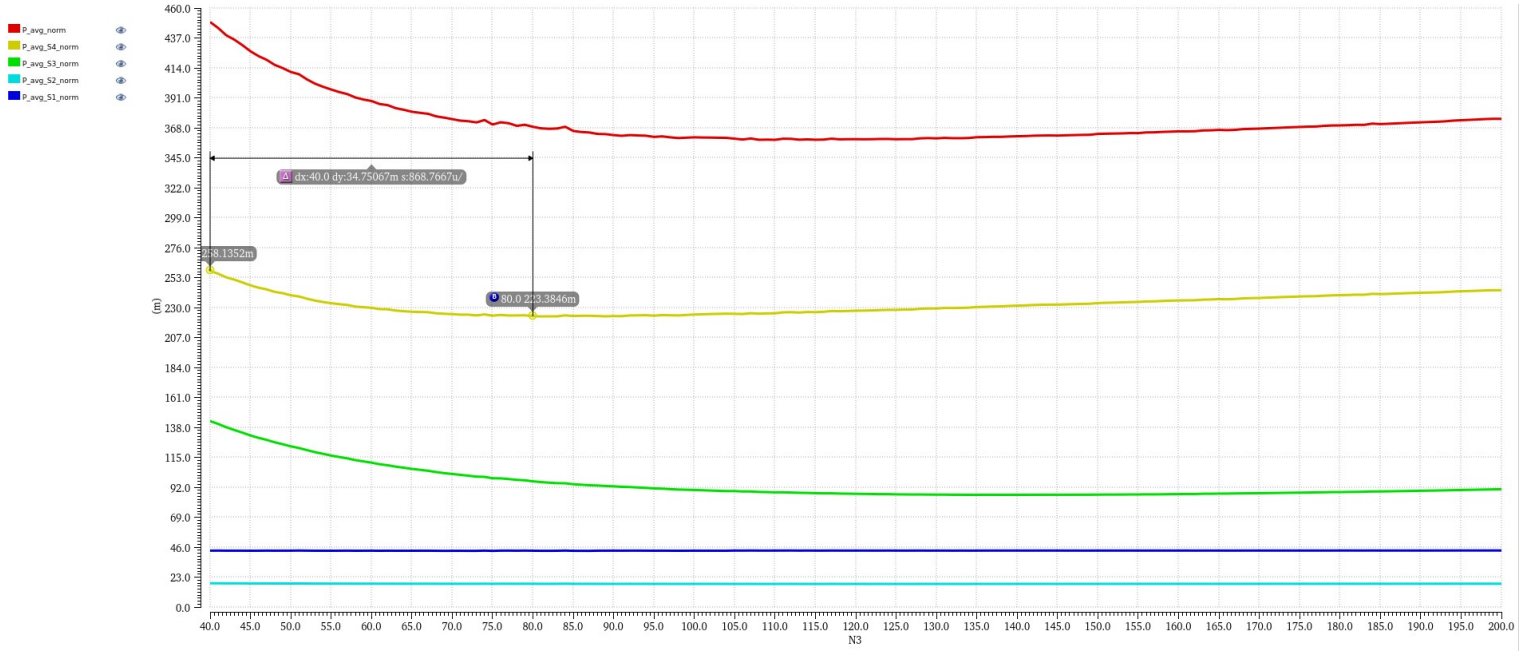

Figure 4.6: Effect of M3 width variation on energy consumption of each switches and of the total energy consumption

energy consumption, as shown in the figure 4.6 (N3: Number of fingers of M3, from 40 to 80 meaning total width from 2mm to 4mm). Ultimately, the sizes of M1 to M4 are set as shown in the table 4.3 to obtain a optimal total switches' energy consumption to approximately $0.362 \cdot f_0 \cdot Cp \cdot V_{TX}^2$. It should be noted that the total width of S1 is fairly large, but this is necessary because changes in the size of S1 have a significant impact on both its own and the overall pulser's energy consumption. This is also acceptable, as S1 is a low-voltage PMOS, and its actual die area can even be smaller than that of S4, which has only $\frac{1}{8}$ of its total width.

Switches number	Nominal energy consumption (%fCV^2)	Total width (mm)
M4	0.223	11.4
M3	0.097	3.04
M2	0.013	4.2
M1	0.042	1.4
Sum	0.362	

Table 4.3: Size and power consumption of each of the switches

4.3 Gate drivers

The gate drivers must be connected between the level shifters and the corresponding switches' nodes they control to enhance the drive strength of the 5V logic signal output from the level shifters. Typically, a gate driver increases its drive strength by cascading an even number

of inverters. Increasing either the inter-stage fan-out f or the number of stages N (i.e. the number of inverters) improves its driving strength. For a first estimate, f is set to 3 for all drivers. As shown in the equations,

$$N = \log_f(F) \quad \text{, where} \quad F = \frac{C_{gg}}{C_{in}} \tag{4.4}$$

by measuring the input and output capacitance of each driver, the required number of stages can be determined. Further fine-tuning of f around 3 is then performed to achieve an optimum between each driver's drive strength and its intrinsic delay. Table 4.4 presents the final settled fan-out and number of stages for all the gate drivers.

Gate driver control signal	Number of stages	Fan-out f	First stage NMOS/PMOS width (μm)
S1	6	3	2/4
S2	6	3.6	2/4
S3	6	3.4	2/4
VA	6	4.8	2/4
Vc	8	3	2/4
Vd	2	2	2/4

Table 4.4: Gate driver sizes of the resonant pulser core

4.4 Logic blocks and TDACs

A set of input signals from an off-chip FPGA with a 20 ns grids enters the chip and passes through logic blocks and TDACs to generate the 1.8V level shifter's front-end signals for S1, S2, S3, V_A, V_c, and V_d. This section will elaborate the implementation of the key TDACs and how they work in coordination with the logic blocks.

4.4.1 Additional trimming steps required to be executed in transistor level circuit

Due to the implementation of several transistor-level switches with stringent requirements, the trimming and calibration cycles described in the previous chapter are not sufficient to ensure the entire circuit operates with precise timing anymore. Therefore, three additional trimming steps need to be executed.

First of all, the behaviors of the rising and falling edges are not entirely symmetrical. During the rising edge, once M4 is turned on, the external inductor L immediately transfers its stored energy to the capacitor Cp through resonance. However, during the falling edge, the reverse energy recycling does not commence immediately after M4 is turned on. Prior to this, Cp needs to balance charges with C_R to equalize the voltages of V_{TX} and V_R. Subsequently, Cp and C_R together transfer the stored charge back to the external inductor L,

initiating the resonance process. As illustrated in figure 4.7, this indicates that the falling edge proceeds in two phases and always requires more time to complete than the rising edge.

Figure 4.7: Falling edge asymmetric to rising edge

Based on this, the rising edge and falling edge need to be trimmed separately; otherwise, the rising edge will never stop at the peak. The falling edge can still be trimmed based on the previously mentioned zero-voltage detection. However, the rising edge requires a new sensing block and algorithm. In fact, when the TDAC setting is too small on the rising edge, the circuit is forced into the "hold phase" before the resonant energy is fully transferred from L to Cp, meaning the rising edge cannot reach the peak amplitude. On the other hand, if the TDAC setting is too large, the resonant energy will partially return back to L after fully transferring to Cp, rather than immediately entering the "hold phase," which means that by the time it enters the "hold phase," V_{TX} is also below the peak amplitude. Therefore, the relationship between the TDAC setting for the rising edge and the pulse amplitude at the end of the rising edge is not monotonic. Thus, the SAR algorithm is no longer applicable for trimming the rising edge by observing the amplitude at the end of the rising edge, a linear search algorithm could be an option.

As amplitude calibration has not yet been performed at this stage, the target amplitude for trimming the rising edge is not fixed. If a linear algorithm is used, a sample/hold capacitor could be a feasible sensing block. At the end of each trimming cycle's rising edge, it samples V_{TX}, stores it through the capacitor, and at the end of the next cycle's rising edge, provides the comparator with a reference from the V_{TX} at the end of the current cycle's rising edge as the other input. A positive comparator output indicates that the TDAC setting is still too short. When the comparator output turns negative, the optimal setting is determined, and

the rising edge trimming is completed.

The other additional trimming steps that require to be executed are the time delay between the rising edges of V_A and V_c, as well as the pulse width of V_c itself. Both of these time intervals are set within the range of several hundred picoseconds to 1 nanosecond. Any PVT (process, voltage, temperature) variations in fixed delay cells, as well as variations from different process corners and propagation delays in the involved level shifters and gate drivers, could easily exceed the nominal value of the delay cell itself.

As discussed in the previous section, there exists an optimal time delay between the rising edges of V_A and V_c that maximizes energy efficiency. This optimal delay can be reflected in the duration of the inductive discharge phase. Consequently, recycling calibration must be embedded within this trimming step. By assessing whether the inductive discharge time increases or decreases over consecutive trimming cycles, the optimal can ultimately be searched. Additionally, the pulse width of V_c must also be tightly restricted. Given that excessively long or short pulse widths for V_c impact energy efficiency in the same way as the time delay between the rising edges of V_A and V_c, achieving an optimal pulse width requires an identical trimming approach.

4.4.2 Logic blocks and critical TDACs implementation

So far, all the trimming and calibration steps have been determined. TDACs need to be implemented to execute these tasks. They control the position of the edges of the system level logic input signals on the 20ns grids, working in conjunction with well-determined logic blocks and delay cells to output appropriate level shifter frontend signals for $S1$, $S2$, $S3$, V_A, V_c, and V_d. As shown in the figures 4.8 and 4.9, five 20ns grid signals — Charge, Recycle, Pulse, Sel_rise&fall, and Sel_S1&S2 — are invoked from the off-chip FPGA and input into the system. Five TDACs are employed to complete a total of six trimming/calibration cycles. The reason for saving one cycle is that the full range and resolution required by the TDAC for both the rising and falling edges can be the same, allowing the same TDAC (TDAC_rise&fall) to be used for both edges, with different settings applied for these two. A 10ns Min_rise&fall delay cell is inserted before the TDAC_rise&fall, as a single rising or falling edge within the tolerance of the resonant frequency can never be finished within 10ns. This fixed delay cell helps prevent the TDAC from wasting search range in unnecessary areas during the setting process.

To prevent the M4 body diode from forward biasing during the transition between the falling edge and the inductive discharging phase, the voltage difference caused by one LSB of the TDAC_rise&fall near the V_{TX} zero voltage crossing should be sufficiently less than 0.7V. Ultimately, its LSB is set to 160ps with a 7-bit resolution. The LSB and resolution for TDAC_charge and TDAC_recycle are set to 3ns and 3 bits, respectively, to ensure amplitude calibration and recycling calibration accuracy within 1V and 10mA.

The greatest implementation challenge lies in the TDAC_Vc_delay, used to control the time delay between the rising edges of V_A and V_c, and the TDAC_Vc_pulse, used to control

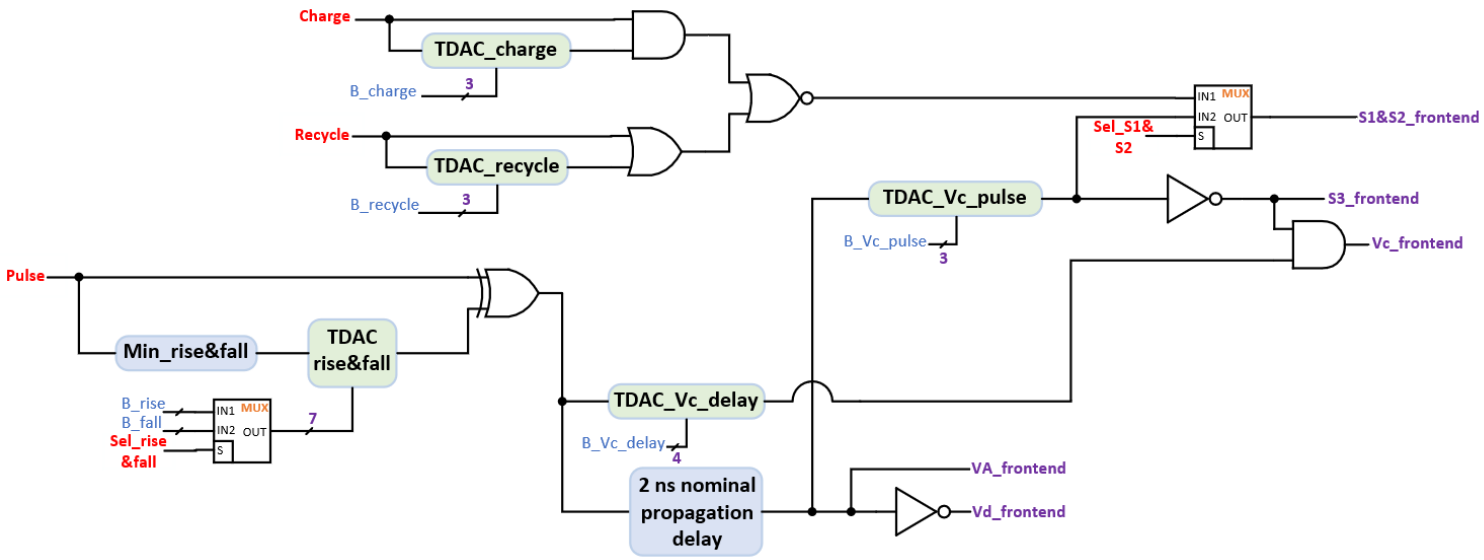

Figure 4.8: TDACs and logic blocks system level work flow

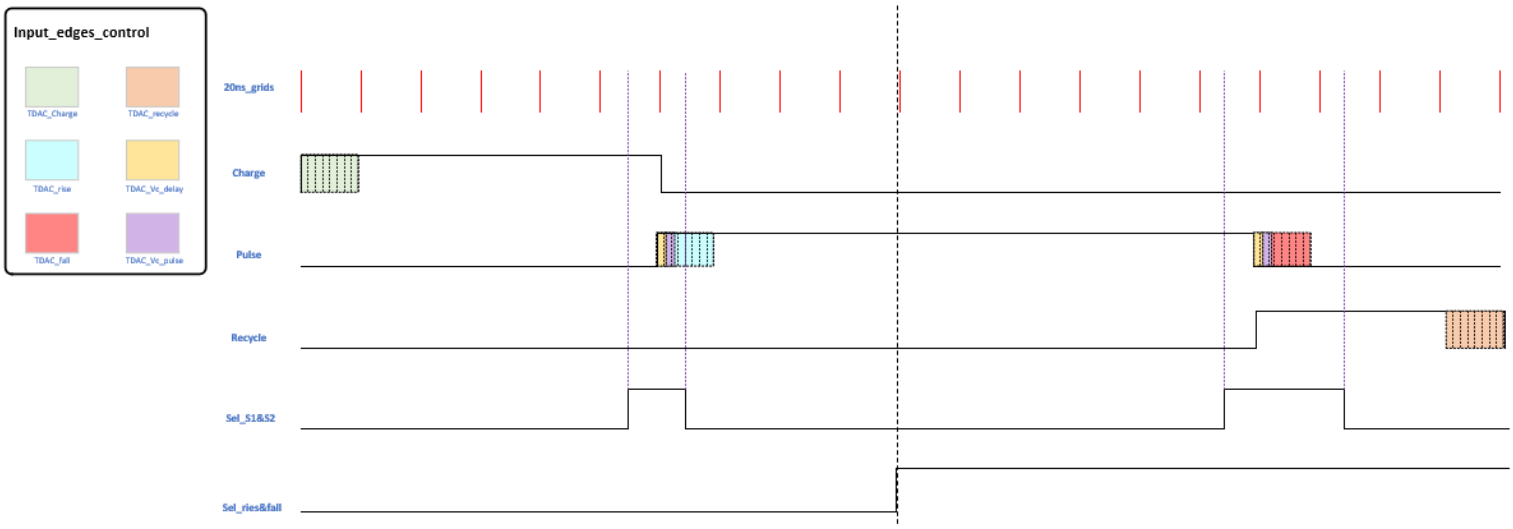

Figure 4.9: 20ns grids system logic inputs and the edges controlled by TDACs

the pulse width of V_c, see figure 4.10. Especially in TDAC_Vc_delay, considering that the charging time of V_B from 0 to 4.3V has been minimized to around 300ps, meaning that the TDAC_Vc_delay must be capable of fine-tuning within this time range.

In 180nm BCD technology, the inherent propagation delay of a precisely designed delay cell, combined with the delays contributed by the DAC's internal logic components (such as decoders and Multiplexer (MUX)), can easily exceed 1ns. This makes it impossible to achieve a delay unit that can fine-tune within approximately 300ps. Therefore, an additional nominal

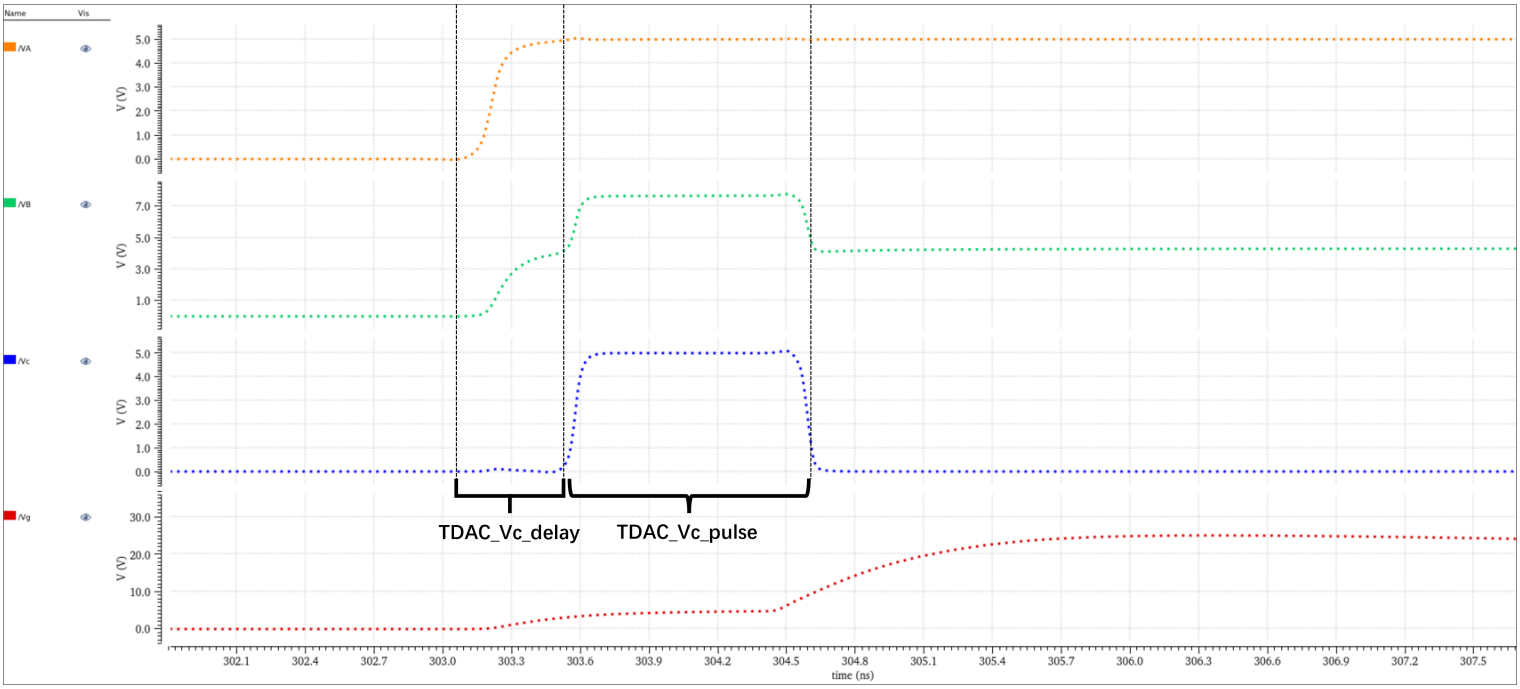

Figure 4.10: The edges of VA, VB, and Vc controlled by TDACs

propagation delay cell with a delay of 2ns is inserted. In this way, it only requires a delay unit capable of fine-tuning around 2.3ns, which, when placed in parallel, allows the output rising edges of V_A and V_c to have a time difference of around 300ps, see figure 4.11.

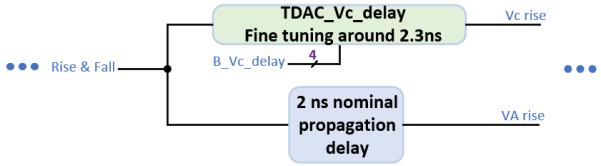

Figure 4.11: Schematic of the realization of a time delay of about 300ps

In addition, to make the TDAC_Vc_delay can be essentially fine-tuned around 2.3ns, the step size must be sufficiently small, targeting 15ps-20ps. Clearly, using two traditional cascaded inverters to build up a delay cell at this level is also impossible. However, as shown in the figure 4.12, by converting the binary input of the DAC into one-hot code and, based on this, selecting the branch with the PMOS having different width to connect to the circuit, the TDAC can provide different delay cells with fairly small time differences between them.

With 16 one-hot code branches, the PMOS widths decrease in steps of 10nm, from 610nm down to 460nm. This implementation allows the TDAC output to achieve fine-tuning in the range of 2.2ns to 2.45ns, with a resolution of 4 and a step size of approximately 15ps in typical corner, see 4.13. The TDAC_Vc_pulse is implemented in a similar manner, with relatively relaxed requirements, the resolution is ultimately set to 3 bits, with an LSB of 70ps,

allowing fine-tuning of the V_c pulse width between 0.7ns and 1.2ns.

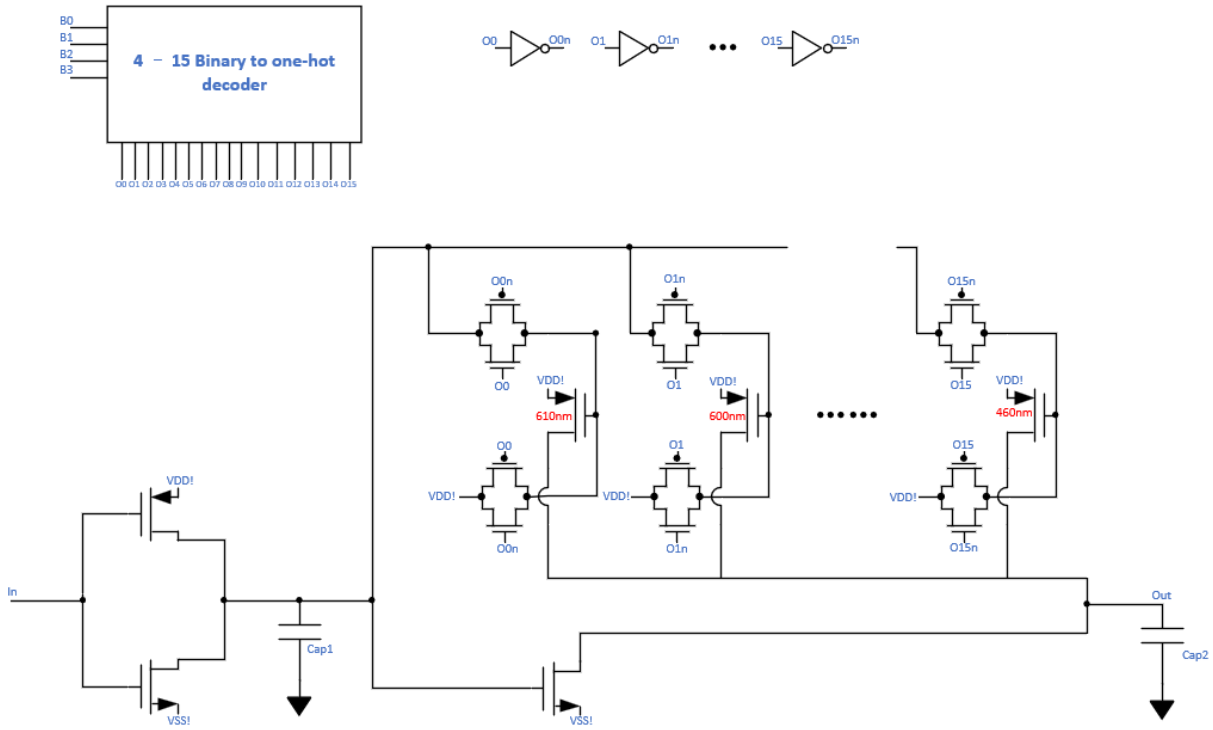

Figure 4.12: TDAC_Vc_delay implementation

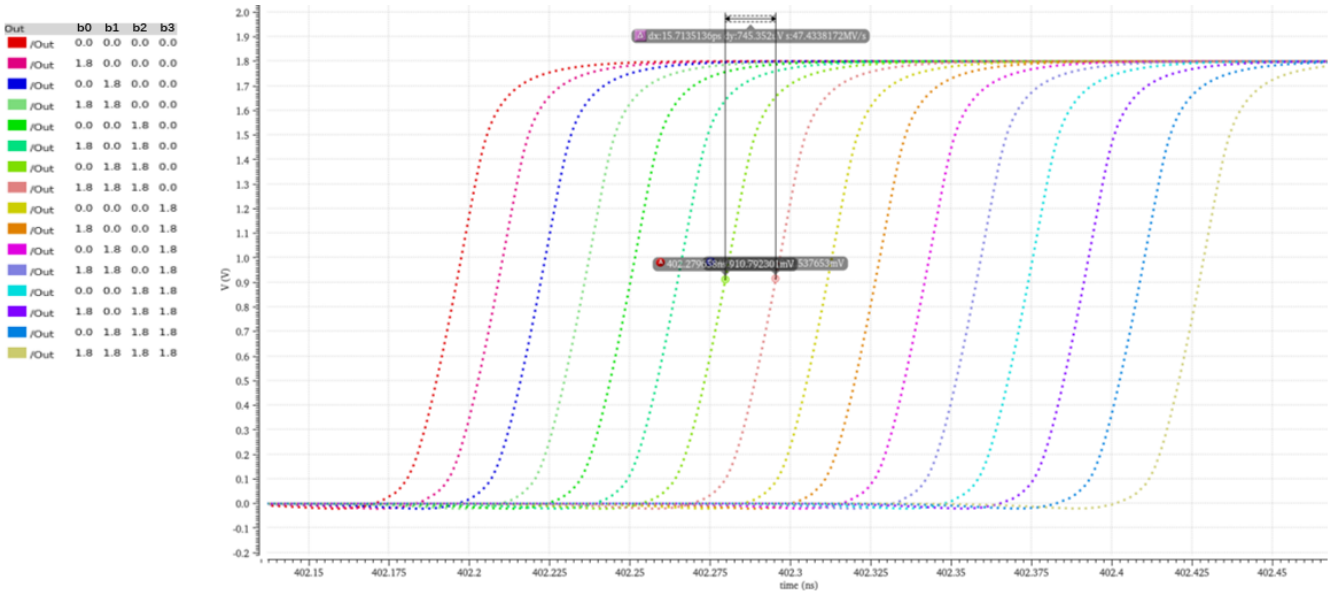

Figure 4.13: TDAC_Vc_delay simulation performance, achieved a step size of 15ps

5 Simulation results

5.1 TDACs to generate a HV pulse train

The TDACs complete all trimming and calibration steps by adjusting the (rising/falling) edge positions of S1, S2, S3, V_A, V_c, and V_d, ultimately enabling the pulser to produce an appropriate 30V pulse train. This section demonstrates the performance of all the TDACs at the transistor level implemented circuit through simulation results.

Set the TDAC_charge to set the position of the edges of S1 and S2, thereby setting the inductor charging time. Adjust the settings by observing the pulse amplitude at the end of rising edge. The step size of TDAC can reach 3 ns. Eventually, V_{TX} is calibrated to be sufficiently close to 30V (green curve), see figure 5.1.

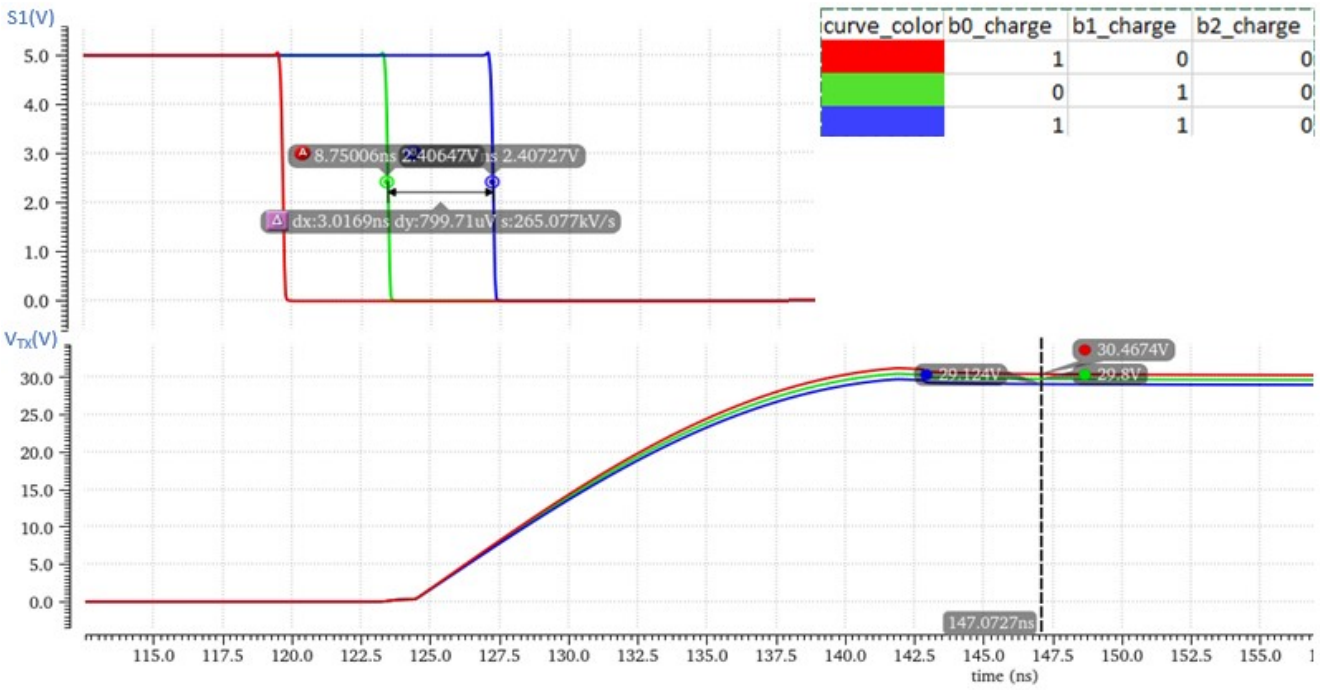

Figure 5.1: Adjust the TDAC_charge setting to shift S1 then further modify the time of inductor charging phase

Set the TDAC_recycle to set the position of the edges of S1 and S2, thereby setting the inductor discharging time. Adjust the settings by observing the inductor current at the end of recycling phase. The step size of TDAC can reach 3 ns. Eventually, I_L is calibrated to be sufficiently close to 0A (green curve), see figure 5.2.

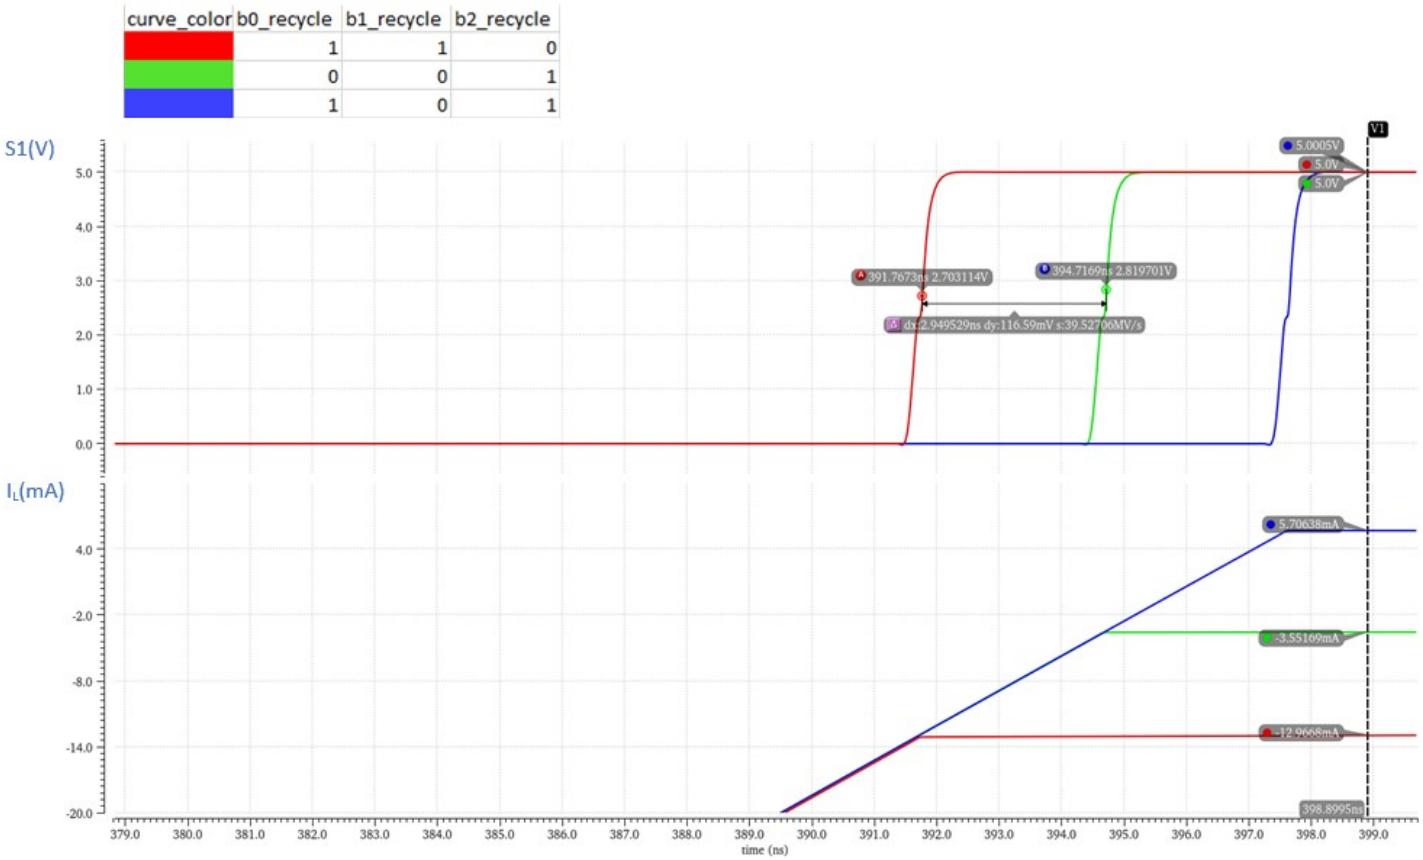

Figure 5.2: Adjust the TDAC_recycle setting to shift S1 then further modify the time of inductor discharging phase

Set the TDAC_rise to set the position of the edges of S3 and V_A, thereby setting the time of rising edge. Adjust the settings by observing the pulse amplitude at the end of rising edge. The step size of TDAC can reach 160 ps. Eventually, V_{TX} is trimmed to its highest possible value (green curve), see figure 5.3.

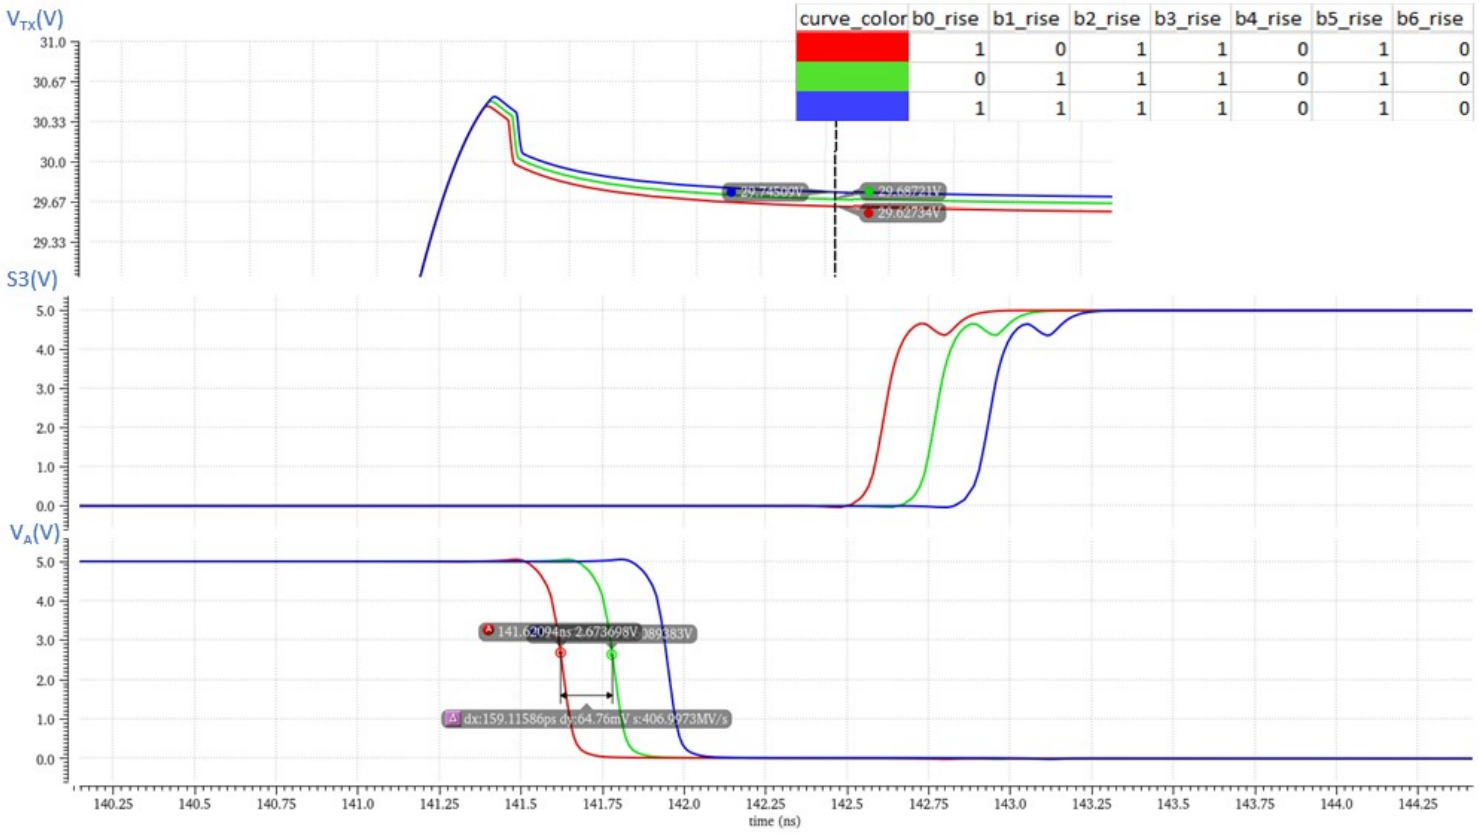

Figure 5.3: Adjust the TDAC_rise setting to shift S3, V_A then further modify the time of rising edge

Set the TDAC_fall to set the position of the edges of S2 and S3, thereby setting the time of falling edge. Adjust the settings by observing the V_{TX} at the end of falling edge. The step size of TDAC can reach 160 ps. Eventually, V_{TX} is trimmed to sufficiently close to 0V (green curve), see figure 5.4.

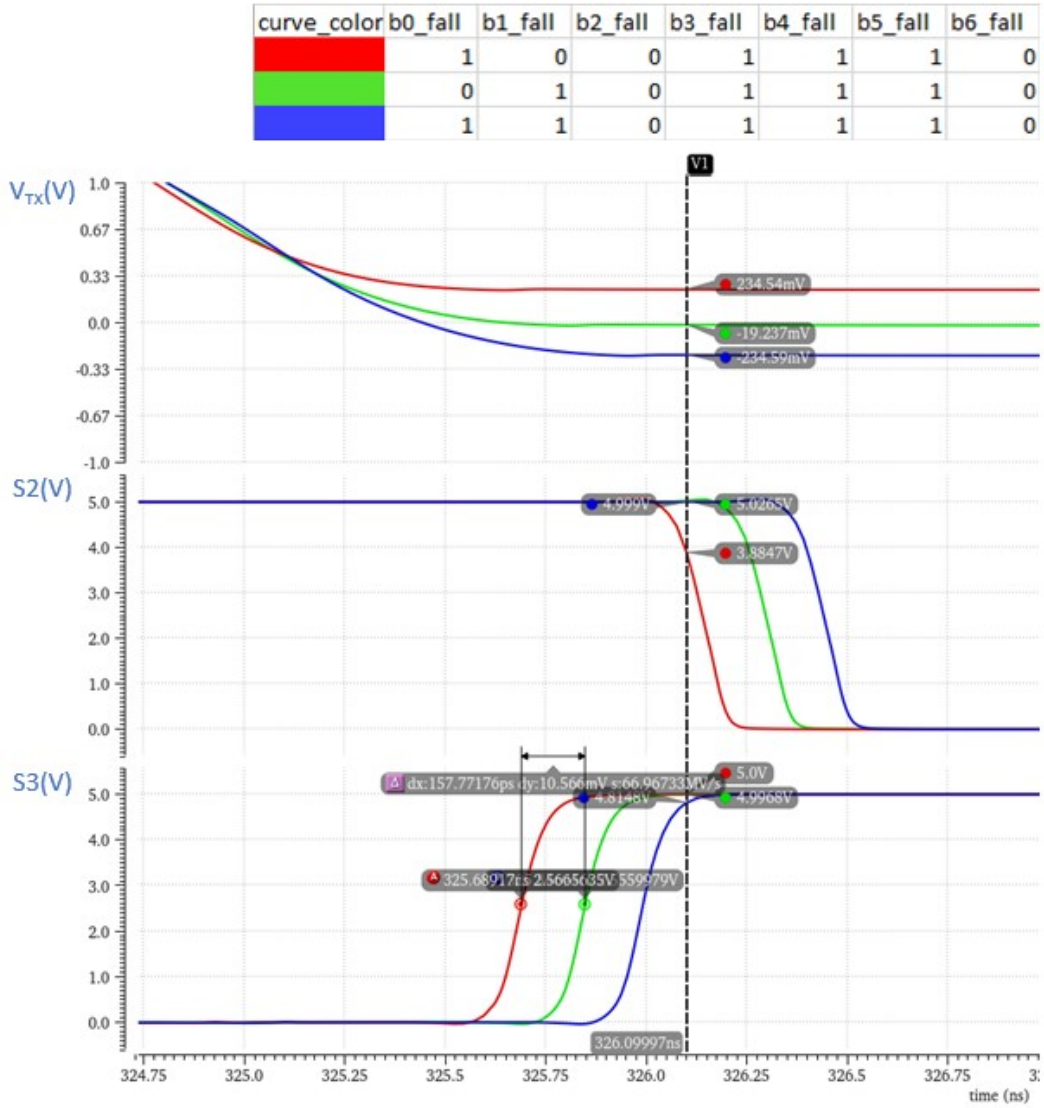

Figure 5.4: Adjust the TDAC_fall setting to shift S2, S3 then further modify the time of falling edge

Set the TDAC_Vc_delay to set the position of the rising edge of V_c, thereby setting the amplitude that V_B can reach, further setting the current level that I_L can reach at the end of falling edge. Adjust the settings by observing the recycling time that is needed for recycling phase. The step size of TDAC can reach 16 ps. Eventually, recycling time is trimmed to its highest possible value (green curve), see figure 5.5.

curve_color	b0_Vc_delay	b1_Vc_delay	b2_Vc_delay	b3_Vc_delay
	1	1	1	0
	0	0	0	1
	1	0	0	1

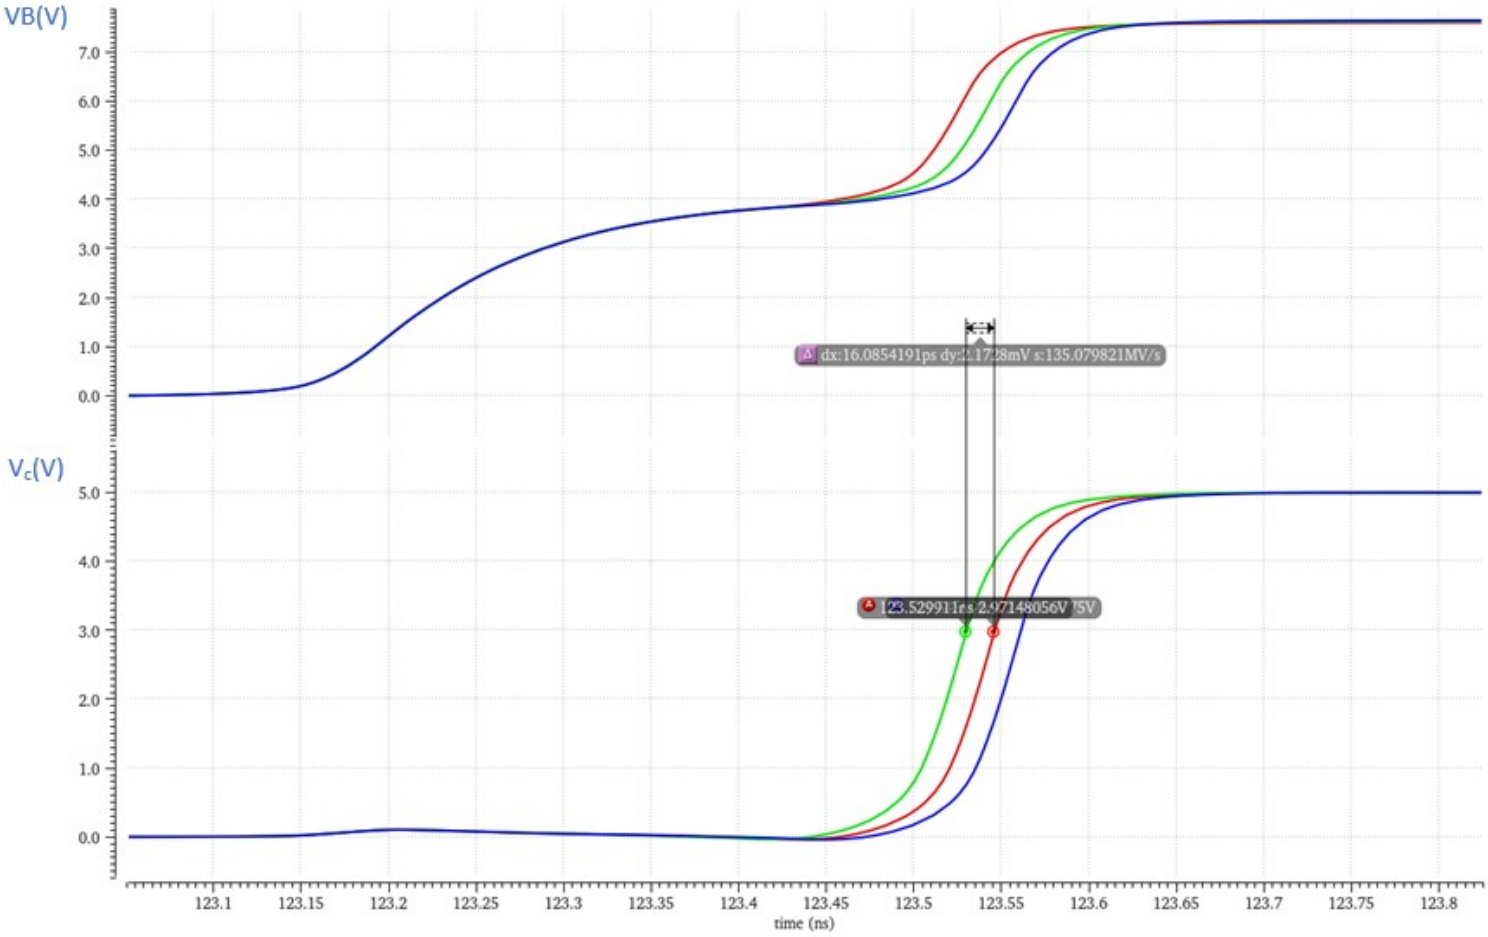

Figure 5.5: Adjust the TDAC_fall setting to shift the rising edge of V_c then further shift V_B

Set the TDAC_Vc_pulse to set the pulse width of V_c, thereby setting the current level that I_L can reach at the end of falling edge. Adjust the settings by observing the recycling time that is needed for recycling phase. The step size of TDAC can reach 16 ps. Eventually, recycling time is trimmed to its highest possible value (red curve), see figure 5.6.

curve_color	b0_Vc_pulse	b1_Vc_pulse	b2_Vc_pulse
	1	0	0
	0	1	0
	1	1	0

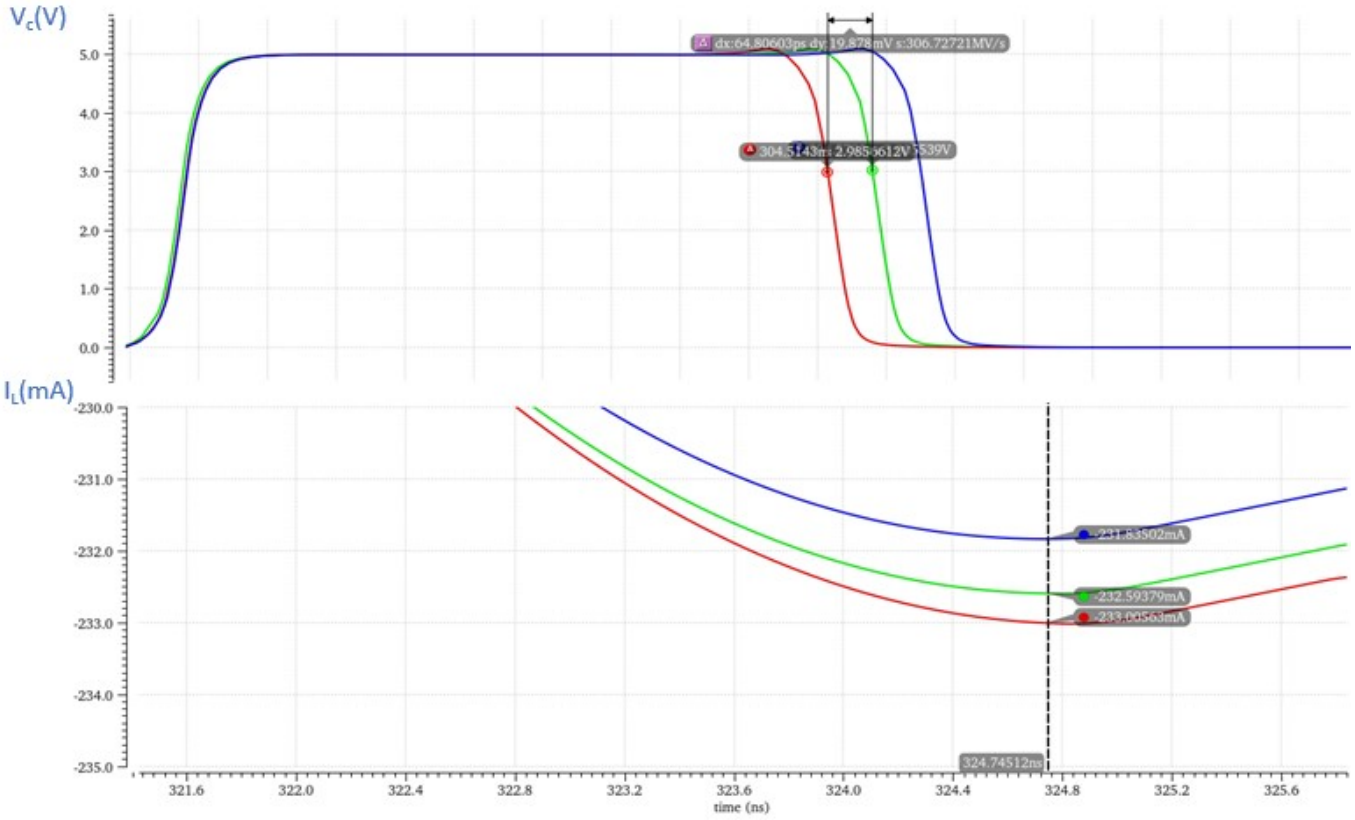

Figure 5.6: Adjust the TDAC_fall setting to modify the pulse width of V_c then further influence the energy efficiency

5.2 The falling edge trimming loop

During the trimming process for the falling edge, an off-chip algorithm determines whether the voltage crosses zero by observing if the continuous-time comparator's output transitions from logic "0" to "1". In the actual circuit, an additional offset is introduced to the comparator to trigger the comparison slightly before the zero-crossing point. This is because the comparator itself has propagation delay, and its output must pass through the level shifter and gate driver to control the core switches and advance the whole circuit to the recycling phase when the falling edge of V_{TX} approaches approximately -0.7V during any of the trimming cycles. The level shifter and gate driver together contribute an additional delay of about 1.5ns.

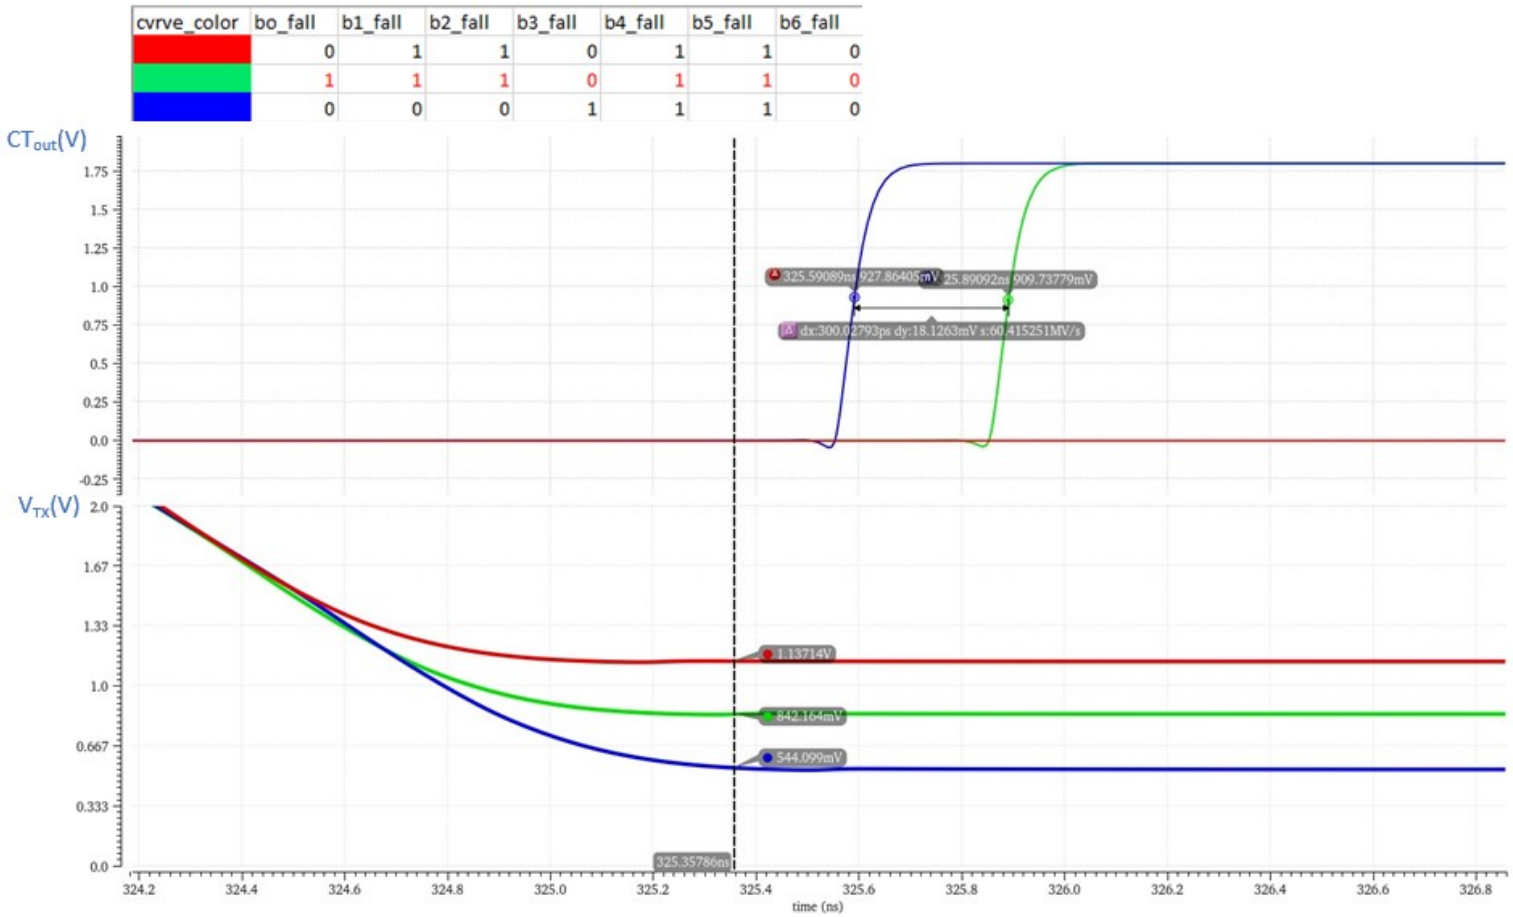

Figure 5.7: The continuous-time circuit observes the V_{TX} and sets the TDAC_fall

As shown in the figure 5.7, a linear search algorithm is hired to determine the TDAC_fall setting. In the case of the red curve, the comparator output remains at logic "0", indicating that the zero-crossing does not occur and that the TDAC setting is insufficient. For the green curve, the comparator output begins to transition from "0" to "1", indicating that the optimal TDAC setting has been identified. It is worth to note that the V_{TX} voltage does not

really fall to zero but rather stops at around 800mV, this is designed to compensate for that delays of the level shifter and gate driver.

If the TDAC setting exceeds the optimal, as shown by the blue curve, the comparator transitions from ”0” to ”1” even earlier (approximately 300ps earlier than the green curve). This occurs because, under the critical condition represented by the green curve, the differential input to the comparator is too small, resulting in a longer time to generate the comparison result. Eventually, the code corresponding to the green curve is identified as the setting for TDAC_fall.

5.3 Final pulse generation

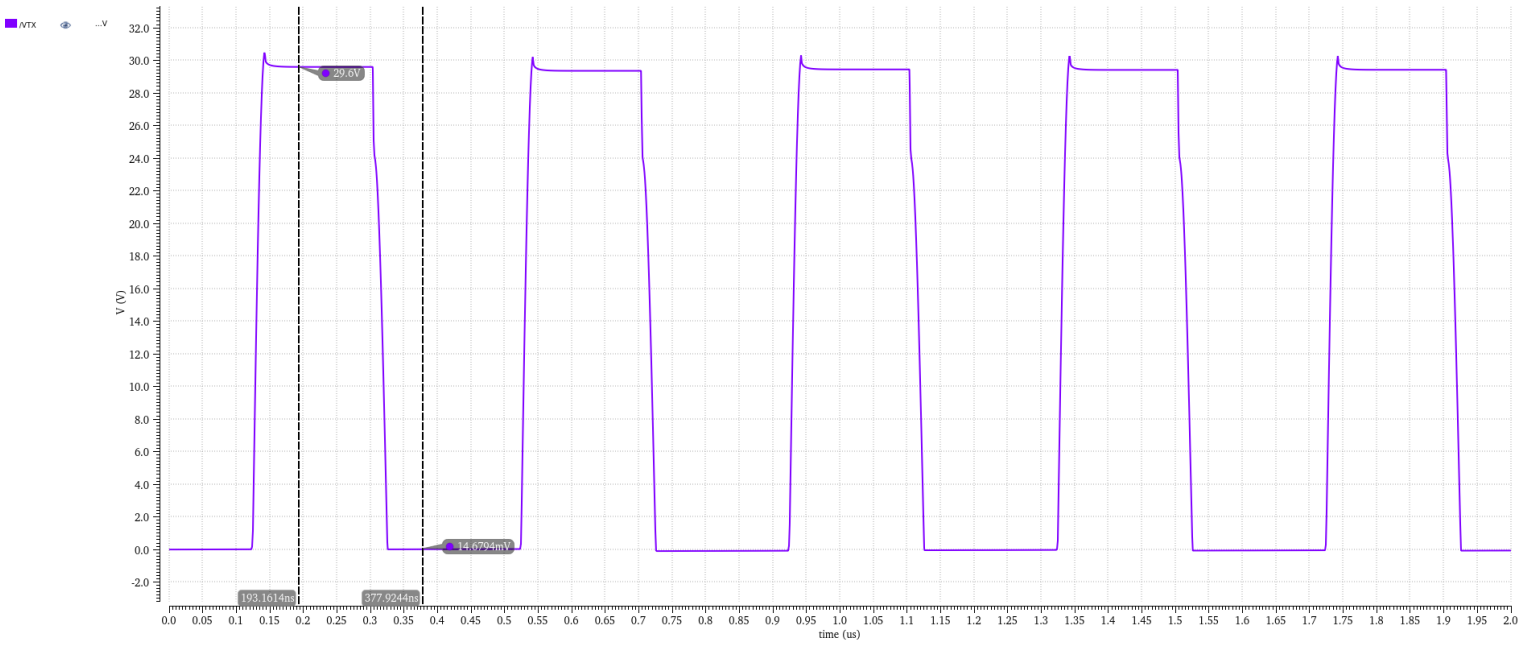

Figure 5.8: The final generated pulse train — after completing the calibration and trimming steps

After completing all the trimming and calibration cycles, the proposed circuit design is capable of generating a HV pulse train of 30V. As shown in the figure 5.8, the pulse amplitude during the ”hold phase” reaches 29.6V after calibration, while the transducer retains only 15mV of energy after the falling edge is completed.

6 Conclusions and future work

6.1 Conclusions

Based on the current research and simulation results, the proposed pulser has been proven feasible for directly converting LV battery supply into HV pulse train. This is primarily attributed to the design of the high-side HV switch, which has the capability to be turned off and back on under the HV conditions even in the absence of a HV supply. Simulation results indicate that the total energy consumption of all the switches can be optimized to below $0.4\%fCV^2$. However, this on the other hand requires the switches to have a very large size, which poses significant challenges in operating them. Since larger switches usually need to be turned on or off within a shorter time frame compare with smaller switches, introducing substantial timing challenges.

6.2 Discussion and future work

To proceed the current research direction, some blocks in the circuit still need to be further implemented at the transistor level to complete the design. Certain sensing blocks for calibration and trimming remain missing, such as the potential S/H circuitry mentioned in chapter 4, which provides an adjustable reference level for the comparator involved in trimming the rising edge.

In addition, Verilog code that captures signals from the sensing circuit and sets the DAC (potential SAR algorithms and linear search algorithms) also needs to be written to enable the circuit to perform calibration and trimming tasks fully automatically. Following this, layout and post-layout simulations need to be conducted to verify whether the circuit's feasibility meets expectations. Finally, the circuit design will need to be taped out and subsequent measuring after the chip is fabricated to evaluate its performance.

Taking a step back, current research highlights the sensitivity of switch operation to timing setting. It further reveals the potential complexity of highly sophisticated sensing blocks and the algorithms that work collaboratively with them. Given the potential impact of PVT variations and parasitic capacitance introduced by layout on circuit performance, it can be anticipated that the actual performance of this circuit might be relatively fragile.

Therefore, in addition to advancing the existing research, future efforts could also consider re-evaluating and redirecting certain aspects of the current approach. For instance, rethinking the transistor-level implementation of the highside HV switch M4 to relax the extremely stringent timing requirements associated with its associated nodes. Similarly, investigating

simpler implementations for the sensing circuits involved in trimming step of Vc_delay_time could prevent any other calibration steps from nesting within the process.

From a more higher level viewpoint, considering the complexity introduced by the current circuit design to generate HV pulses, it might be a wiser trajectory to explore reducing the amplitude of the pulses required to excite wearable devices. Alternatively, in applications like image processing, where only a decent SNR is required, it could be a viable option to abandon pulse train excitation in favor of traditional measurement/averaging schemes. On the other hand, for applications such as Doppler measurements, continuous-wave pulsing, ultrasonic power stimulation, or those aimed at beamforming as the ultimate goal, HV pulse trains are indispensable. In these cases, tolerating the complexity of the circuit design becomes an unavoidable trade-off.

Bibliography

[1] Kailiang Chen et al. "Ultrasonic imaging transceiver design for CMUT: A three-level 30-Vpp pulse-shaping pulser with improved efficiency and a noise-optimized receiver". In: *IEEE Journal of Solid-State Circuits* 48.11 (2013), pp. 2734–2745.

[2] Krzysztof Iniewski. *Medical imaging: principles, detectors, and electronics*. John Wiley & Sons, 2009.

[3] Yaohua Zhang and Andreas Demosthenous. "Integrated circuits for medical ultrasound applications: Imaging and beyond". In: *IEEE Transactions on Biomedical Circuits and Systems* 15.5 (2021), pp. 838–858.

[4] Imad Bellouki et al. "6.4 A Resonant High-Voltage Pulser for Battery-Powered Ultrasound Devices". In: *2024 IEEE International Solid-State Circuits Conference (ISSCC)*. Vol. 67. IEEE. 2024, pp. 106–108.

[5] Muyang Lin et al. "A fully integrated wearable ultrasound system to monitor deep tissues in moving subjects". In: *Nature Biotechnology* 42.3 (2024), pp. 448–457.

[6] Chao Chen and Michiel AP Pertijs. "Integrated transceivers for emerging medical ultrasound imaging devices: A review". In: *IEEE Open Journal of the Solid-State Circuits Society* 1 (2021), pp. 104–114.

[7] Jaesuk Choi et al. "Energy-efficient high-voltage pulsers for ultrasound transducers". In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 68.1 (2020), pp. 19–23.

[8] John D Larson et al. "Modified Butterworth-Van Dyke circuit for FBAR resonators and automated measurement system". In: *2000 IEEE ultrasonics symposium. proceedings. an international symposium (Cat. No. 00CH37121)*. Vol. 1. IEEE. 2000, pp. 863–868.

[9] Hanyun Zhou, SH Huang, and Wei Li. "Electrical impedance matching between piezoelectric transducer and power amplifier". In: *IEEE Sensors Journal* 20.23 (2020), pp. 14273–14281.

[10] Linas Svilainis, Andrius Chaziachmetovas, and Vytautas Dumbrava. "Efficient high voltage pulser for piezoelectric air coupled transducer". In: *Ultrasonics* 53.1 (2013), pp. 225–231.

[11] Dongning Zhao et al. "High-voltage pulser for ultrasound medical imaging applications". In: *2011 International Symposium on Integrated Circuits*. IEEE. 2011, pp. 408–411.

[12] Yidi Xiao, Hassan Rivandi, and Tiago L Costa. "An Energy-Efficient High-Voltage Pulser for High-Frequency Ultrasound Medical Applications". In: *2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE. 2023, pp. 1–5.

[13] Lars J Svensson and Jeff G Koller. "Driving a capacitive load without dissipating fCV/-sup 2". In: *Proceedings of 1994 IEEE symposium on low power electronics*. IEEE. 1994, pp. 100–101.

[14] Jonas Pelgrims, Kris Myny, and Wim Dehaene. "A 36V Ultrasonic Driver for Haptic Feedback Using Advanced Charge Recycling Achieving 0.20 CV 2 f Power Consumption". In: *ESSCIRC 2021-IEEE 47th European Solid State Circuits Conference (ESSCIRC)*. IEEE. 2021, pp. 159–162.

[15] Jihee Lee et al. "A 36-channel auto-calibrated front-end ASIC for a pMUT-based miniaturized 3-D ultrasound system". In: *IEEE Journal of Solid-State Circuits* 56.6 (2021), pp. 1910–1923.

[16] Richard J Przybyla et al. "3D ultrasonic rangefinder on a chip". In: *IEEE Journal of Solid-State Circuits* 50.1 (2014), pp. 320–334.

[17] Saifullah Amir, Ronan van der Zee, and Bram Nauta. "A self-oscillating boosting amplifier with adaptive soft switching control for piezoelectric transducers". In: *IEEE Journal of Solid-State Circuits* 54.1 (2018), pp. 253–265.

[18] Simon Chaput, David Brooks, and Gu-Yeon Wei. "21.5 A 3-to-5V input 100V pp output 57.7 mW 0.42% THD+ N highly integrated piezoelectric actuator driver". In: *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE. 2017, pp. 360–361.

[19] Eric J Carlson, Kai Strunz, and Brian P Otis. "A 20 mV input boost converter with efficient digital control for thermoelectric energy harvesting". In: *IEEE Journal of Solid-State Circuits* 45.4 (2010), pp. 741–750.

[20] Mingliang Tan et al. "A 64-channel transmit beamformer with±30-V bipolar high-voltage pulsers for catheter-based ultrasound probes". In: *IEEE Journal of Solid-State Circuits* 55.7 (2020), pp. 1796–1806.