StimDust: Miniaturized Ultrasonic Peripheral Nerve Stimulator

Ka Yiu Li

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-16
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-16.html

May 1, 2019

StimDust: Miniaturized Ultrasonic Peripheral Nerve Stimulator

by Ka Yiu Li

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Professor Rikky Muller
Research Advisor

11 May 2018

(Date)

****

Professor Michel Maharbiz
Second Reader

11 May 2018

(Date)

Acknowledgments

First, I would like to thank my advisor, Professor Rikky Muller, for the opportunity to work on this project and her insightful guidance. I would also like to thank Professor Benjamin Johnson for his crucial contributions to StimDust and the help during the tape-out. Finally, I would like to thank my family for supporting me throughout my undergraduate and graduate career.

Abstract

Neuromodulation has been proven successful as an alternative to patients who failed to respond to pharmaceutical treatments. However, commercially available peripheral nerve stimulators typically require batteries and wires, limiting the lifespan of the implants and possibly causing complications. Radio frequency (RF) powered neurostimulators have been demonstrated in the past, yet, they are unable to scale below mm scale, thus they are also difficult to implant into the body.

This technical report presents the design of a wirelessly powered and communicated chip for peripheral nerve stimulator using ultrasound instead of RF. Ultrasound has several advantages over RF. The power transmission limit and the path loss are both superior for ultrasound, allowing ultrasonic neurostimulators to scale below mm-scale. In this chip, Ultrasound is used to power, control, and monitor the stimulator while maintaining a small form factor of the chip. The objective of this project is to explore the lower limit of the area of the chip and make improvements over the previous generation.

The chip is fabricated in 65nm process, occupying $500\mu m \times 200\mu m = 0.1mm^2$ that is 1/10th of the previous generation.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

The history of peripheral nerve stimulation can be traced back as early as 1962 [1]. Since then, it has been developed as an alternative treatment for diseases such as neurogenic detrusor overactivity [2], chronic migraine [3], shoulder pain [4], and more. There are also positive outcomes for using peripheral nerve stimulation to restore sensory feedback for individuals using prostheses [5]. The recent advancements in recording techniques, such as mm-scale wireless recording device [6] and optogenetics, have greatly increased the recording resolution, potentially enabling scientists to create an accurate mapping of the nerve fibers and decode neural signals. Thus, a wide range of clinical and research applications of peripheral nerve stimulator is expected.

Despite a wide range of applications and a long history of deploying peripheral nerve simulators, such bioelectronic medicine has not been widely used. Looking at commercially available peripheral nerve stimulators, most of them fall into at least one of the following categories: (i) powered by a battery (ii) bulky size (iii) require wire connections. The drawback of (i) is that it limits the lifespan of the simulators, requiring surgical replacement for long-term usage, while (ii), (iii) increase the chance of infection and make the surgery lengthy and challenging [7]. Thus, a small, batteryless, and wirelessly controlled peripheral nerve stimulator is needed for the sake of safety, durability, and ease in implanting.

1.2 Ultrasound Versus Electromagnetic Wave

To achieve wirelessly powered and controlled implants, one option is to use electromagnetic (EM) wave. EM is widely used for wireless communication and energy transmission. However, because of health concerns, the transmission power of EM wave is limited by regulations and standards. For example, IEEE recommends an upper limit of $0.1mW/mm^2$ at a frequency of 3GHz [8]. Moreover, the attenuation of EM wave by body tissue, and the low efficiency of power transfer with a small footprint make EM-based implants very difficult to scale [9].

Alternatively, we can use ultrasound. Ultrasound has been widely used for medical imaging and is considered to be safe. It has several advantages over EM. First, the path loss of ultrasound in tissue is much lower [9]. Second, the permissible transmission power of ultrasound recommended by FDA is $0.94mW/mm^2$ for most applications, and this limit increases to $7.2mW/mm^2$ if the device follows the Output Display Standard [10]. Assuming the stimulator is used in a closed-loop without an output display, the transmission power will be limited by $0.94mW/mm^2$. This limit is still 9.4 times higher than that of the EM. Because of these reasons, ultrasonic based peripheral nerve stimulation can be scaled down well below mm scale while still harvesting enough power.

1.3 Previous Generation

To demonstrate the possibility of an ultrasonic-based peripheral nerve stimulator, the first StimDust was fabricated, where I involved in designing the layout and verifying the chip. The stimulator consists of an integrated circuit (IC) for power harvesting, communication, and control; a piezocrystal to convert ultrasound pressure to voltage; a 0201 capacitor for power storage. The system was tested *in-vivo*, showing its ability to induce motor unit action potential. Figure 1.1 shows the StimDust mote and the in vivo setup. The IC occupies an area of $1mm^2$ and the stimulator has a total volume of $6.5mm^2$ with 82% peak efficiency [11]. Figure 1.2 and Figure 1.3 shows the StimDust mote and the die of the IC respectively.

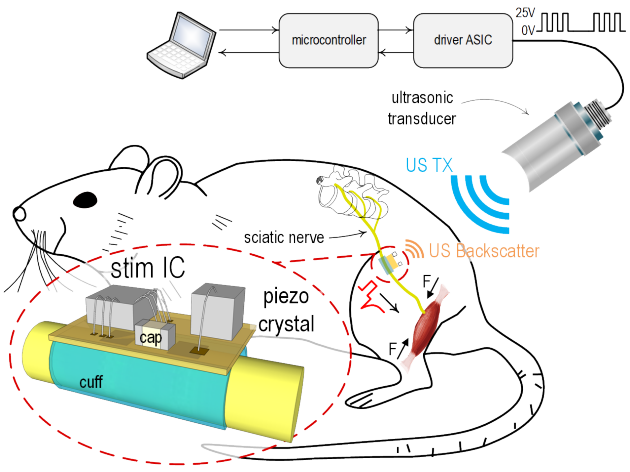

Figure 1.1: The system of the StimDust. The StimDust mote is cuffed to the sciatic nerve. External ultrasonic transducer is used to control and monitor the stimulation.

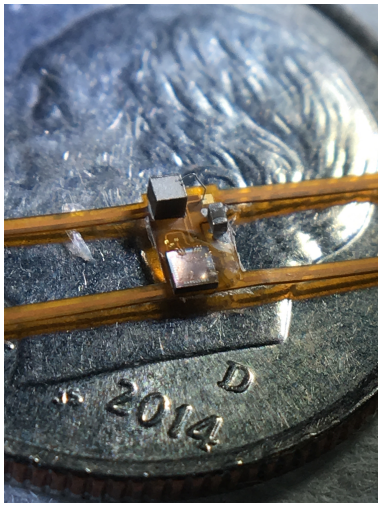

Figure 1.2: The photo of the StimDust mote on top of a dime. The three components on the mote from top to bottom are piezocrystal, 0201 capacitor, and the IC. [11]

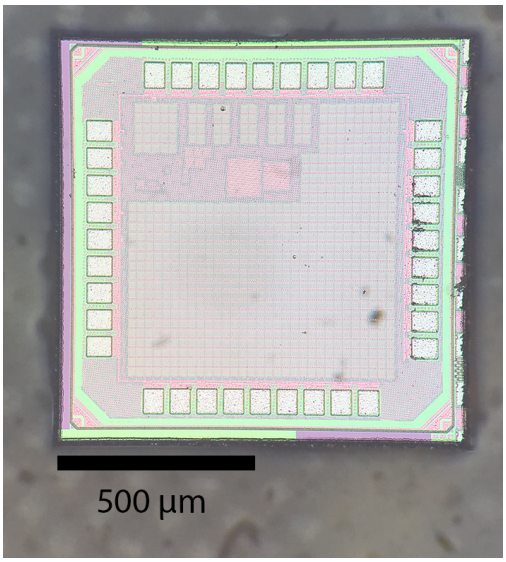

Figure 1.3: The IC under a microscope. [11]

Although the chip has been working as expected, there are few improvements we want to make. First, the active area is only a fraction of the size of the IC. As seen from Figure 1.2, the chip occupies the most area amongst the three components on the mote. If we can reduce the size of the IC, it will further reduce the volume of the stimulator. Second, there is no safety feature to stop the stimulation when, for reasons such as noise, coupling, piezo ringing, state-flip, etc, the IC considers "stim" as "recharging" or verse versa. In general, the recharging time is much longer than the stimulating time. Thus, in the event of misinterpretation, the amount of charge from the stimulator will potentially damage the nerve. This is especially important for a device that is aiming for medical use. Although we can monitor the status of the StimDust by looking at the amplitude of the reflected wave, we strongly believe that an additional safety feature is needed such that when one safety feature fails, the other will kick in and correct the problem. Third, current peaking was observed at the beginning of the stimulation, as shown in Figure 1.4. This limits the accuracy of the total output charge. For applications where short pulses are desired, the extra charge due to current peaking will be significant. Therefore, the goal of this new StimDust is to minimize the area while adding safety feature and improving current accuracy.

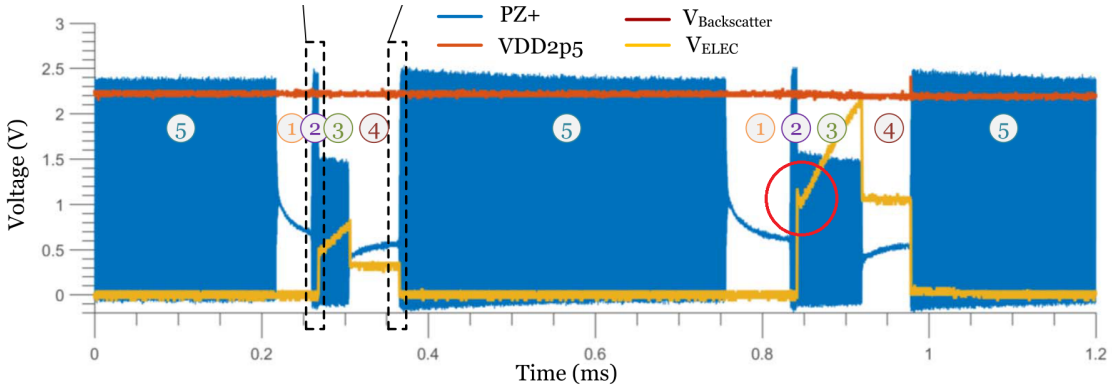

Figure 1.4: The output waveform of the previous generation StimDust [11]. The current peaking is circled in red. The amount of peaking is worst when the output current is at the maximum level.

Chapter 2

Powering and Communication Protocol

The protocol of the IC is inherited from the previous generation, with a few changes, which will be highlighted in this chapter and discussed in detail in Chapter 3.8. The protocol contains six phases, as in the previous generation. They are shown in Figure 2.1. Compared with the previous generation, phase 2 now requires one additional clock and the chip can transit from phase 3 to phase 5 directly. We will explain each phase in their numerical order, and these changed will be highlighted again when we explain phase 2 and phase 3 of this protocol.

During phase 0, the chip harvests the energy from the incoming ultrasound and stores it into a capacitor. When the energy across this capacitor is sufficient, the circuits will be biased at its designed operating point. Since most of the digital logic in this process are reset by active low, the POR (Power-On Reset) signal is initially low to reset the state of the chip. When the VDD is high enough, the POR goes high to release the reset. The chip is ready for stimulation when POR goes high.

After phase 0, stimulation command (stim command) can be sent to start the stimulation. The stim command consists of phase 1-5, as shown in Figure 2.1. The stim command starts with an Ultrasound-Free Interval (UFI), phase 1, where no ultrasound is being sent to the chip, to set the current amplitude of the stimulator. To properly determine whether there is incoming sound, the watchdog signal is designed to go high after $\approx 3\mu s$ of Ultrasound-Free

Interval (UFI). When the watchdog signal goes high to acknowledge this Ultrasound-Free Interval, it will start the time-to-digital converter(TDC). For every $\approx 10\mu s$, the output of the TDC increases by one and the current amplitude will increase by $50\mu A$, up to $400\mu A$. When ultrasound resumes, the chip immediately transits to phase 2.

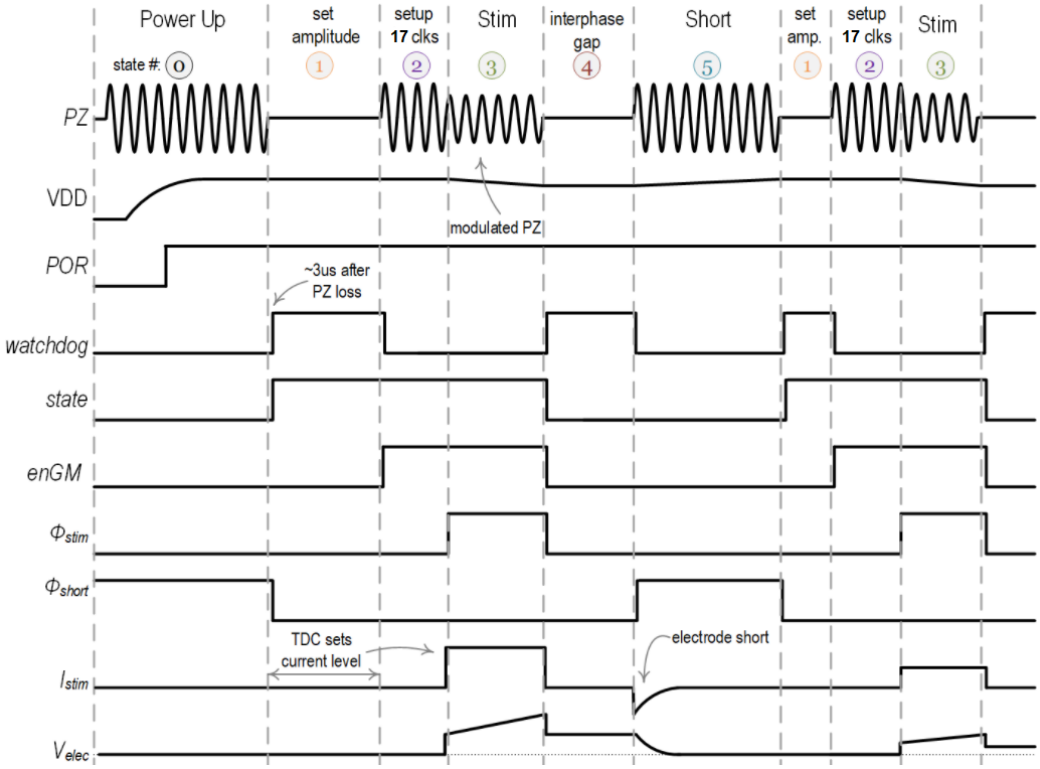

Figure 2.1: Timing diagram of the protocol

The duration of phase 2 is precisely 17 cycles of the incoming ultrasound instead of 16 cycles in the previous generation. As in the previous generation, these cycles allow the biasing and current reference for the DAC to settle. However, by increasing this by one, we can transit to phase 3 by looking at only the MSB of the counter. This frees up the 4 LSBs to be used for the safety timer during stimulation to prevent prolong simulation. Once the counter counts 17 cycles, the chip transits to phase 3 during which the output of the DAC is enabled and the reflected wave is modulated for external monitoring. The details of increasing the number of cycles of phase 2 will be discussed in Chapter 3.8.

Depending on the duration of phase 3, the chip will transit to different states instead

of always transiting to phase 4 in the previous generation. During normal operation, the duration of phase 3 should be less than what the safety timer allows, which is 1.1ms. In this case, the chip simply transits to phase 4. However, if the duration of phase 3 is longer than 1.1ms, the safety timer will be triggered and the chip will immediately transit to phase 5. To illustrate the effect of such transition, consider we send two pulses, with different phase 3 duration, as shown in Figure 2.2. When the safety timer is triggered, the chip will immediately short the output to ground and treat the next pulse as stimulation. This essentially creates a flip in the state of the chip. Typically, the length of phase 5 for recharging the mote is longer than 1.1ms. Thus, the state flip ensures that the chip interprets the short pulse as stim (phase 2 and 3) and the long pulse as recharge (phase 5), but not the other way around.

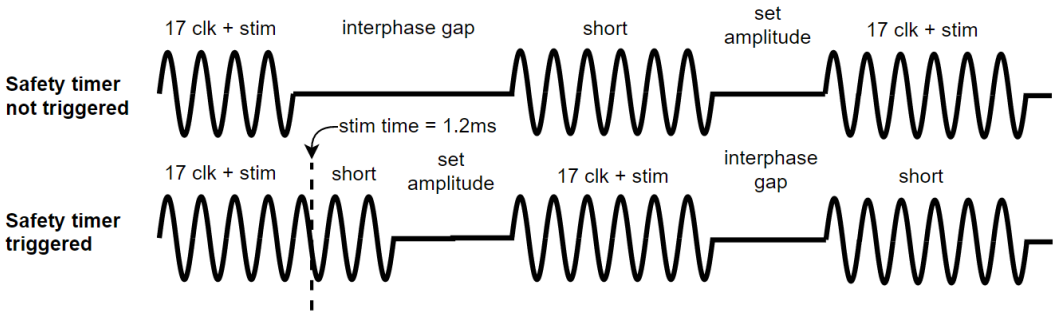

Figure 2.2: Timing diagram to demonstrate the states after triggering safety timer.

Assuming the safety timer is not being triggered, the chip transits to phase 4 $3\mu s$ after the ultrasound is lost. This creates the necessary interphase gap for energy efficient stimulation. When the next ultrasound pulse arrives at the chip, the chip transit to phase 5 and two things happen: (i) the output node is shorted to ground to balance the charge. This creates a charge balanced biphasic waveform with delayed fast reversal as shown as I_{stim} in Figure 2.1 (ii) the chip is being recharged. After phase 5, the chip is ready to start another stimulation.

The complete finite state machine of the chip is shown in Figure 2.3. The changes are highlighted in blue in the figure. Figure 2.4 also shows the finite state machine of the previous chip for comparison. In general, the state of the FSM should transit clockwise in the figure.

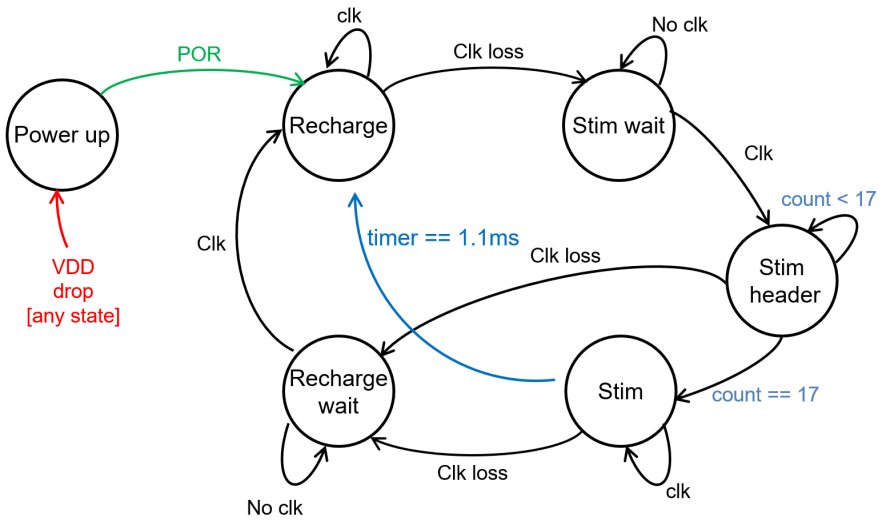

Figure 2.3: Finite state machine of this chip. The changes are highlighted in blue in the figure.

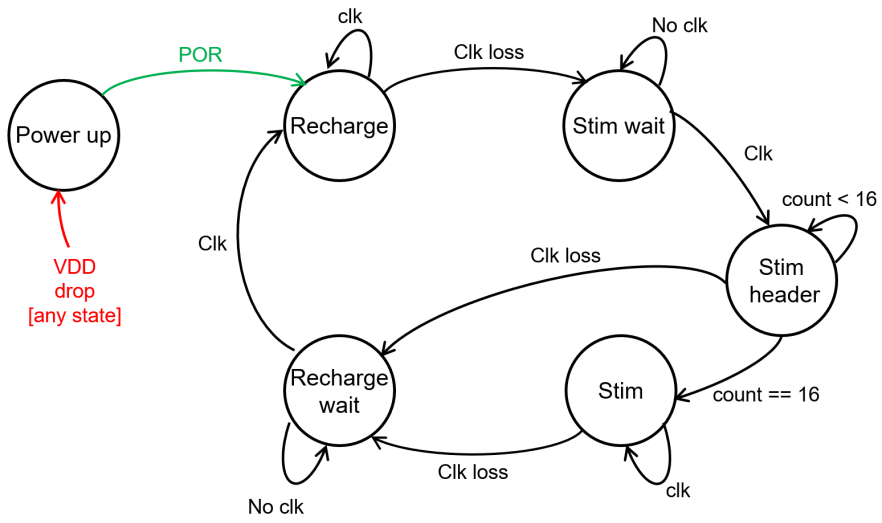

Figure 2.4: Finite state machine of the previous chip.

Chapter 3

Circuit Architecture

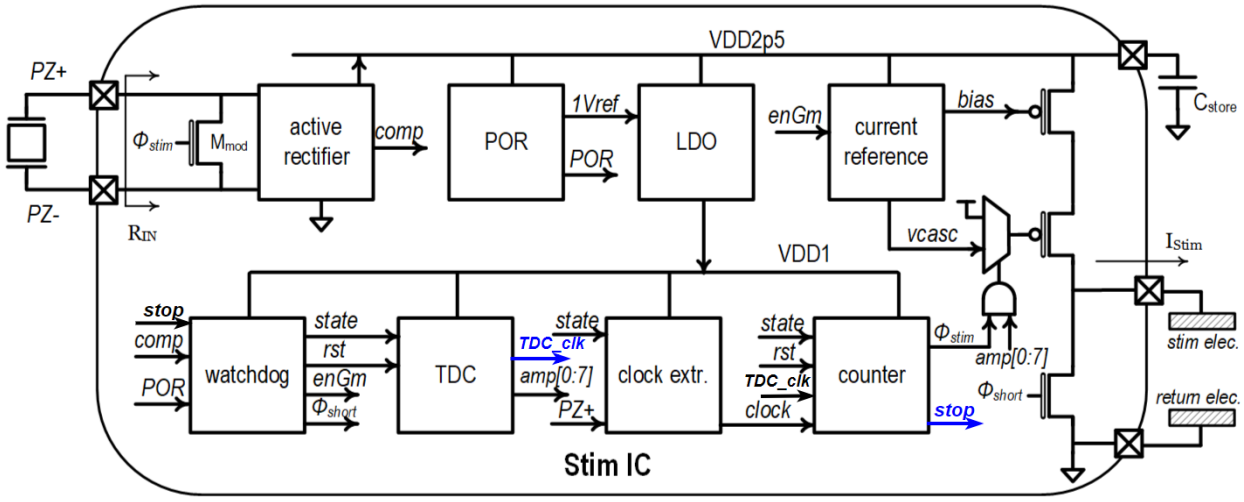

Figure 3.1: Circuit block diagram of the chip. The signals in blue are new for this chip to support the safety timer.

Since the protocol of this chip is similar to its predecessor, most of the circuit blocks that are designed by Dr. Benjamin C. Johnson in [11] are reused. The blocked being reused were simulated again to verify their performance and layouts were adjusted to fit into the smaller area. The counter and TDC were redesigned to support safety timer and state flip as mentioned in Chapter 2. The current starved buffers in TDC are also resized to limit their leakage. Since the names of some of the circuit blocks and signals are similar, the names of the signals are italicized to avoid confusion in the following discussion. The discussion below

will begin with some overview of the system and each circuit block. Then in Chapter 3.8 and 3.9 we will discuss the design to accomplish the goals in Chapter 1.

The StimDust consists of an active rectifier, POR, LDO, current reference, watchdog, TDC, clock extract, counter, current DAC, and a modulation transistor. Figure 3.1 shows the circuit block diagram of the design. The safety timer is embedded into the TDC and counter by using the TDC's output, *TDC_clk*, as an input to the counter. When enough *TDC_clk* is seen by the counter, *stop* will be low briefly to reset the *state* in watchdog. Typically, the active rectifier, POR, LDO, and watchdog are active all the time. The TDC is active only during set amplitude (phase 1) and stim (phase 3); the current reference, clock extract, counter, and the current DAC are on during stim (phase 3). When there is an incoming ultrasound, the active rectifier harvests the energy and stores it into an external capacitor that is connected across *VDD2p5* and *VSS*. The voltage across this capacitor will continue to increase until it reaches slightly below the peak level of the input signal. Once the voltage of *VDD2p5* reaches 1.7V, which is the minimum designed operating voltage of the StimDust, the POR releases the reset signal of all the digital circuities, by pulling *POR* to high (digital logics are reset by low signal in this process). Beyond this voltage, the circuits enter into its normal operation and are ready to receive stimulation command.

As described in Chapter 2, the stim command is composed of set amplitude (phase 1), setup (phase 2), stim (phase 3), interphase gap (phase 4) and short (phase 5). The role of the watchdog is to determine whether there is an incoming ultrasound and pull the *watchdog* to high when there is no incoming ultrasound. This *watchdog* is used to clock the *state* into the correct value which then turns on the TDC during set amplitude (phase 1) to convert the Ultrasound-Free Interval to digital code. After phase 1, the clock extract will extract the incoming ultrasound as *clock* and send it to the counter. At the same time, the current reference will be turned on. The counter will count 17 clocks from clock extract, after which it will generate ϕ_{stim} to start the stimulation. Thereafter, the interphase gap (phase 4) toggles the *watchdog* to high again to flip the *state*. This disables the current reference and current output of the DAC. At last, the chips enter into short (phase 5), which shorts the output to ground and recharges the external capacitor.

To conserve power, the system has two voltage domains. The *VDD2p5* is designed to stay at 2.5V for analog circuits where higher headroom is needed; the *VDD1* is designed to

stay at 1V for digital circuits. The 1V voltage supply is generated using LDO and *1Vref* (1V voltage reference) from the POR.

The output of the StimDust is a 3-bit thermometer current DAC. Each of the unit current source provides $50\mu A$ and is in cascode configuration. The biasing and sizing are optimized to preserve the headroom at the output.

During stimulation, the modulation transistor M_{mod} is being turned on. This changes the loading seen by the piezocrystal, thus changing the reflective coefficient. This creates a 1-bit modulation of the reflected ultrasound. By looking at the amplitude of the reflected ultrasound, we can monitor the status of the mote.

3.1 Active Rectifier

To convert AC power to DC power, an active bridge rectifier is used. In lots of application, passive full wave rectifier is used for this purpose. However, the diode drops of a PN junction diode limits the maximum voltage at the output, thus limiting the power conversion efficiency. To improve the efficiency of rectification, a full wave active rectifier is used in this chip.

The principle of an active rectifier is to actively turn on the transistors to minimize the voltage drops across the transistors. As shown in Figure 3.2, when the voltage of *PZ+* is higher than *PZ-*, M1 and M4 are being turned on, while M2 and M3 remain off. When the voltage of *PZ-* is higher than *PZ+*, as shown in Figure 3.3, M2 and M3 are being turned on, while M1 and M4 are kept off. In both cases, the current flows from *VDD2p5* to *VSS*, thus the external capacitor is being charged.

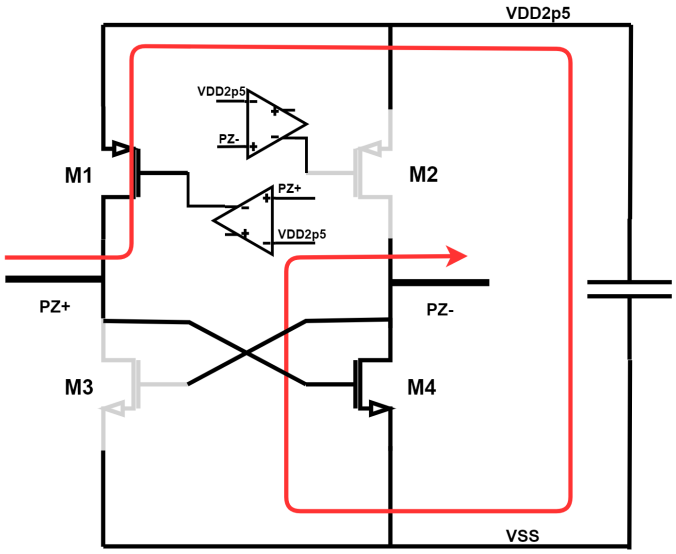

Figure 3.2: Signal path of the active rectifier when $PZ+ > PZ-$. The transistors in gray are off. The transistors in black are on.

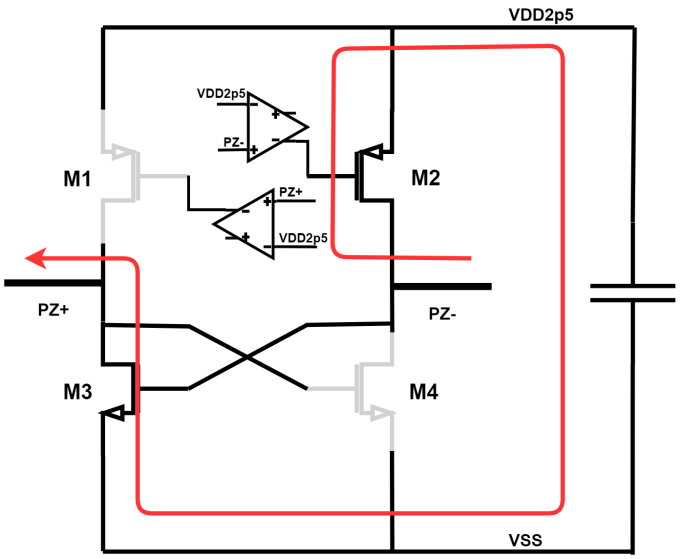

Figure 3.3: Signal path of the active rectifier when $PZ- > PZ+$. The transistors in gray are off. The transistors in black are on.

To properly drive the gates of the transistors, the gate of M4 in Figure 3.2 and 3.3 is connected to $PZ+$ and the gate of M3 is connected to PZ-. The gates of M1 and M2 are

each controlled by an amplifier-based comparator. The comparator is designed such that the non-inverting output is low when the voltage of the inverting input is above *VDD2p5*. Referring to Figure 3.2, when *PZ+* has a higher voltage than *VDD2p5*, the output of the comparator is low. As a result, M1 and M4 are being turned on.

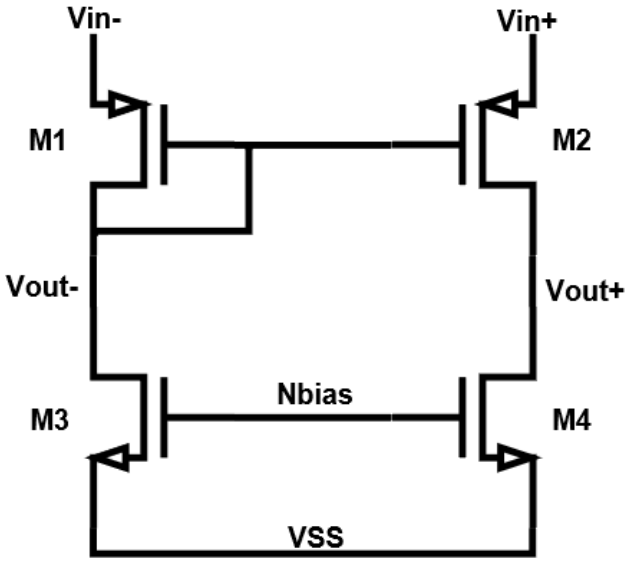

Figure 3.4: Schematic of the amplifier-based comparator controlling the gates of the PMOS of the rectifier.

Figure 3.4 shows the amplifier-based comparator. In this figure, the inverting input V_{in}^{-} is connected to the source of the diode connected PMOS, M1. Thus, the gate voltage of M1 and M2 is directly controlled by the voltage of V_{in}^{-}, and being kept around $V_{in}^{-} - V_{Th}$. When $V_{in}^{-} > V_{in}^{+}$, the $|V_{GS}|$ of M2 drops below $|V_{Th}|$, thus being turned off. In this case, $V_{out}^{-} \approx V_{in}^{-} - V_{Th}$ and $V_{out}^{+} \approx VSS = 0V$. As V_{in}^{-} continue to decrease, V_{out}^{-} continue to decrease and V_{out}^{+} continue to increase. Figure 3.5 shows the V_{out}^{+} and V_{out}^{-} at different V_{in}^{-}.

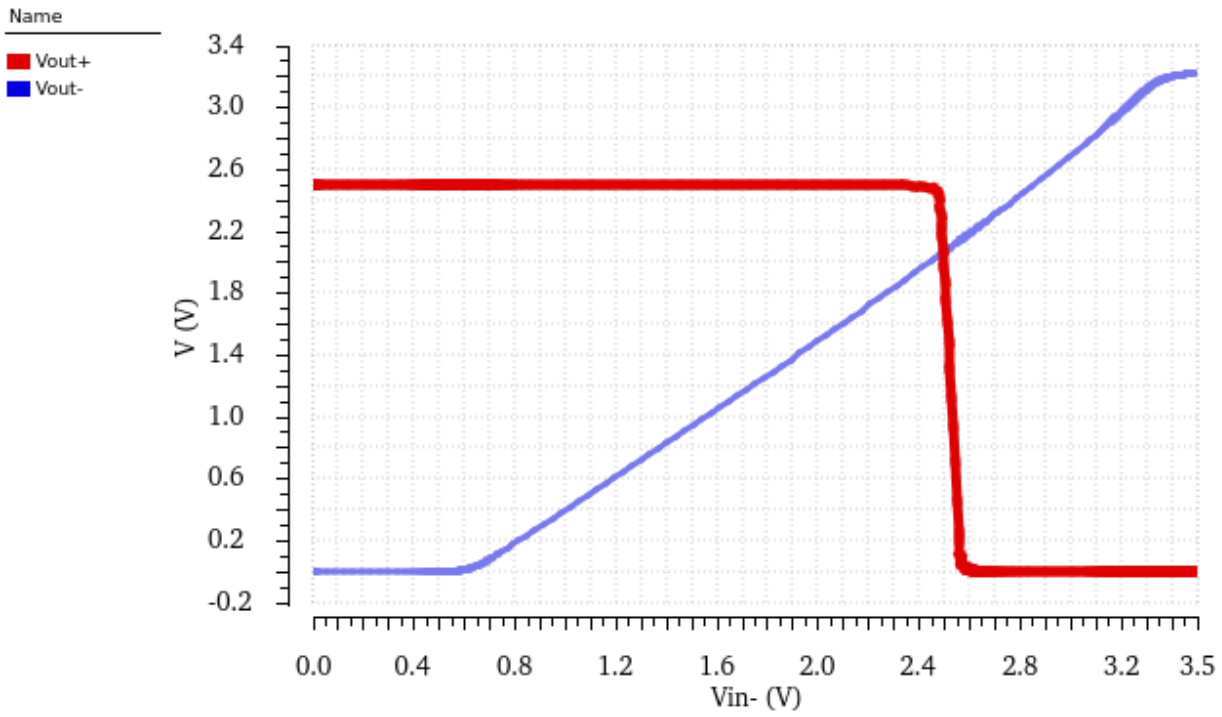

Figure 3.5: V_{out}^{+} and V_{out}^{-} vs V_{in}^{-} at $V_{in}^{+} = 2.5V$

3.2 POR

The functionality of POR is to reset the circuits and the state machine of the system when *VDD2p5* is below its minimum operating voltage. It also provides a 1V reference to the LDO to generate the 1V supply voltage for digital circuits.

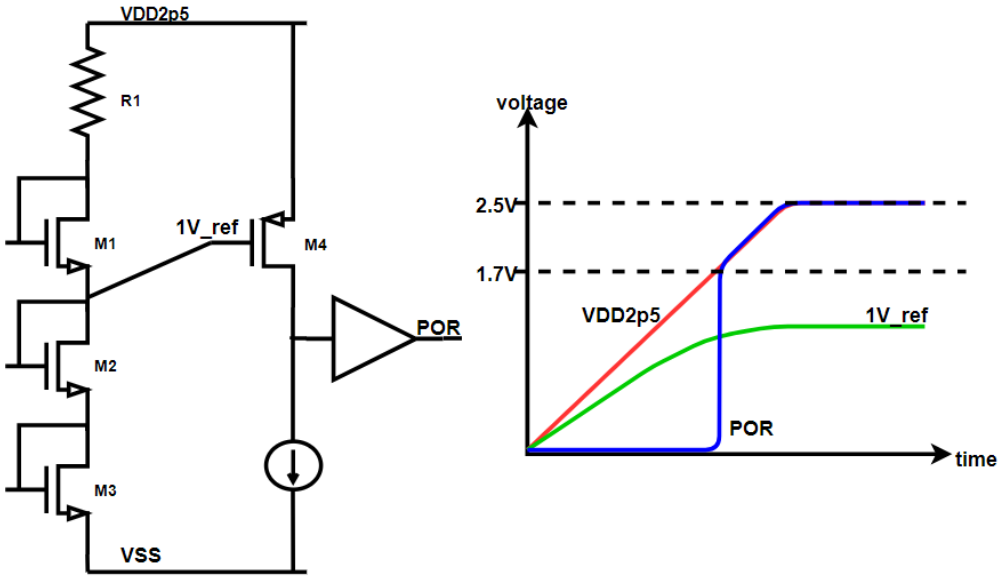

Figure 3.6: Circuit for 1V_ref and POR, and their corresponding voltages as VDD2p5 slews up

As shown in Figure 3.6, the *1V_ref* is provided by three diode-connected transistors, M1-M3, and R1. M1-M3 are sized such that they are operating in the subthreshold region when conducting 100nA of current. In subthreshold region, current is proportional to $e^{\frac{V_{GS}}{nV_T}}$, where V_T is the thermal voltage and n is some constant. With this exponential relationship, the V_{GS} of each diode is relatively stable even when the *VDD2p5* (and thus the current) fluctuates. The purpose of R1 is then to set the current of M1-M3. Since each diode connected transistors drops about 500mV, the resistance value of R1 can be determined by $\frac{2.5V-1.5V}{100nA} = 10M\Omega$.

The *POR* is initially pulled down by a current source. However, as seen from the plot in Figure 3.6, *1V_ref* rises much slower than VDD2p5. As VDD2p5 slews up, $|V_{GS4}|$ increases. At VDD2p5=1.7V, M4 is strong enough to pull POR to high. Figure 3.7 shows the simulation result of the circuit.

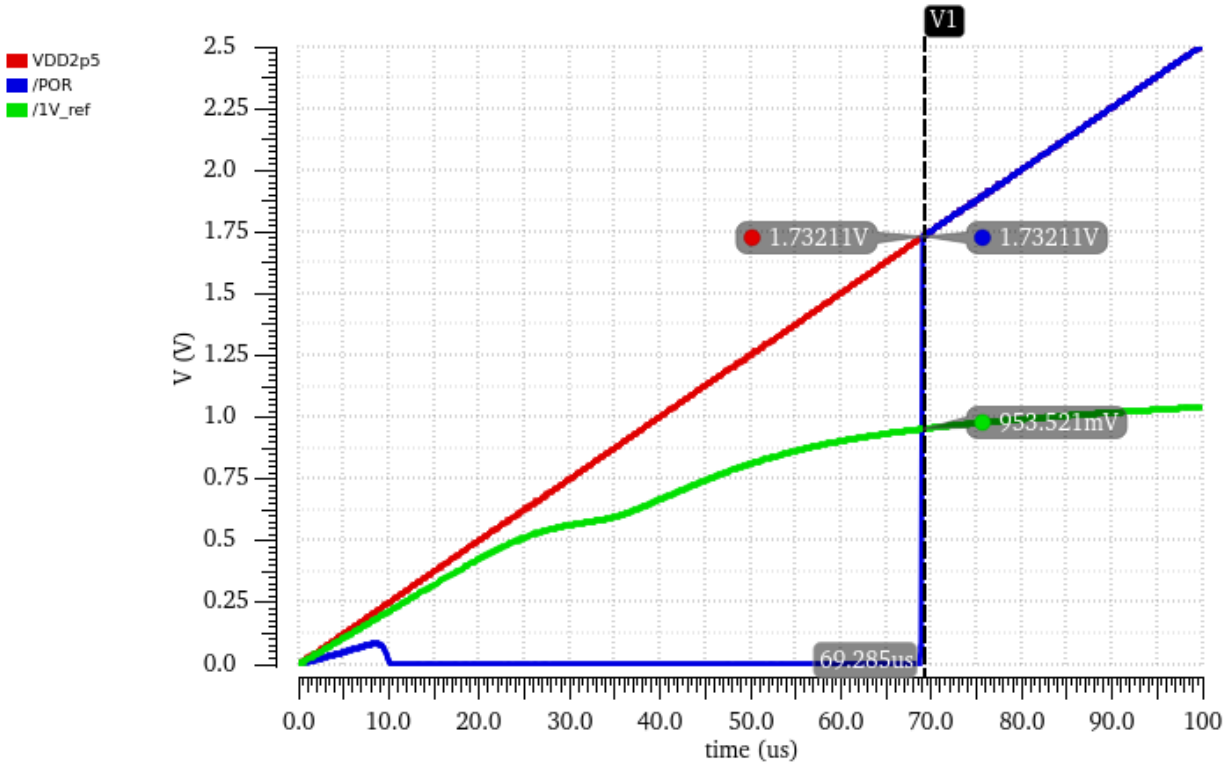

Figure 3.7: Simulation result of POR. The POR is high and 1V_ref=953mV when VDD2p5=1.73V

3.3 LDO

The LDO buffers the *1V_ref* to supply the current for the digital circuits. It consists of a differential to single-ended amplifier and a source follower in a feedback configuration. The source follower provides a low output impedance to drive a capacitive load. Figure 3.8 shows the schematic of the LDO.

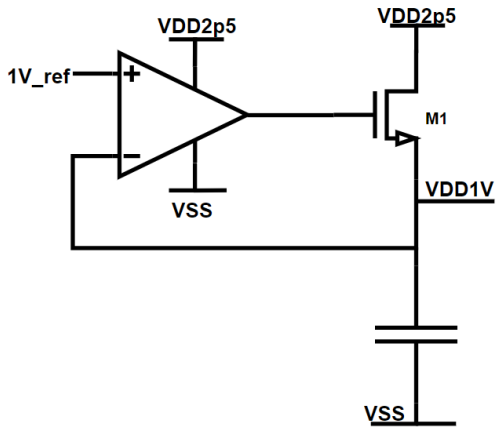

Figure 3.8: Schematic of the LDO

3.4 Current Reference

In modern analog circuits, bandgap reference is usually used to generate a temperature stable reference. However, since this peripheral nerve stimulator is for warm-blooded (endothermic) animals, the requirement for temperature constant is slack. Moreover, since the supply voltage of the StimDust is dependent on the input and output power, the current reference should be independent of the supply voltage. Thus, the main requirement of the current reference is to have good rejection of the supply voltage variation, with moderate temperature constant. For this reason, self-biasing current reference is used. Compared with bandgap reference, the self-biasing circuit requires less number of transistors and resistors, thus saving the area used.

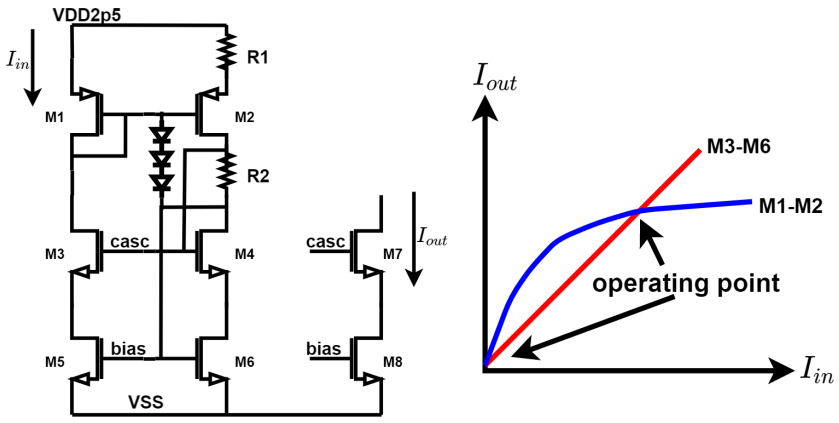

Figure 3.9: Schematic of the current reference circuits, and its input and output relation

Figure 3.9 shows the schematic of the circuit. The current reference is composed of Widlar current source (M1-M2, and R1), cascode current mirror (M3-M6), and start-up circuit (three diodes). The blue curve and red curve in Figure 3.9 shows the input and output relation of the Widlar current source and the cascode current mirror. Since these two current mirrors are connected, they must operate at the same point of the curves. Thus, there are two possible operating points of the circuit as shown in the figure. To avoid the current reference to be stuck at $I_{in} = I_{out} = 0$, three diodes are used as a start-up circuit to turn on M3-M6. To bias the current running through M1-M6, R1 and R2 are designed according to the following relationship [12],

$$\sqrt{I_{in}} = \frac{-\sqrt{\frac{2}{k'(W/L)_2}} + \sqrt{\frac{2}{k'(W/L)_2} + 4R_1 V_{ov1}}}{2R_1}$$

$$V_{DS6} = I_{in} R_2$$

During startup, the voltage at the gate of M1-M2 starts to rise. As this voltage rises above the three diode drops, the gate of M3-M6 are being charged, thus turning on M3-M6 to avoid the zero current operating point.

3.5 Watchdog

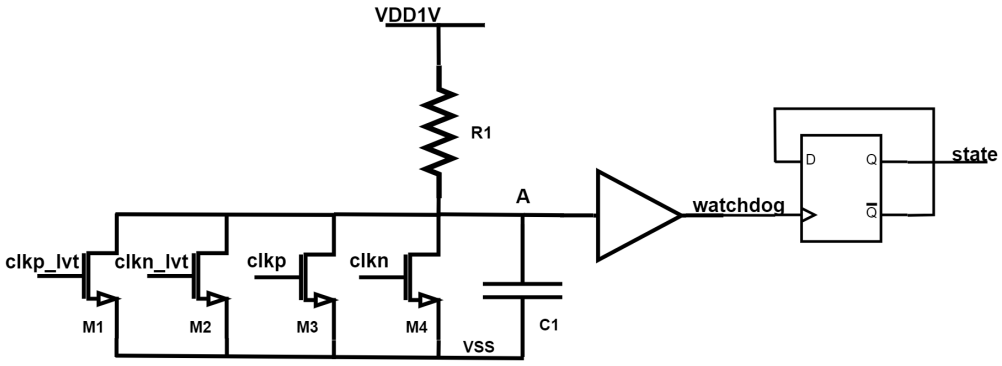

Figure 3.10: Schematic of the watchdog

The watchdog monitors the incoming ultrasound and generate the *watchdog* to clock the *state* into the right value. Figure 3.10 shows the schematic of the circuit. The *watchdog* signal remains low when there is incoming ultrasound, and switches to high when there is no incoming ultrasound. This is done by a series RC circuit and M1-M4. The ultrasound, depending on the state of the StimDust, is used to drive the gates of M1-M2 or M3-M4, which discharge the capacitor C1 to pull *watchdog* signal to low. During the time when there is no ultrasound, the RC circuit is charged in typical RC response, which in turn pull the *watchdog* signal to high. This *watchdog* signal is then used to clock the flip-flop that stores the state signal.

As mentioned above, either M1-M2 or M3-M4 are used to discharge the capacitor when there is incoming ultrasound. This is because when the StimDust is stimulating (Phase 3), the modulation transistor, M_{mod} is being turned on. This lowers the loading impedance seen by the piezocrystal, thus reducing the voltage seen by the chip. To continue to discharge the capacitor during phase 3, low V_T device is used for M1-M2. They remain off in all other phases.

3.6 TDC

Time-to-Digital converter, or TDC, is used to measure the duration of Ultrasound-Free Interval (phase 1) to set the amplitude of the stimulation. The heart of the TDC is a chain of current-starved buffers. Each of them has a fast fall time but a long rise time. Figure 3.11 illustrates the timing of the chain of current-starved buffers when the input is switched from low to high. At the output node of each buffer, a moscap is used to boost the capacitance to slow down the rise time. The simulation result is shown in Figure 3.12. Each of the current starve buffer provides about $10\mu s$ delay. The propagation delay from the In to T8 (for T8 to reach VDD/2) is about $93\mu s$ in simulation.

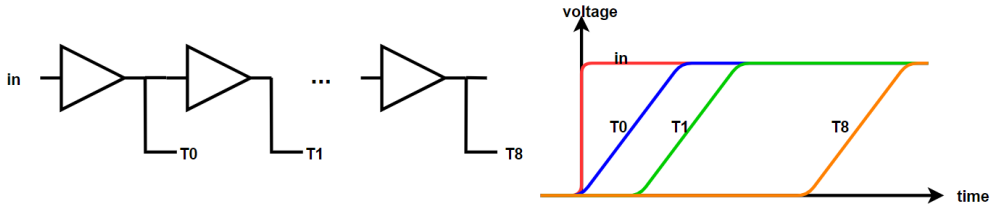

Figure 3.11: Timing diagram of the TDC when the input switches from low to high.

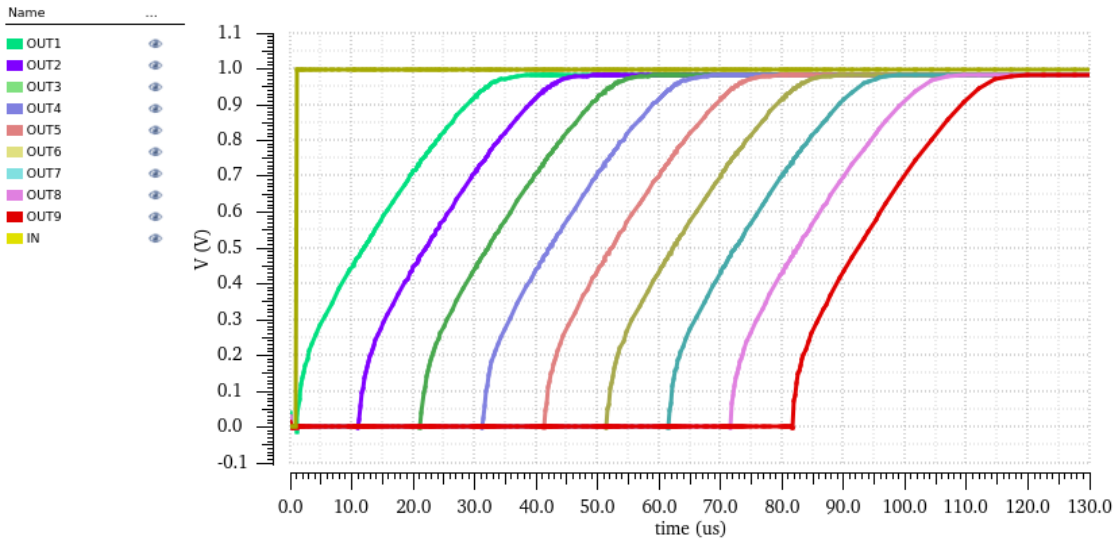

Figure 3.12: Simulation results of the TDC.

At the beginning of phase 1, the input to the TDC switches to high. The output of the current-starved buffers start to rise one after another. When the system transit from phase 1 to phase 2, the output of each buffer is latched to a flip-flop. The values of the flip-flops are used as the input to the current DAC to set the stimulation current amplitude.

3.7 Clock Extract

The core of the clock extract is a CMOS inverter with resistive feedback. The resistive feedback stabilizes the input and output biasing voltages to about $VDD/2$. At this bias voltage, the inverter has the highest gain, thus is sensitive to the input. A capacitive divider is used to extract and divide the $PZ+$ down to a lower voltage before passing to the input of the clock extract. The CMOS inverter then amplifies this signal and produce a clock. Figure 3.13 shows the schematic of this circuit.

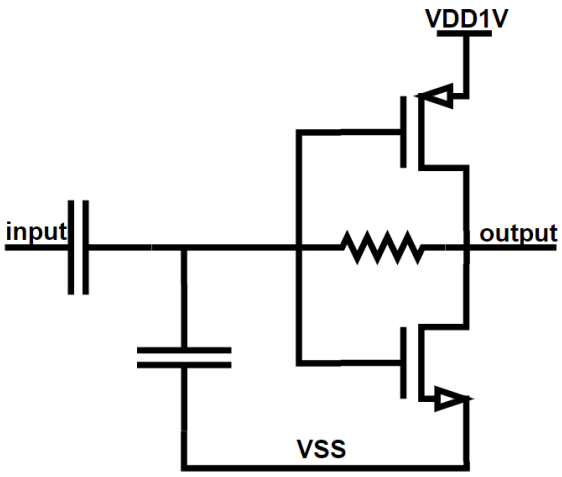

Figure 3.13: CMOS inverter in resistive feedback used to extract AC signal as clock.

3.8 Safety timer

This is the first new feature that is added to this chip and was discussed at the end of Chapter 1. As discussed in chapter 1 and 2, the safety timer limits the stimulation time to $\approx 1.1ms$. Thereafter, the chip stops the current output and immediately short the output to ground. To implement such function without increasing significant amount of area, the

counter and TDC are being reused. The output of the last buffer of the TDC is used to generate *TDC_clk* to be used as a clock to the counter during stimulation. Recall from Chapter 3.6, the propagation delay of each current starved buffer is about $10\mu s$. In this implementation, eight buffers in the TDC are used as a timer, thus the propagation delay is only about $80\mu s$, which is ≈ 14 times less than the 1.1ms target. Thus, the TDC will need to be able to recycle itself and a four bits counter is needed to count 14 cycles of TDC.

During stimulation (phase 3), the input to the buffer chain of the TDC is high if and only if the output of the last buffer, T8, is low. When the stimulation starts (phase 3), all the buffers' outputs are low. Thus, the input to the buffer chain in TDC will be high and the buffers' outputs will start to rise. Eventually, the output of T8 will rise to high. When T8 is high, the input to the buffer chain is pulled low. The buffers' outputs then start to fall. When T8 is back to low, the input is high again to start the second cycle. This process continues until the stimulation stops either because the safety timer is triggered or the circuit transit normally to another phase. The output of T8 is buffered and used as a clock to the counter. Figure 3.14 shows the timing of the TDC input and outputs during stimulation phase.

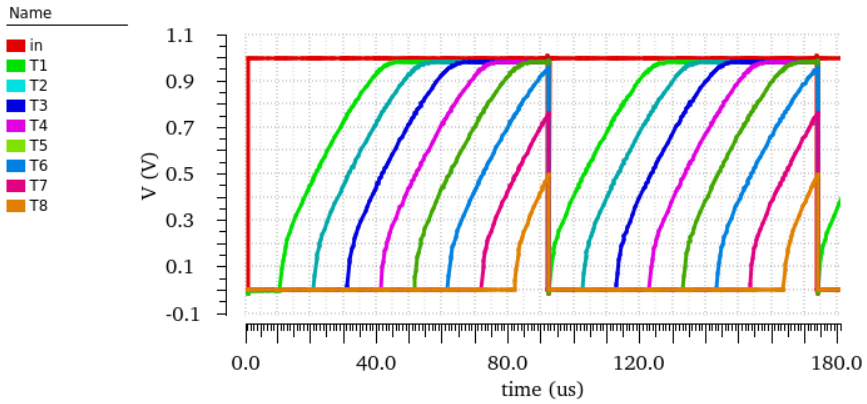

Figure 3.14: The TDC buffer chain rises until T8 is high. When T8 is high, the input becomes low to pull all the buffers' output to low. Once T8 is low again, the next cycle starts.

To reuse the counter, the protocol is changed to spare four LSBs in the counter. As mentioned in chapter 2, phase 2 (setup) now required 17 clocks instead of 16 clocks. In the previous generation, four flip-flops are used as a counter. To count 16 clocks before

stimulation starts, the first clock is used to release the reset signal, and the remaining 15 clocks are used to increment the counter. At the end of the 16 clocks, all four flip-flops are in state "1". Thus, ϕ_{stim} (when ϕ_{stim} is high, the stimulation starts) is simply determined by whether all the flip-flops are in state "1". In this generation, we added one more flip-flops to count 17 clocks. At the end of 17 clocks, the MSB flip-flop is "1", but the rest of the flip-flops are "0". Thus, ϕ_{stim} is simply dependent on whether the MSB is "1" or "0". This frees up the four LSBs of the counter to be used for the safety timer. When 14 *stim_stop* are being counted, the *stop* will be low briefly to reset the flip-flop that stores the *state* of the circuits.

3.9 Low peaking current source

This is the second new feature as discussed in Chapter 1 as an improvement to the previous design. Figure 3.15 shows the cascode current source that is being used. Assuming the gate of M4 is always connected to VDD for now, so M4 is always off (the purpose of M4 will soon be clear). In general, there are two spikes when the gate of M2 is switched from VDD to Vcasc during the time when the chip is transiting from setup (phase 2) to stimulation (phase 3).

The first spike is characterized by a short, moderate and instant peak when the chip transit to phase 3. During phase 2, the reference current is settling and the output is disabled by connecting the gate of M2 to VDD. Therefore, only M2 and M3 are off during phase 2. Since M1 is on, the drain voltage of M1 is at VDD. The output node, on the other hand, is at VSS, as it was shorted to VSS by M3 during power-up (phase 0) and short (phase 5). When the gate of M2 is being switched to Vcasc during phase 3, the drain of M1 (V_{D1}) starts to fall. However, since the initial $|V_{GS2}| = V_{DD} - V_{casc}$ is big, the parasitic capacitors at the drain of M1 quickly discharges through M2 to the output, causing a short and instant spike.

The second spike is, in general, longer in duration and larger in magnitude. It is due to a second order effect when V_{D1} quickly discharges, causing V_{G1} to drop below its nominal voltage. The C_{GD1} and $C_{GS1} + C_{GS5}$ forms a voltage divider from V_{D1} to V_{DD}. Since the current flowing at V_{G1} is limited by the current reference, not enough charge can be drawn

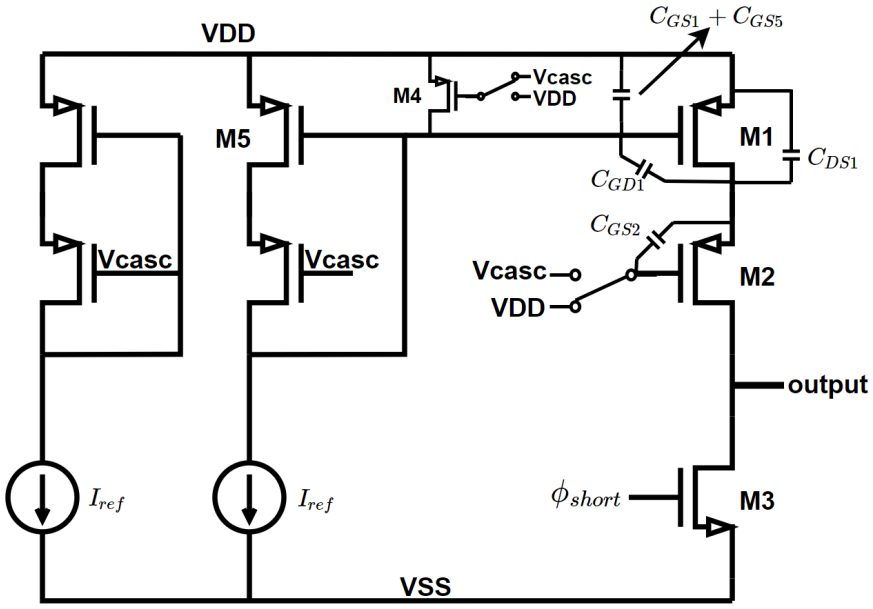

Figure 3.15: The schematic of the cascode current source being used in the DAC. The parasitic capacitors that is responsible for current peaking are also shown.

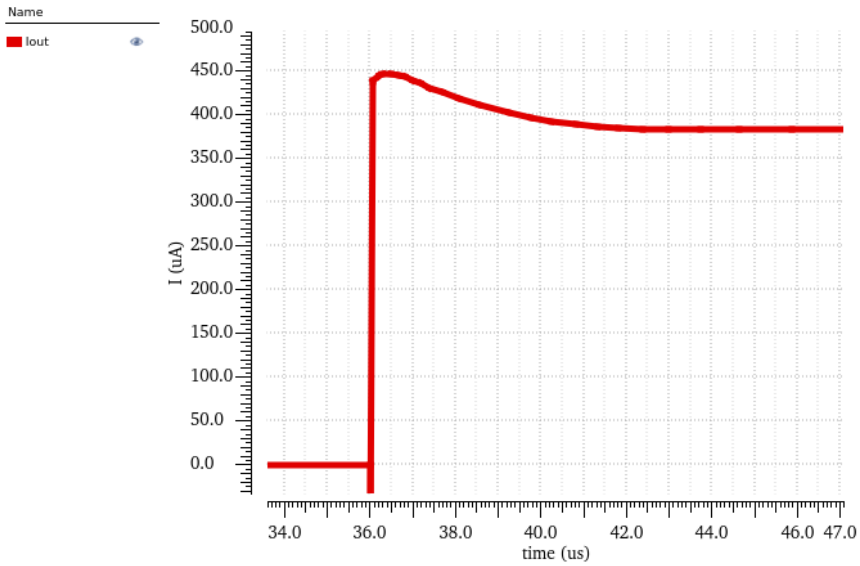

Figure 3.16: The simulation result showing the peaking of the output current. The initial overshoot is caused by the first spike. After that the second spike kicks in.

to maintain the voltage at V_{G1} when the drops at V_{D1} is coupled into V_{G1}. Figure 3.16 shows the simulation result of these spikes.

To minimize the current peaking due to the C_{GD1}, M4 is being used. During phase 2, M4 is being turned on. The effect of M4 is that it slightly lowers the $|V_{GS1}|$ during this phase and take the coupling as an advantage to quickly get to the desired bias voltage when transiting to phase 3. Figure 3.17 illustrates this idea. In general, this technique does not completely eliminate peaking, but it reduces peaking significantly without much penalty in area, settling time and power.

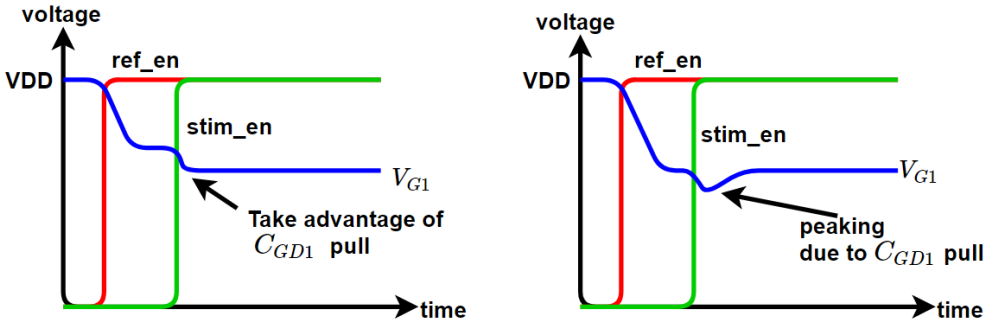

Figure 3.17: The M4 is used to bias V_{G1} slightly higher than its desired voltage to limit the peaking due to coupling from C_{GD1}

Chapter 4

Layout

As reported in [11], the active area in the previous StimDust was $150 \times 400\mu m^2$. With only $200 \times 500\mu m^2$ for this chip, the circuit-under-pads structure is used, meaning that the pads are directly on top of the circuits. Five pads are placed along the center line of the layout for PZ+, PZ-, VDD, VSS, and stim_output, as shown in Figure 4.1.

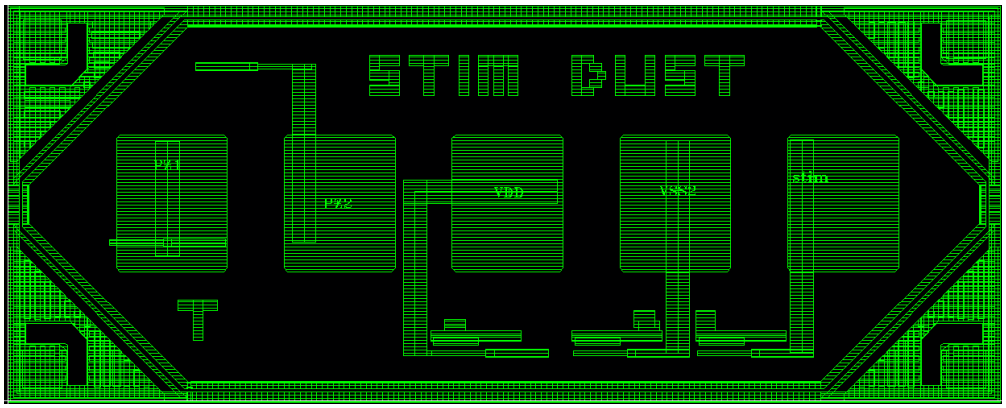

Figure 4.1: The top metal layer of the layout. The pads are evenly distributed along the center line of the layout. The structure around the pads is the sealring.

To minimize the coupling between the ac input pads (PZ+ and PZ-) and the circuits, circuits that are sensitive to external coupling are placed away from the area underneath the PZ+ and PZ- pads. For this design, the rectifier and the clock extract are sensitive to coupling and are placed away from the PZ+ and PZ- pads.

Moreover, compared with typical pad structure with a stack of metals, the circuit-under-pads structure has less support under it. Since transistor performances are sensitive to stress, if the structure is not stable, the wire bonding process can exert enough force that will reduce the performance of the transistors, or even damage the chip. For this reason, MIM capacitors are used directly under the layer beneath the pads for (1) decoupling the circuits, (2) provide shielding to the circuits underneath the pads, (3) consolidate the structure. Along with MIM capacitors, MOM capacitors are also used to fill most of the empty space of the layout for further decoupling and consolidation. Figure 4.2 shows that over half of the areas of the chip are filled with MIM or MOM capacitor.

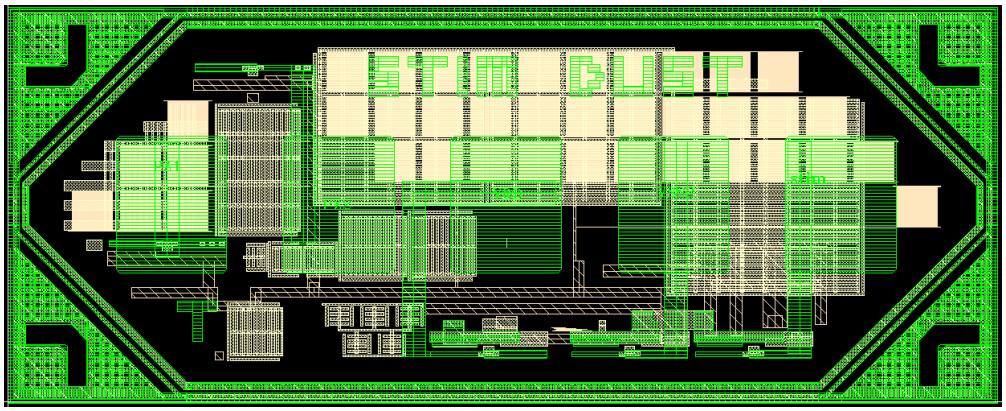

Figure 4.2: The layout showing that the chip is heavily decoupled using MIM and MOM capacitors to consolidate the structure. Green: top metal; Beige: MIM and MOM capacitors.

Chapter 5

Benchtop Measurements

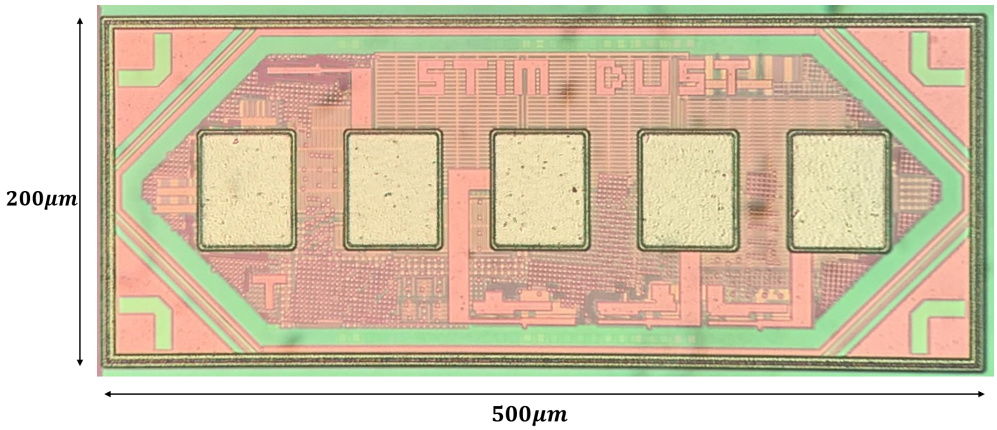

Figure 5.1: The die photo of the chip.

The design is fabricated in TSMC 65nm LP process. Figure 5.1 shows the die photo of the circuit. After the chip has been fabricated, five chips were packaged and characterized on benchtop. The functionality of the chips are first being tested to show its ability to response to stim commands and output the current. The chip is connected to a function generator that generates the stim command. The output of the chip is connected to a 5050Ω resistor in series to ground for measuring the output current. Figure 5.2 shows the schematic of the setup and Figure 5.3 shows the waveforms during stimulation. It shows that the chips are functioning as expected. After this sanity check, we characterized the power-up time, TDC transfer characteristic, current DAC transfer characteristic, safety timer, current peaking.

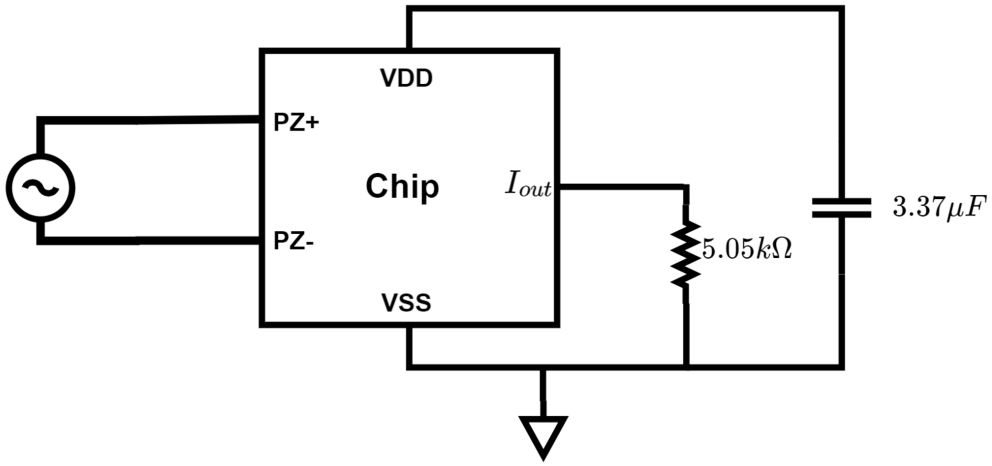

Figure 5.2: The test setup for functional testing.

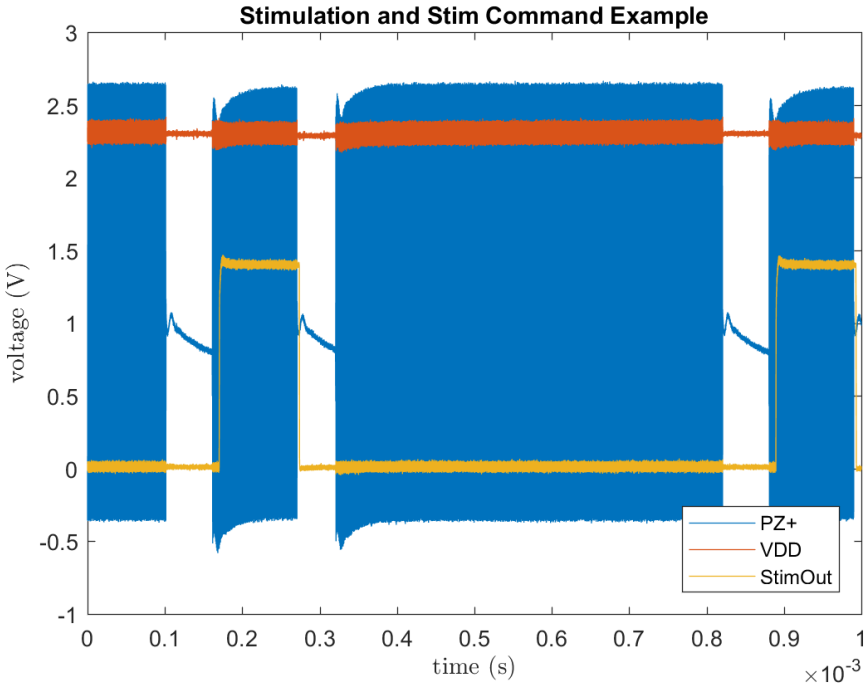

Figure 5.3: The waveform of the PZ+, VDD, and StimOut during stimulation.

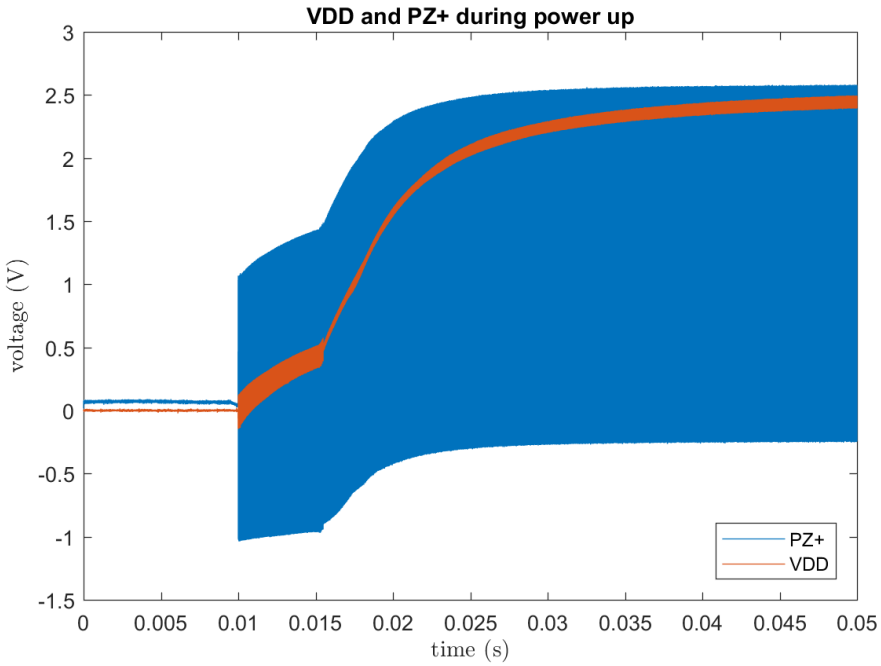

Figure 5.4: The waveform of the VDD of one chip during power-up measured by the oscilloscope.

We define the power-up time as the time for the VDD of the chip to reach 1.8V from 0V. Typically, the capacitance of a 0201 capacitor can be as high as few μF. As a result, we used a $3.37\mu F$ capacitor between VDD and VSS for this measurement. The setup is the same as the functional test mentioned above, except the input signal is a pure sinusoid instead of the stim command. The measured power-up time is 12.1ms for input Vpp (at PZ+ and PZ-) of 3V. To reach VDD=2.3V, it requires additional 13.4ms. Figure 5.4 shows the VDD waveform of one of the chips during power-up.

The TDC and current DAC transfer characteristics are measured by sweeping the duration of Ultrasound-Free Interval (UFI) in $0.1\mu s$ step. We measured the TDC trip points and the I_{stim} for all five chips and take their average to eliminate the effect of process variations. Figure 5.5 shows the average transfer characteristics of the five chips we have measured. The unit current sources match accurately with our design where $I_{unit} = 50\mu A$. However, the $T_{DC,unit} = 8\mu s$ is 20% less than that in simulation. One possible explanation is that the transistors are not intended to be used as a capacitor in this process. Thus, the model of

the capacitance of the transistor when used as MOSCAP may not be accurate. Although undesired, the users can select the appropriate current output by adjusting the Ultrasound-Free Interval. The stimulation frequency can also be preserved by changing the duration of the short phase (phase 5).

The linearity of the TDC and the current DAC within each chip are also measured by characterizing their differential nonlinearity (DNL) and integral nonlinearity (INL). Figure 5.6 and Figure 5.7 show the DNL and INL of the TDC and the current DAC respectively. The small DNL and INL show that both the TDC and current DAC are highly linear. The performances of the TDC and current DAC are summarized in Table 5.1.

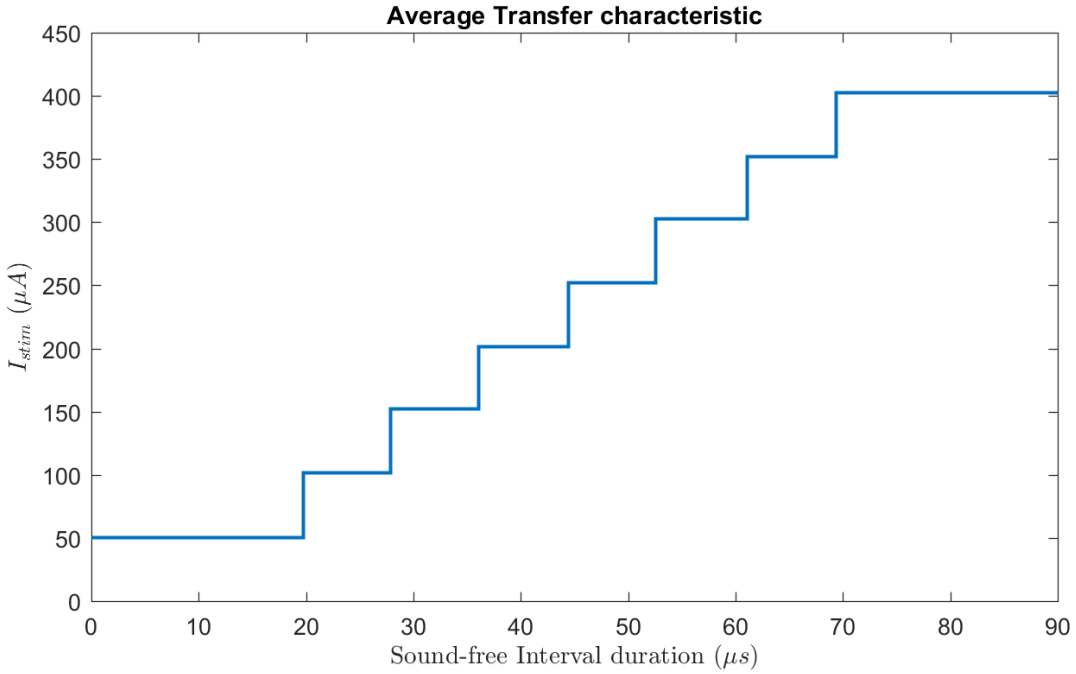

Figure 5.5: The average transfer characteristics of five chips.

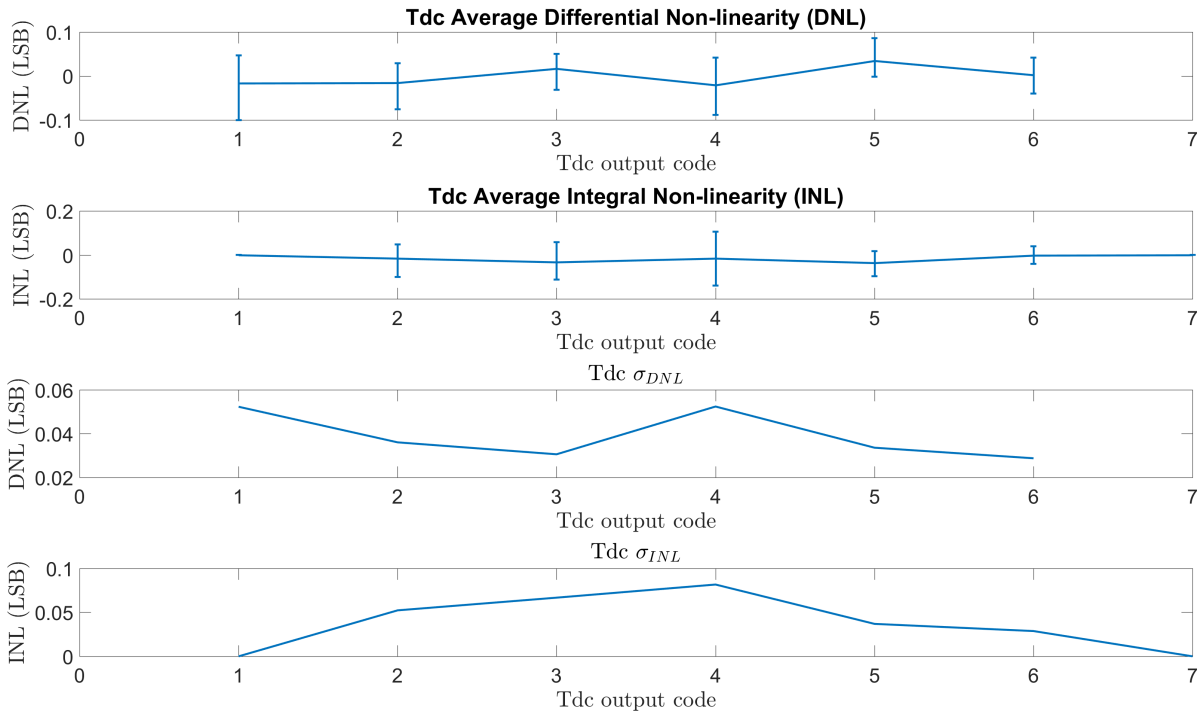

Figure 5.6: The nonlinearity of the Tdc. The error bar shows the range of DNL and INL across the five chips.

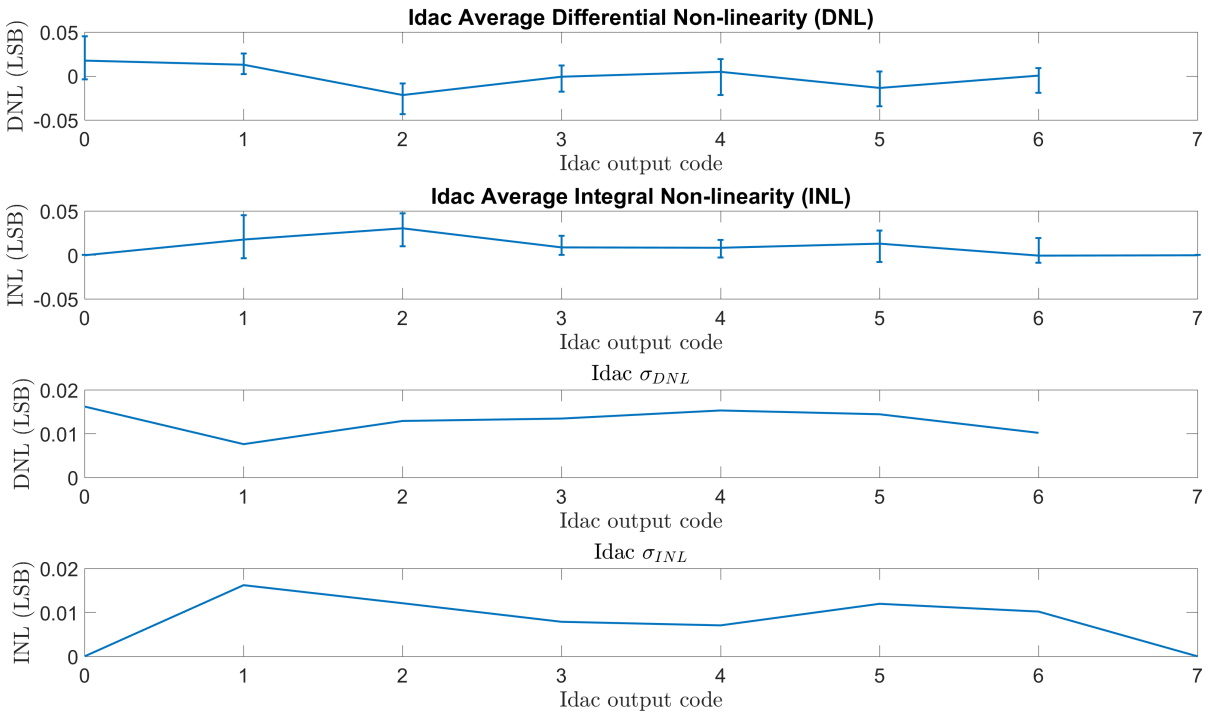

Figure 5.7: The nonlinearity of the Idac. The error bar shows the range of DNL and INL across the five chips.

	TDC	I_{DAC}
ΔT_{avg}	$8.28\mu s$	-
ΔI_{avg}	-	$50.2\mu A$
$\sigma_{\Delta T_{avg}}$	$< 0.4349\mu s$	-
$\sigma_{\Delta I_{avg}}$	-	$< 1.5822\mu A$
σ_{DNL}	$< 0.0524LSB$	$< 0.0162LSB$
σ_{INL}	$< 0.0816LSB$	$< 0.0162LSB$

Table 5.1: Summary of the performances of TDC and current DAC.

While the variation of the TDC and current DAC within each chip is small, some chip to chip variations are expected as the current references for TDC and current DAC are generated using polysilicon resistors that are sensitive to process variations. We quantified

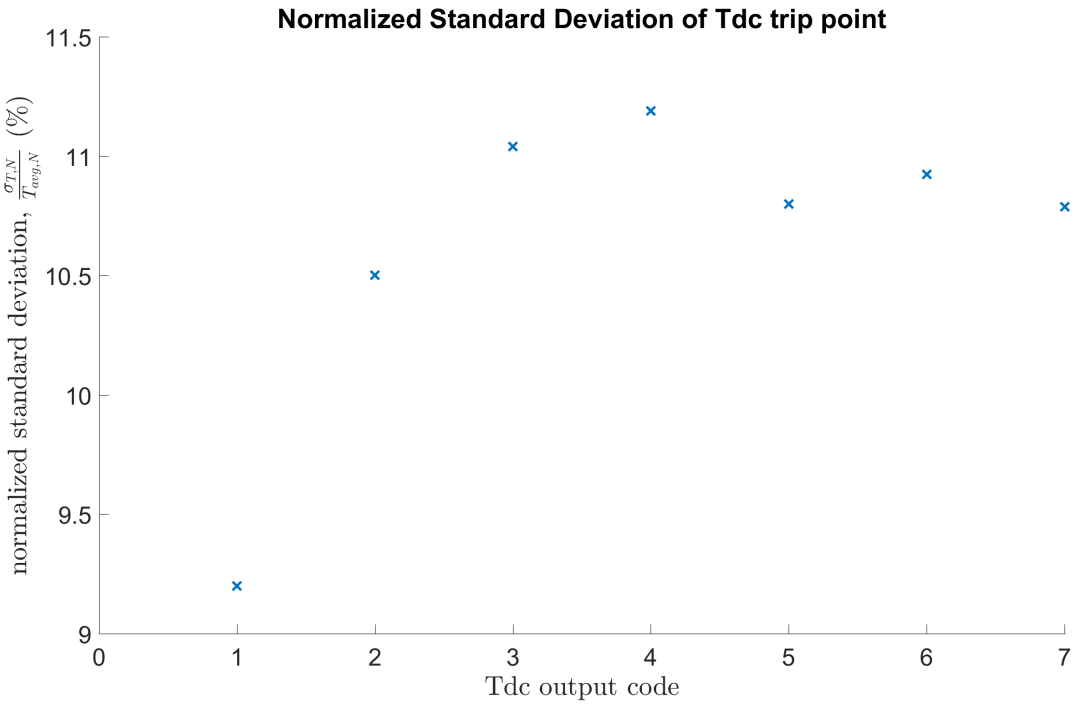

Figure 5.8: The standard deviation of the step size of TDC at each level, normalized by the average step size at that level.

the chip to chip variation by calculating the standard deviation of the step size of each TDC and current DAC level, then normalize the standard deviation by the average value of its corresponding level. Figure 5.8 and Figure 5.9 show the results of such measurements. As seen from the figures, the standard deviation of the TDC and current DAC are quite similar as the same type of resistor is used.

The safety timer is measured to be triggered at $888\mu s$ in average, which is less than the $1.1ms$ in simulation by 20%. This can be explained by the same argument that the TDC is running faster than that in simulation. Figure 5.10 shows the waveforms when the safety timer is being triggered. In this example, the stimulation stops after 1.04ms. The next stream of ultrasound starts the stimulation as expected.

Current peaking is also improved. The peaking is measured to be 5% in average when the maximum output current is used. This corresponds to the worst case scenario since the more unit current sources are being turned on, the stronger the coupling to the biasing will

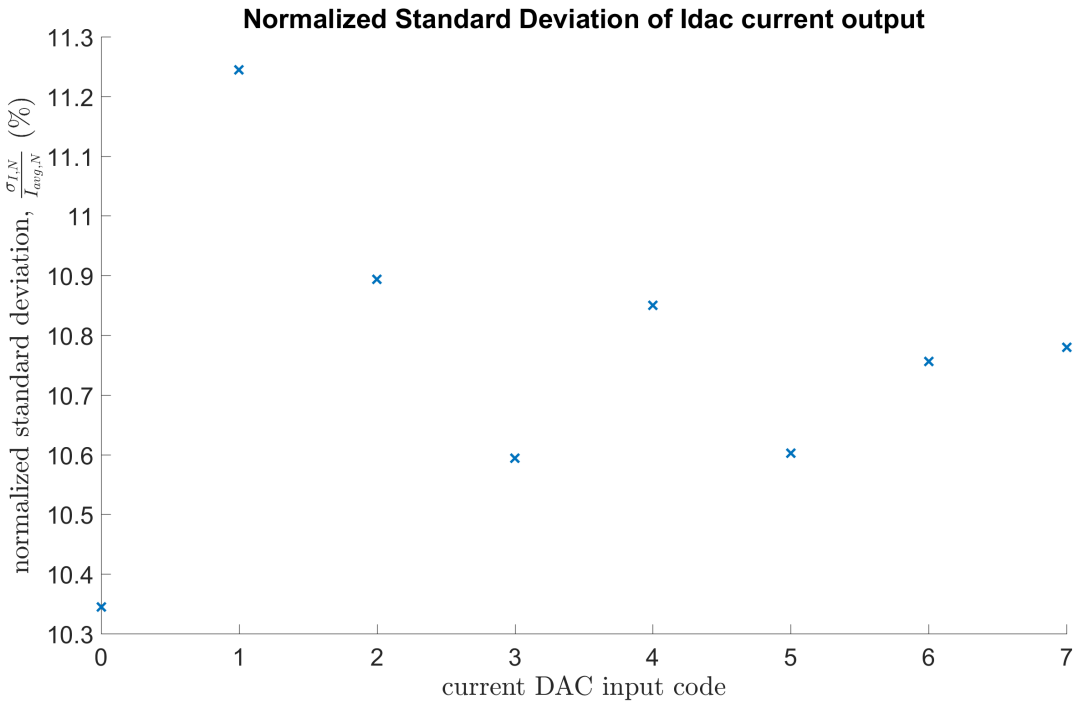

Figure 5.9: The standard deviation of the step size of the current DAC at each level, normalized by the average step size at that level.

be. For the same amount of peaking, the technique being used can greatly reduce the size of the decoupling capacitor used to stabilize the bias voltage. Figure 5.11 shows the current waveform normalized by the static current amplitude.

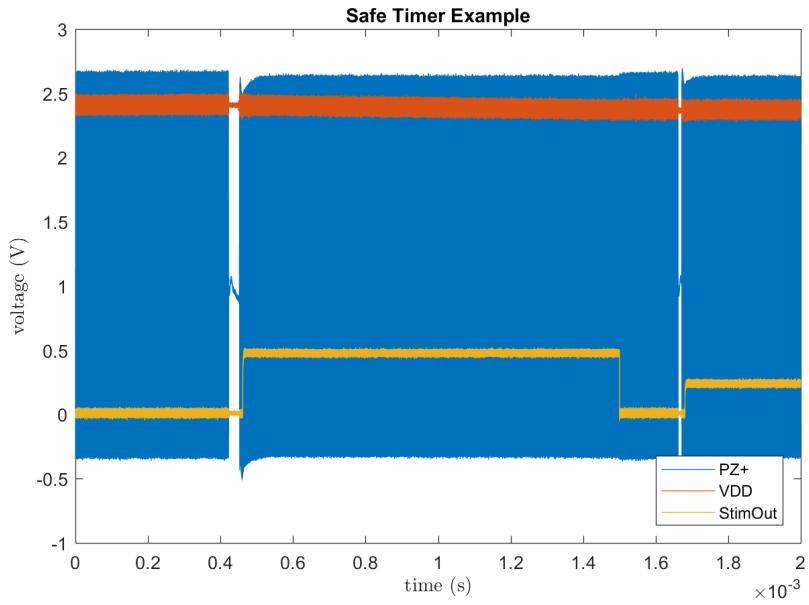

Figure 5.10: The waveforms example showing that the safety timer stops the stimulation after 1.04ms. The state of the chip is flipped such that the next pulses of ultrasound immediately starts the stimulation.

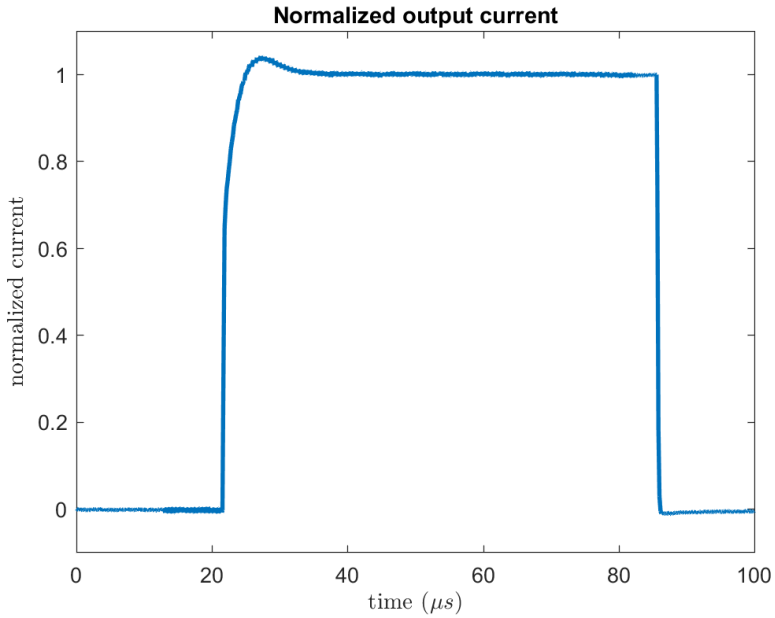

Figure 5.11: The waveform of the maximum current output normalized by the static current amplitude. The peaking is less than 5% in this example.

Chapter 6

Conclusion

The chip is tested on the benchtop and shows promising results. With a size of μm scale, it will enable a sub-mm mote, which has not been achieved to date. The next step for this chip will be to assemble the mote and test it *in-vivo*. The next generation of StimDust can improve compatibility in multi-mote operation and have better current and voltage references across process variation. For the former, the downlink should be changed to allow data transfer to provide proper addressing for each command and also increase the resolution of output current DAC, thus allowing multi-mote operation and increasing the simulation resolution. Low power data transfer has been demonstrated in the past. Typically, amplitude-shift keying (ASK) is suitable for wirelessly powered implants, as it can maintain a certain level of power transfer during data transmission. Moreover, since the piezocrystal will continue to ring for some time after the external pressure wave stops, some modulation scheme such as on-off keying (OOK) that requires a sharp transition of the waveform will be difficult to provide enough data rate. Thus, ASK modulator for the downlink seems to be a good choice. One promising example is low modulation depth ASK modulator shown in [13]. This technique can provide enough data rate without significant penalty in power transfer.

The current and voltage references of this chip can also be improved. As shown in 3.4, the current reference is generated by referencing to a polysilicon resistor. The polysilicon resistor varies significantly over process corners, thus changing the stimulation current from the designed values. To fix this, either trimming circuit or a better current reference is

needed. For more accurate current reference, MIM or MOM capacitor can be used instead of the polysilicon resistor, as they vary less over corners and can be fabricated with high accuracy.

Bibliography

[1] Konstantin V. Slavin. "History of Peripheral Nerve Stimulation". In: *Progress in Neurological Surgery* 24 (2011), pp. 1–15.

[2] Dennis J Bourbeau et al. "Genital nerve stimulation increases bladder capacity after SCI: A meta-analysis". In: *The Journal of Spinal Cord Medicine* (2017), pp. 1–9.

[3] Yen-Fu Chen et al. "Occipital Nerve Stimulation for Chronic Migraine-A Systematic Review and Meta-Analysis". In: *PLoS ONE* 10(3): e0116786 (2015).

[4] Richard D. Wilson et al. "Peripheral Nerve Stimulation Compared to Usual Care for Pain Relief of Hemiplegic Shoulder Pain: A Randomized Controlled Trial". In: *American journal of physical medicine & rehabilitation / Association of Academic Physiatrists* 93(1) (2014), pp. 17–28.

[5] Matthew Schiefer et al. "Sensory feedback by peripheral nerve stimulation improves task performance in individuals with upper limb loss using a myoelectric prosthesis". In: *Journal of Neural Engineering* 13.1 (2016), p. 016001.

[6] Dongjin Seo et al. "Wireless Recording in the Peripheral Nervous System with Ultrasonic Neural Dust". In: *Neuron* 91.3 (2016), pp. 529–539.

[7] J. H. Schulman. "The Feasible FES System: Battery Powered BION Stimulator". In: *Proceedings of the IEEE* 96.7 (July 2008), pp. 1226–1239. ISSN: 0018-9219. DOI: 10.1109/JPROC.2008.922588.

[8] "IEEE Standard for Safety Levels with Respect to Human Exposure to Radio Frequency Electromagnetic Fields, 3 kHz to 300 GHz". In: *IEEE Std C95.1-2005 (Revision of IEEE Std C95.1-1991)* (Apr. 2006), pp. 1–238. DOI: 10.1109/IEEESTD.2006.99501.

[9] Dongjin Seo et al. "Neural Dust: An Ultrasonic, Low Power Solution for Chronic Brain-Machine Interfaces". In: *arXiv* (July 2013).

[10] FDA. "Guidance for Industry and FDA Staff: Information for Manufacturers Seeking Marketing Clearance of Diagnostic Ultrasound Systems and Transducers". In: (2008).

[11] Benjamin C. Johnson et al. "StimDust: A 6.5mm^3, Wireless Ultrasonic Peripheral Nerve Stimulator with 82% Peak Chip Efficiency". In: *2018 IEEE Custom Integrated Circuits Conference (CICC)* (Apr. 2018).

[12] Paul R. Gray. *Analysis and Design of Analog Integrated Circuits, Fifth Edition*. Wiley, 2009.

[13] Y. Wang et al. "A 13.56MHz Wireless Power and Data Transfer Receiver Achieving 75.4% Effective-Power-Conversion Efficiency with 0.1% ASK Modulation Depth and 9.2mW Output Power". In: *2018 IEEE International Solid - State Circuits Conference - (ISSCC)*. Feb. 2018, pp. 142–144. DOI: 10.1109/ISSCC.2018.8310224.