

Chapter 7

Conclusions and Future Work

7.1 Thesis Conclusions

This thesis presents a theoretical method of drastically reducing the power consumed in tissue stimulation, as well as a prototype implementation. Of the three main sections of the prototype system, the RF power transmission system of Chapter 3 worked as planned, the power handling circuitry of Chapter 4 worked sufficiently well, with a few design flaws and a few surprises due to layout flaws, and the electrode switching circuitry of Chapter 5 worked sufficiently well, with minor flaws due to lack of sufficient modeling and simulation, and substrate pickup from the 125 KHz magnetic fields.

Table 7.1 summarizes the power savings of the system presented in the preceding chapters. In this table, the term overhead refers to the power consumed by the system when it is not driving electrodes. This includes references, control circuitry, and any standby power consumption of the stimulation system.

It is worth exploring, to the extent possible, the power losses associated with this system. First, the measured power into the electrodes, $49 \mu W$, is 2.3 times the theoretical minimum for delivering the same charge. Theoretical calculations with ideal steps showed that a 4-step system could deliver the charge using only 15% more power than the minimum. However, the steps used here are equal steps, and do not perfectly account for the IR drop offset in the electrode voltage. Another possible

Theoretical	Minimum Possible Power	21.2 μW
This System	Measured Power Into Electrode	49 μW
This System	Net Power During Stimulation (Not Including Overhead)	125 μW
This System	Gross Power During Stimulation (Including Overhead)	148 μW
Aggressive Current Source System	Net Power During Stimulation (Not Including Overhead)	271 μW
Typical Current Source System	Net Power During Stimulation (Not Including Overhead)	373 μW

Table 7.1: Power per Electrode Comparison

reason for this extra power loss would be a change in electrode impedance over time. However, the electrode impedance was measured after several weeks of testing with this system, and the average resistance and capacitance were 2.28 $K\Omega$ and 450 μF , practically the same resistance as before, with slightly less capacitance.

The bulk of the rest of the losses occur between the coil and the electrode connection. Starting from the electrodes, this includes the electrode switches and their drivers, extra power consumed in the control circuitry due to electrode loading, extra power consumed in the rectifier switches and drivers, and leakage paths from the storage capacitors. The electrode switches were large enough (PMOS $250\lambda / 2\lambda$, NMOS $100\lambda / 2\lambda$, $\lambda = 0.8 \mu m$) that they added negligible resistance compared to the electrode resistance, yet they were not so large that their drivers consumed large switching power. The control circuitry consumes practically no more power during stimulation than during standby, and its total power is less than 200 μW . The likely source of extra power consumption is the rectifier switches, and the reason for this power loss is the simplistic algorithm used to charge the storage capacitors. The rectification durations were set conservatively to supply more than enough power for the electrodes, and the durations came out longer than expected due to process variations in the one-shot delay circuit. This long rectification time allows a large voltage to develop across the rectifier switch, burning unnecessary power. Finally, the leakage resistors (very long transistors operating in their linear regions) on the storage capacitors were designed to leak very little current compared with that consumed by the

electrodes, and judging from the standby (no electrode drive) power consumption, they function as designed.

Table 7.1 shows where power was burned in each portion of the system, as a measure of total power when driving 15 electrodes. The minimum possible power is far smaller than the total power, so there should be room for improvement. The entry labeled extra power into the electrodes is the difference between the calculated electrode power and the theoretical minimum. The standby or overhead chip power was taken directly from the measurements in Table 6.4. The power consumed in the chip due to stimulation is the total measured chip and electrode power from Table 6.4 (2.18 mW) minus the standby power and power into the electrodes. The coil power numbers are taken directly from Table 6.4.

Theoretical Minimum Possible Power	$15 \times 21.2 \mu\text{W} = 318 \mu\text{W}$
Extra Power Into Electrodes	$15 \times (49 - 21.2) \mu\text{W} = 417 \mu\text{W}$
Standby (Overhead) Chip Power	$335 \mu\text{W}$
Power Consumed in Chip Due to Stimulation	$(2.18\text{ mW} - \text{above}) = 1.11\text{ mW}$
Standby (Overhead) Coil Power	$2.79 \mu\text{W}$
Coil Power Due to Stimulation	$44.4 \mu\text{W} - 2.79 \mu\text{W} = 41.6 \mu\text{W}$
Total Consumed Power	2.22 mW

Table 7.2: System Power Consumption, 15 Electrodes

From the numbers in the table, one in particular is noteworthy. The power consumed in the chip due to stimulation is quite high, half of the total power for the entire system. This is likely due to the voltage developed across the rectifier transistors due to the long rectification times.

While the system as a whole has some room for efficiency improvement, it meets the goal of using far less power than the competing stimulation systems listed so far. However, it is possible that such systems could be designed more efficiently. A current source based system with lower voltage rails will use less power. If the neural tissue responds to total charge levels, and exact current control is not necessary, then the voltage supplies may be made quite low. This approach may preclude the use of cascode transistors, and may drastically reduce the output impedance of current

sources, but those are acceptable tradeoffs for the potentially vast reduction in power. Another means of reducing current source power is to simply switch the electrode directly to its return for a brief time at the beginning of the second current phase, and then drive positive current into the electrode at the end of the second current phase to ensure charge balancing. The one caveat for these methods is that if the voltage rails get too close to the electrode voltage, the supply voltage may be too low for operation of many of the control circuits. In addition, voltage rails which are too close to the voltage of the charged electrode may need to move if the electrode impedance shifts over time. Moving the supply rails can affect performance of control circuitry designed to run from a specific supply. The best compromise is to separate the control circuitry supply, using a moderate voltage to supply the relatively low-current controls, and a low voltage to supply the high-current electrodes. However, while the highest supply can easily be generated from a simple diode rectifier, the lower voltage supply (which must vary with electrode impedance) must be generated from some form of synchronous rectifier with controls to keep it at the correct voltage. So as a current source based stimulator moves toward lower supply voltages, it begins to resemble the system described in this thesis.

7.2 Future Work

While this stimulation system worked well, it had a few bugs and a few areas where the general design could be improved.

7.2.1 Changes to This Implementation

The first change to this version of the system would be the substrate and well isolation of the clocked comparators and their clocks, as described in Section 4.3.9. In general, much of the circuitry needed to have better substrate contacts. Specifically, the digital clock divider discussed in Section 5.2.1 was too sensitive to substrate noise. A more traditional D flip-flop clock divider, while taking more area, may have been more reliable for this application.

7.2.2 Broader Changes

The first area to target design changes is the section of the system which burns the most power: the synchronous rectifier. While the predictive comparator works perfectly and the rectifier switches are quite large (PMOS $1250\lambda/2\lambda$, NMOS $500\lambda/2\lambda$), the switches are simply turned on for too long a time. This allows a voltage to develop across the switches, costing power. In addition, significant current develops in the coil, which rings when the switch is turned off. If this is done, a timing method is needed which is more precise and more controllable than the one-shots used here. The one-shot delays may be made both more precise and controllable by using current-starved inverters, but some feedback method must be used to set their timing. The ideal method is to have a linear differential amplifier monitoring each capacitor and its associated reference voltage, with a low-pass filter to take a time average over the last few waveform cycles. If the capacitor voltage slips a few tens of mV below the reference voltage, the rectification time is increased. Also, adding a brief fixed duration rectification period onto the back side of the waveform would spread out the charging over time, reducing the required peak current. This can be implemented easily with an over-compensated predictive comparator, which switches before its inputs are equal.

Another method of improving control over the rectification timing is to use a clock which is much faster than the 125 KHz system clock. This requires a phase-locked loop, which may be more complexity and power than is necessary, but it would allow completely digital control of the rectifier timing. In order to get timing resolution that would compete with the analog system suggested above, the on-chip oscillator would have to operate at 20-50 MHz, consuming substantial power.

The method of leaking charge from the storage capacitors described in Section 4.3.7 is somewhat ad hoc, and should be replaced with some intelligent method of maintaining capacitor voltage. Using a charge pump for the rectifier switch gates will work, but is not worth the effort, especially the effort of building a below-ground charge pump, which requires floating diodes. Since the power lost through this path, even

with leakage transistors described in Section 4.3.7, is rather small, a simple means can be implemented of leaking off some charge when the capacitor voltage exceeds the reference voltage by too great a margin. The same linear differential amplifier described above could be the basis for this measurement.

Switching off the current through the coil during rectification caused a large voltage swing and resonance on the AC voltage, which contributed to substrate noise. This noise may be abated with a simple snubber circuit, like a series resistance and capacitance, across the secondary coil. The addition of a snubber may slightly increase power, so this should be done in conjunction with substantial redesign of the rectification timing controls.

One circuit change may make this design more tolerable to nerves by generating more nearly uniform stimulation current without substantially increasing the power consumption. Current source transistors may be added to each voltage level, limiting the output current immediately after the electrode is switched to it. As the electrode charges to nearly the voltage on that storage capacitor, the transistor begins to act like a switch instead of a current source. This would spread the current throughout the time during which the electrode is switched to that capacitor, while delivering only slightly less charge per level than the purely switch-based method.

Finally, as mentioned in Section 3.4, some measurement of the chip V_{dd} should be sent telemetrically to whatever circuit is driving the primary coil. This is potentially a difficult problem, since the data rate should be on the same order of magnitude as the power frequency, 125 KHz. Specifically, the V_{dd} data should be sent at approximately 125 KHz divided by the Q of the primary coil resonance system, or in the tens of KHz range. The back telemetry could be included as part of a higher frequency data communication system.