Intelligent Ultrasound Imaging System

Master of Engineering Capstone

Project Report
BSAC - Liwei Lin Lab

Team 50 Authors:
Tofic Esses, Ryan Johnson, Linda Liu, Suraj Chamakura

May 4, 2026

Contents

1 Executive Summary

Ultrasound imaging has the potential to transform manufacturing, healthcare, and robotics, much as LiDAR revolutionized 3D perception in self-driving cars. However, the key limitation of current ultrasound systems is that they require a large number of sensors and heavy computation to produce a clear 3D image, which makes them expensive and slow.

The project tackles the limitation by designing both the hardware and software using a technique called compressed sensing (CS). Compressed sensing allows us to reconstruct high-quality images from far fewer sensors than traditionally required as long as those measurements from the sensors are carefully designed.

Instead of relying on thousands of sensors, we built a compact system with eight ultrasound sensors on a single printed circuit board (PCB). By optimizing the signals transmitted by each sensor, the system generates unique measurements, which ultimately enable accurate reconstruction even with significantly limited hardware.

On the software side, coupled with a computationally efficient forward model and GPU-accelerated reconstruction, the system reconstructs a sparse $1\,\mathrm{m}^3$ scene at $1\,\mathrm{mm}$ resolution in seconds rather than hours.

Overall, this work highlights the potential of compressed sensing in ultrasound to dramatically reduce hardware requirements while maintaining high-resolution reconstruction.

2 Limits of Traditional Ultrasound and the Need for Compressed Sensing

Traditional ultrasound imaging systems require dense sensor arrays and a high sampling rate to obtain a reconstruction with high resolution. This results in complexity in hardware and a significant computational power usage during reconstruction [1]. This hence creates a scalability issue as the field of view and resolution increase.

Compressed sensing (CS) offers an alternative by enabling the reduction of the number of measurements to reconstruct a high-resolution image. In theory, CS can significantly reduce both sampling rates and hardware complexity while preserving final image quality [2].

Although CS has been successfully adopted in various fields and has recently been explored in ultrasound imaging, the existing applications still face challenges in reconstruction efficiency and quality in practice [3].

In addition, most mature applications of ultrasound nowadays are restricted to clinical settings. Within clinical settings, system size, cost, and reconstruction efficiency are less constrained compared to the reconstruction quality. In contrast, applications in various autonomous systems, such as robotics and manufacturing, are relatively underexplored. The reason largely lies in the difficulty of building compact and real-time ultrasonic imaging systems that provide reliable reconstruction images. Hence, these limitations and attempts of CS in ultrasound motivate a hardware-software co-designed system for ultrasound imaging.

3 Initial State of Compressed Sensing for Ultrasound in Literature

Compressed sensing has been applied extensively to optical imaging and has recently been applied to ultrasound imaging systems. The work by Kruizinga et al. demonstrates that 3D ultrasound imaging can be achieved by reducing hardware by using compressed sensing theories [4]. They used a single ultrasound sensor combined with a mechanically rotating aperture to recover 3D spatial information. They show that 3D reconstruction is possible from a limited number of acquisitions by exploiting sparsity in the scene and encoding diverse measurements.

Their design, however, reveals key challenges, which include the need for sufficiently diverse measurements and efficient reconstruction algorithms. Our work eliminates the mechanical aperture by using a synthetic aperture, where measurement diversity is achieved through optimized transmission profiles. Hardware design, array geometry, and transmission profiles are optimized to further improve measurement diversity and reconstruction quality.

4 System Level Overview

The full system as seen in Figure 1 includes the target object, the electronic system, and the computer with the compressed sensing reconstruction algorithm.

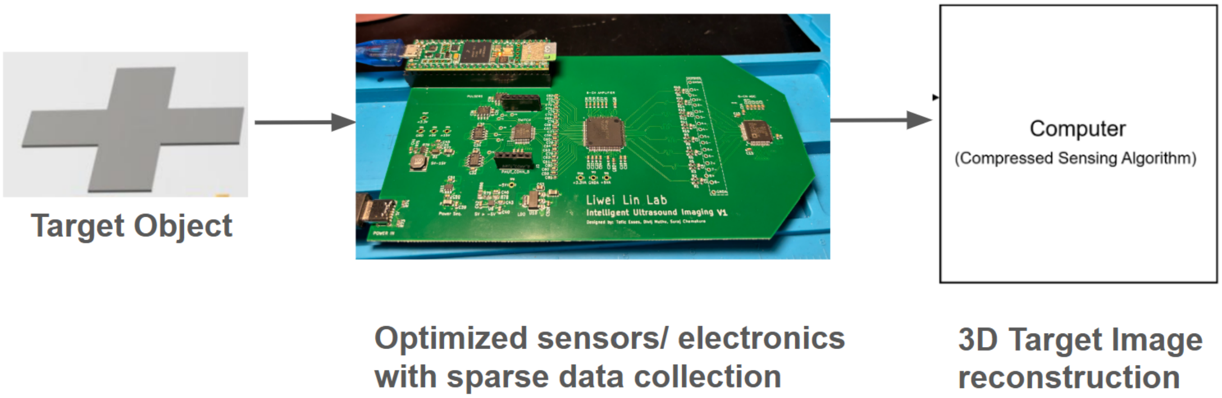

Figure 1: System Overview

4.1 Miniaturization of Electronics from Prior Works

A substantial portion of this project was miniaturizing the electronics of the old design seen in the left hand side of Figure 2 and designing a compact system as seen in the right hand side of Figure 2. In both cases, the system accomplishes the same task, but the old system is not portable, is messy, and is difficult to debug.

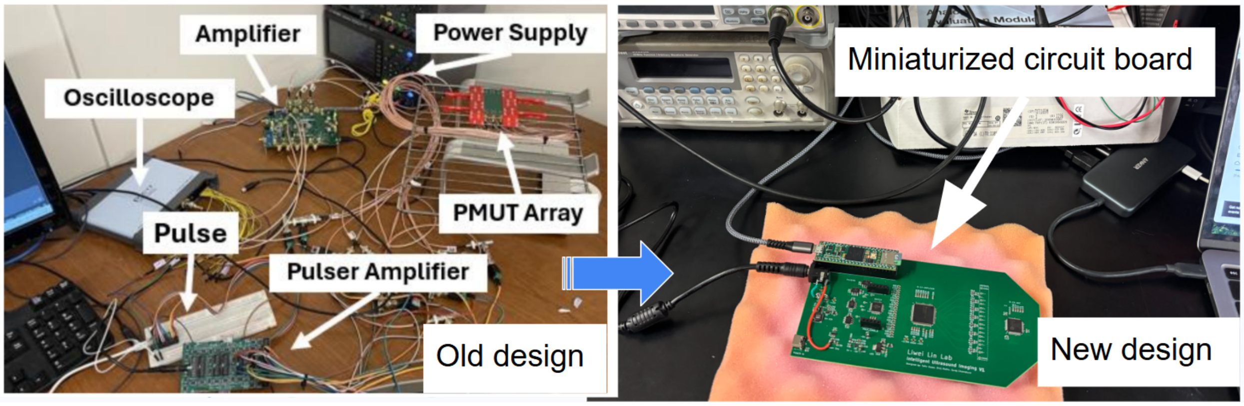

Figure 2: Before and After of Electronic System

5 Electronic System Overview

The electronic system is responsible for enabling the Piezoelectric Micromachined Ultrasonic Transducers (PMUTs) to transmit ultrasound, and to also receive ultrasound. This can be divided into two pipelines - transmitting and receiving. The transmitting pipeline contains pulsers that apply a high voltage (15V) pulsing signal to the PMUTs. This allows the PMUT to vibrate at the desired frequency. Similarly, when the PMUTs receive the echo, the mechanical vibrations induce an electrical signal. This electrical signal is captured by the receiving pipeline. The receiving pipeline contains an amplifier to amplify the low-voltage signals received from the PMUTs. This then passes through an Analog-to-Digital Converter (ADC) that converts it into a digital signal via quantization. This is necessary for a computer to interpret and process the signal. Note that both pipelines are multiplexed using a transmit/receive switch. This behavior is realized by the system level diagram in Figure 3.

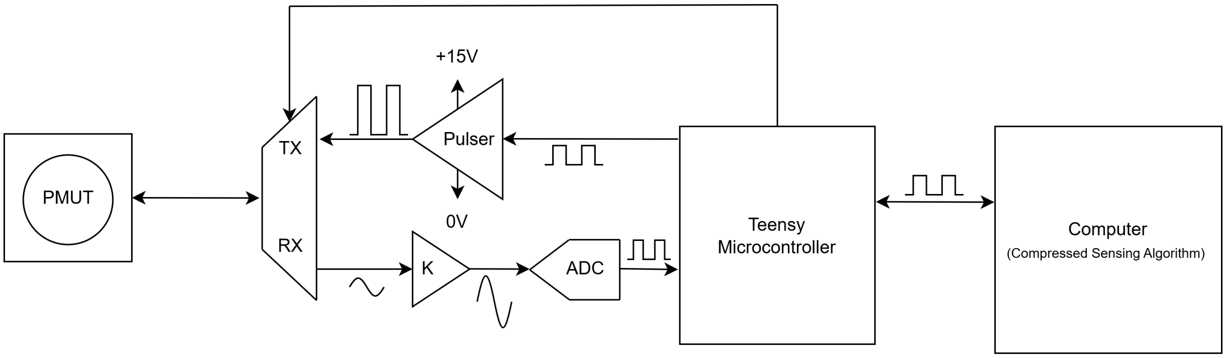

Figure 3: Electronic System Diagram

5.1 Piezoelectric Micromachined Ultrasonic Transducers (PMUTs)

The PMUT is a Micro-Electro-Mechanical System (MEMS) that is capable of converting a mechanical vibration into an electrical signal and an electrical signal into a mechanical vibration. This gives it transducive properties that enable this device to transmit and receive ultrasound and to be interfaced by an electrical system.

5.2 Printed Circuit Board Implementation

5.2.1 Electronic Schematics

The top level schematic is the full electronic system at a higher level of abstraction. This can be seen in Figure 4. By dividing the system's top level schematic this way, it allows for isolating each part and for working on it independently.

.

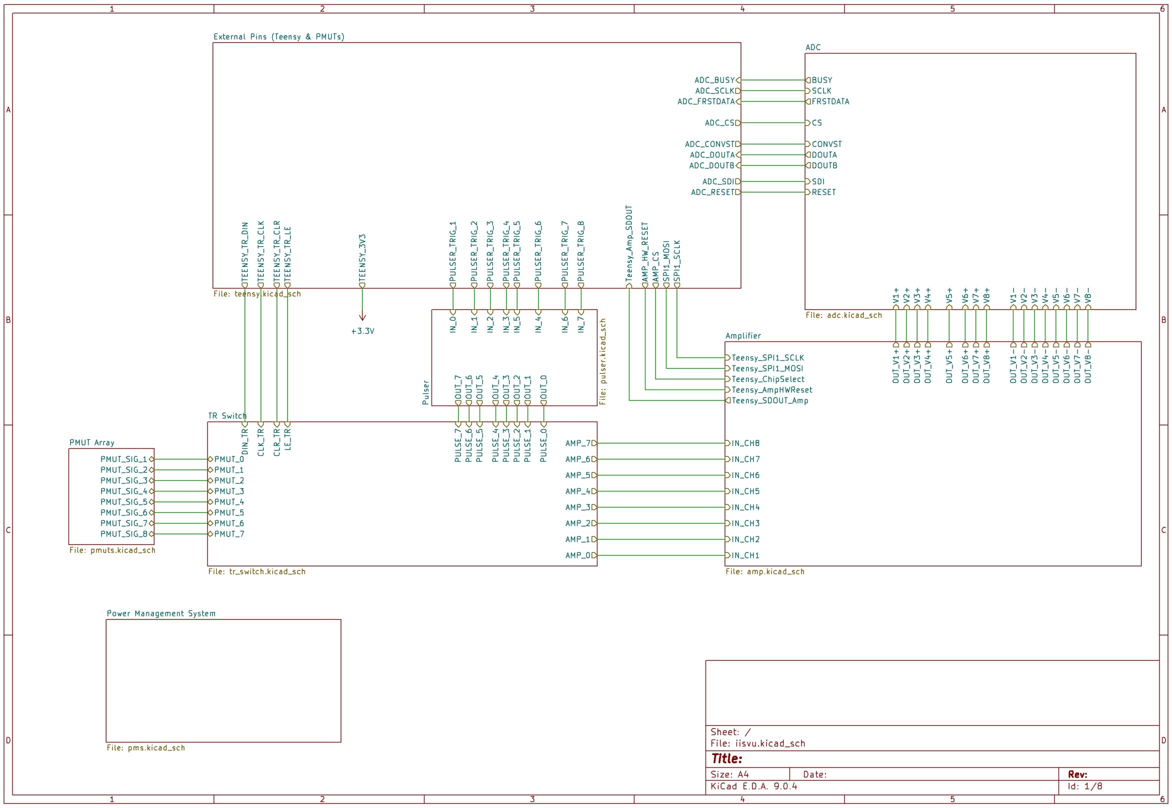

Figure 4: Top Level Schematic

Each schematic sub-sheet in the top sheet corresponds to a subsystem of the electronic system. The Teensy, Amplifier, ADC, Pulsers, TR-Switch, PMUT Breakout schematics, and Power Management System can be found in Fig. 5-11.

The Teensy Connections schematic defines the I/O mapping of the microcontroller to the peripheral circuitry for control and data collection. The microcontroller controls the pulsers, TR switch, amplifier, and ADC. Similarly, the data collection occurs entirely from the output of the ADC.

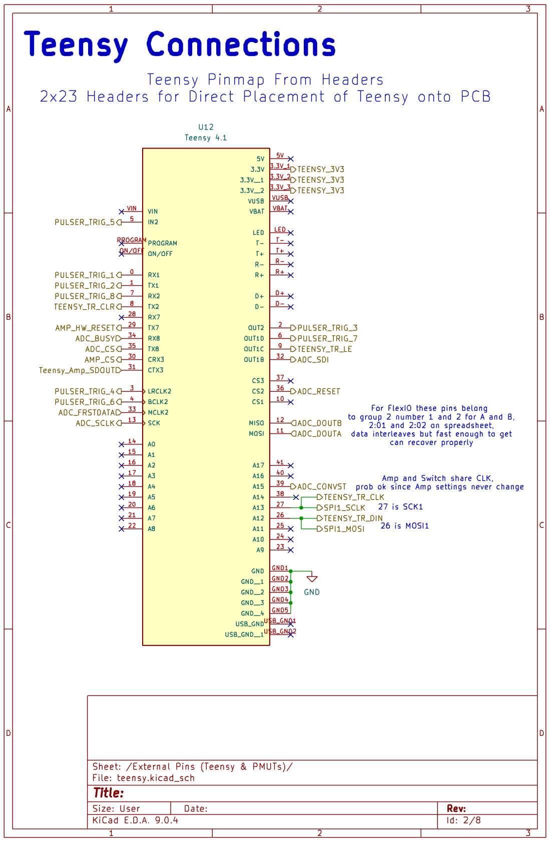

Figure 5: Teensy Schematic

The VCA5807 is an 8-channel integrated voltage controlled amplifier (VCA) specifically designed for ultrasound applications. By choosing this component, it ensures that minimal area is taken on the electronic board.

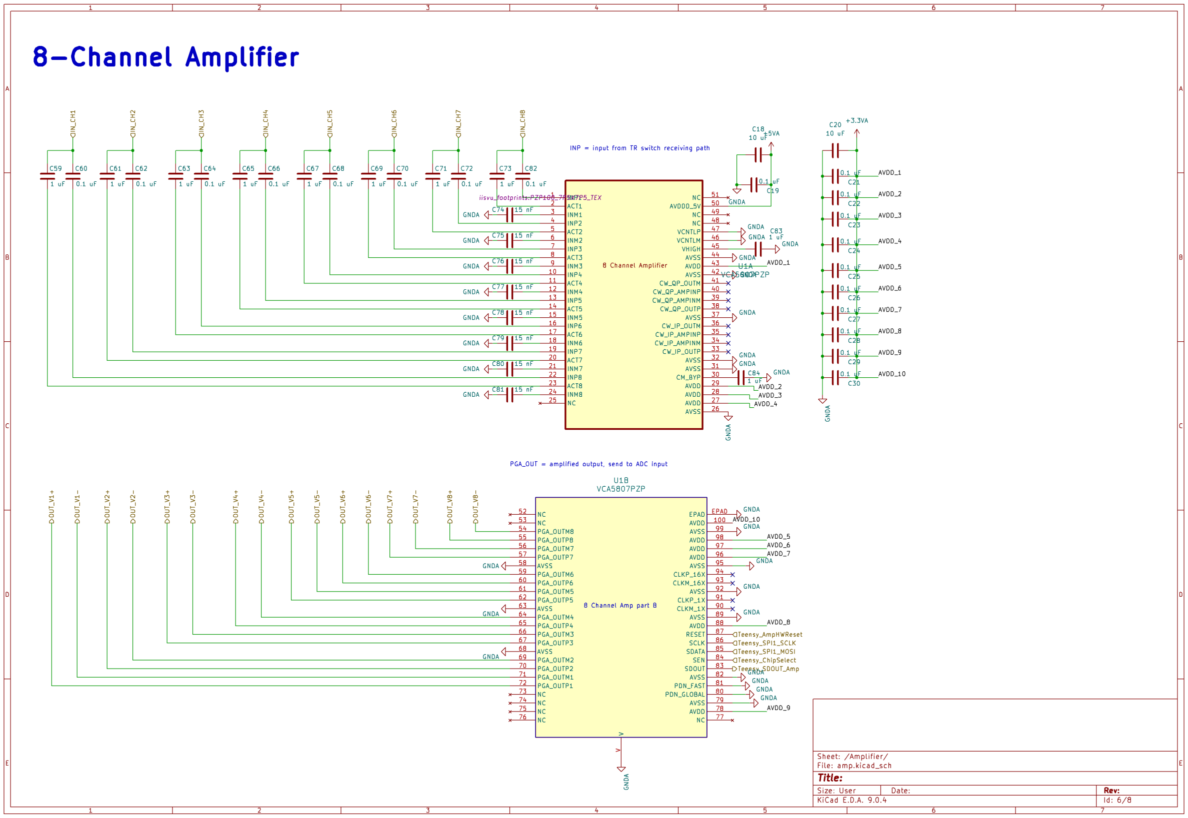

Figure 6: Amplifier Schematic with 8 inputs

The AD7606C is an 8 channel simultaneous sampling ADC with a low-pass filter at the input side to ensure undesired high-frequency signals, such as noise, do not contaminate the data.

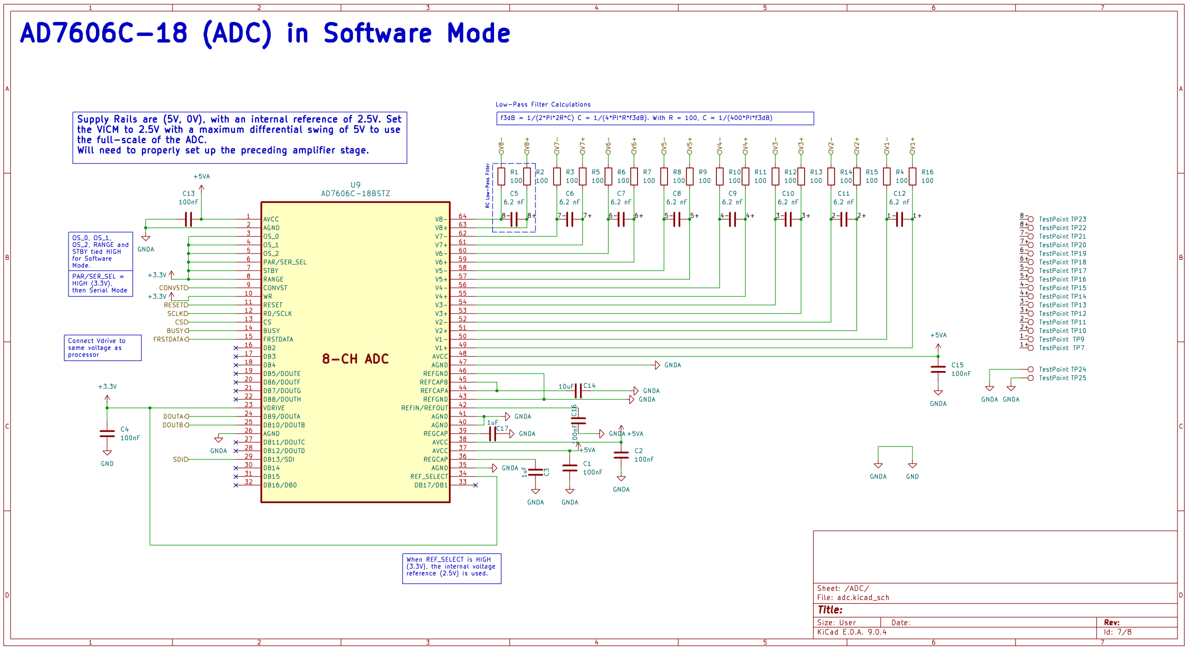

Supply Rails are (5V, 0V), with an internal reference of 2.5V. Set the VICM to 2.5V with a maximum differential swing of 5V to use the full–scale of the ADC.
Will need to properly set up the preceding amplifier stage.

f3dB = 1/(2*PI*2R*C) C = 1/(4*PI*R*f3dB). With R = 100, C = 1/(400*PI*f3dB)

OS_0, OS_1, OS_2, RANGE and STBY tied HIGH for Software Mode.

PAR/SER_SEL = HIGH (3.3V), then Serial Mode

Connect Vdrive to same voltage as processor

When REF_SELECT is HIGH (3.3V), the internal voltage reference (2.5V) is used.

Figure 7: ADC Schematic and Low-Pass Filters with 8 inputs

The TC4427 pulsers are dual-channel pulsers. Since the system requires 8 channels, 4 of these were needed. In the future we may opt for a fully integrated 8 channel pulser.

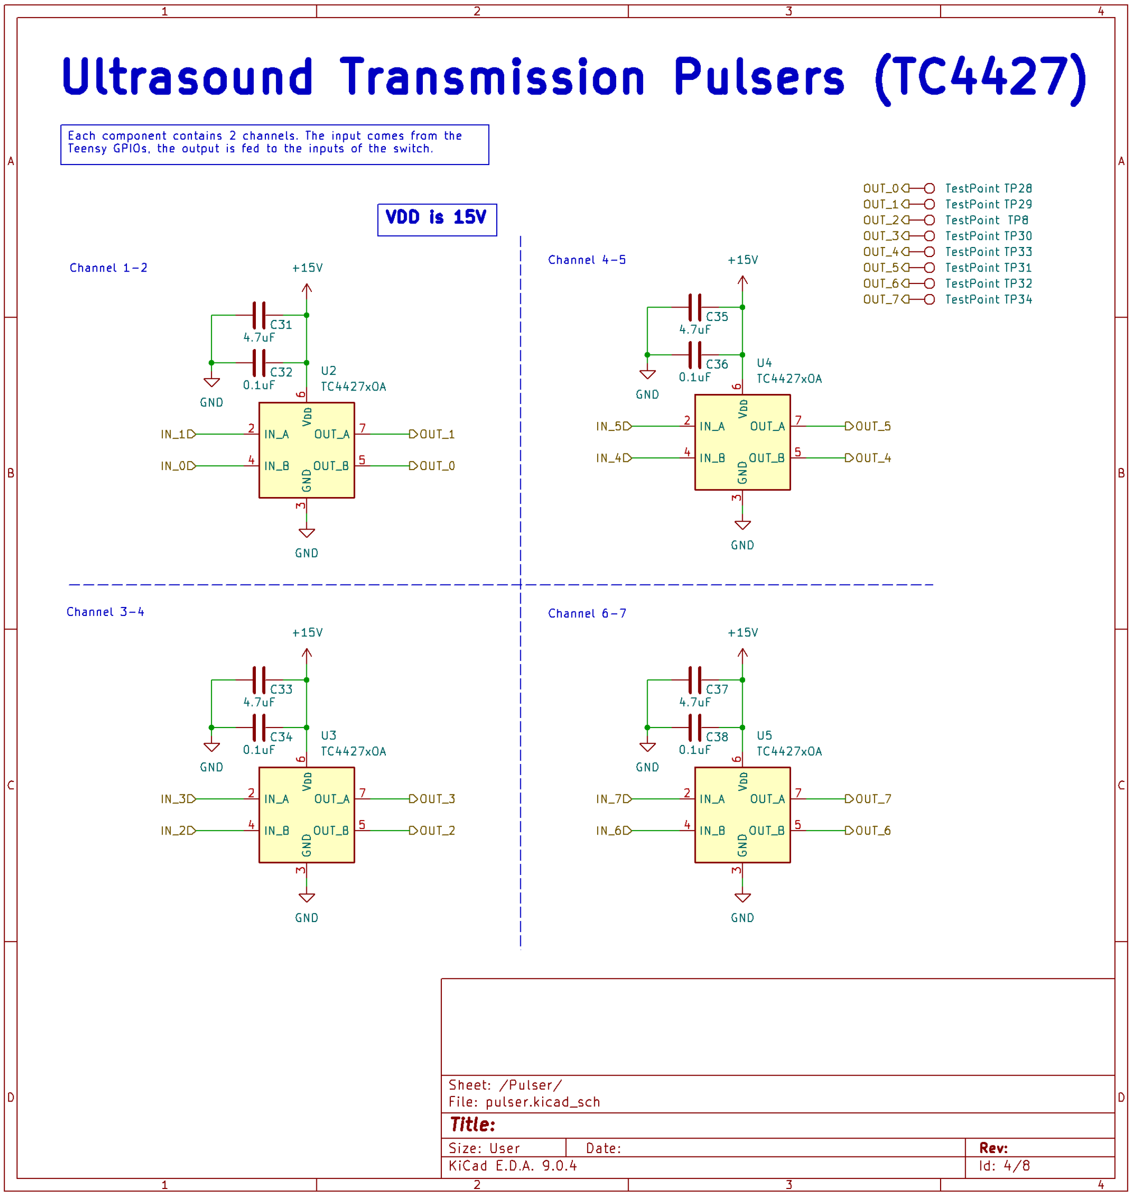

Figure 8: 8 Channel Pulsers Schematic

The HV2707-C TR Switch switches between the transmitting and receiving mode. This is a fundamental block to the system.

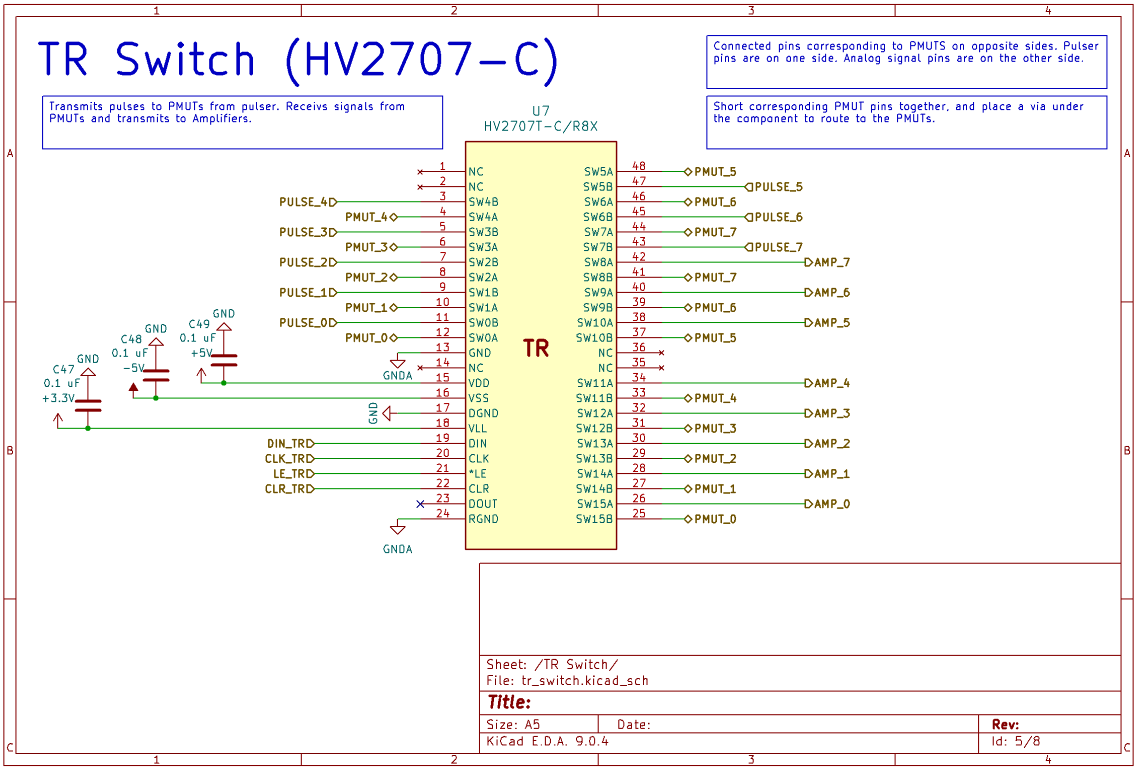
TR Switch (HV2707-C)

Transmits pulses to PMUTs from pulser. Receivs signals from PMUTs and transmits to Amplifiers.

Connected pins corresponding to PMUTS on opposite sides. Pulser pins are on one side. Analog signal pins are on the other side.

Short corresponding PMUT pins together, and place a via under the component to route to the PMUTs.

Figure 9: TR-Switch Schematic for Transmitting/Receiving

The PMUT Breakout Headers allow for the different PMUT boards to connect to the system to easily test various PMUT arrangements.

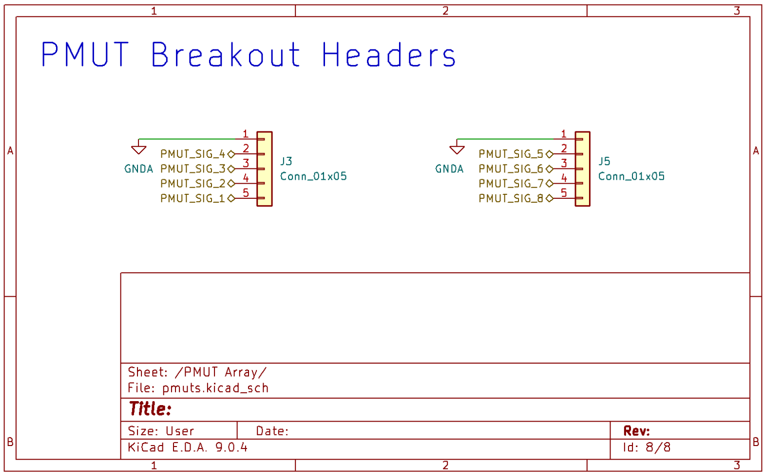

Figure 10: PMUT Breakout Headers Schematic

Finally, to power all the necessary components of this electronic system, a dedicated power management system is required. Since all components have different voltage requirements, multiple integrated circuits, such as LDOs, DC Boosters, and Inverters ensure that the correct voltage is delivered to the system.

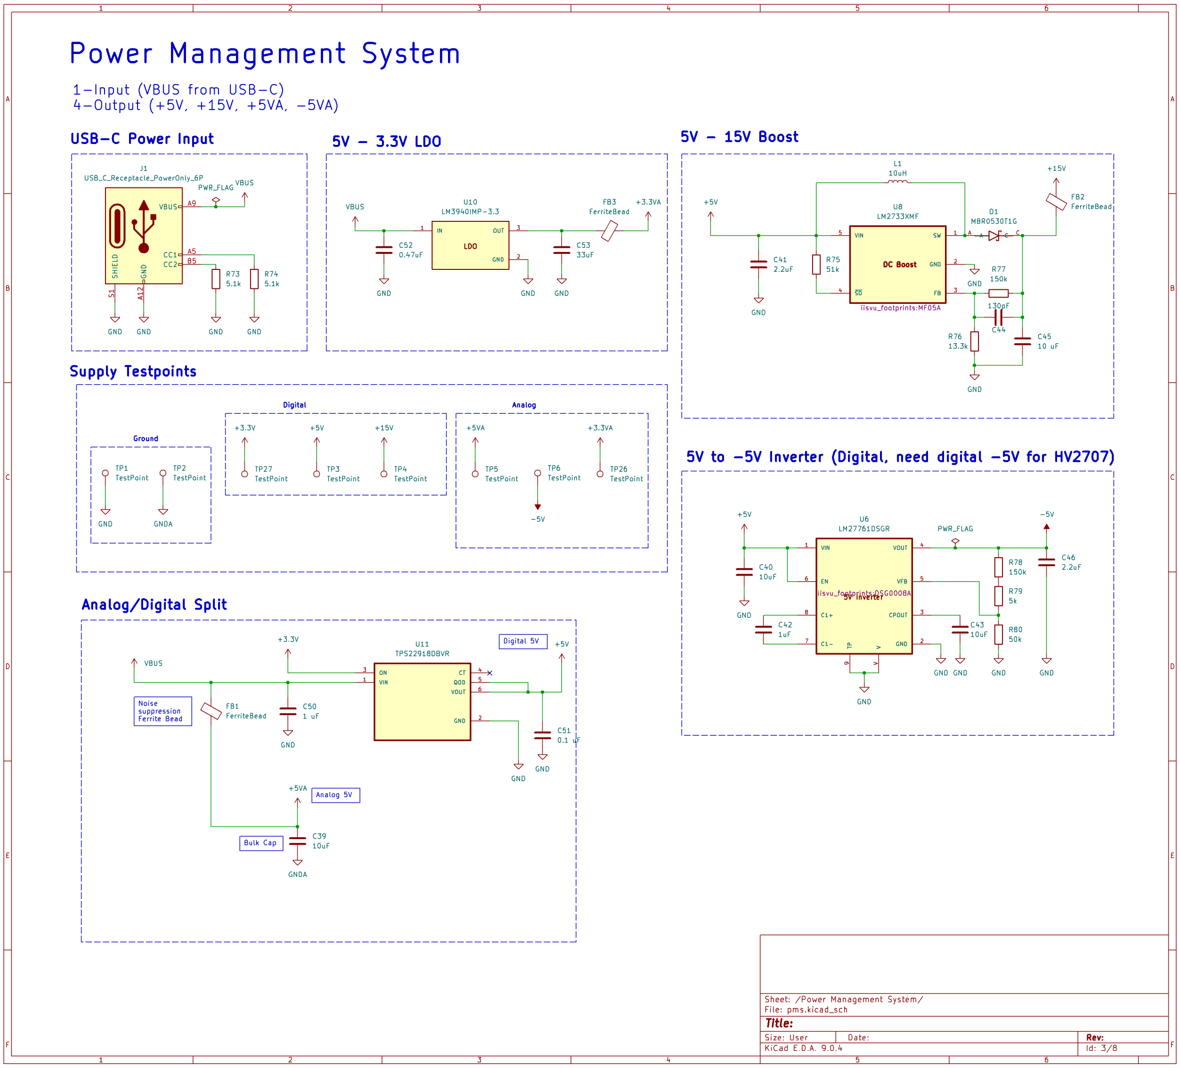

Figure 11: Power Management System Schematic

5.2.2 Electronics Components

The system uses several key integrated circuits for amplification, switching, data acquisition, and power management. The VCA5807 provides low-noise variable gain amplification for the analog front-end [6], while the AD7606C-18 performs simultaneous multi-channel analog-to-digital conversion [5]. High-voltage signal routing is handled by the HV2707 analog switch [8], and the TC4427 MOSFET driver is used to drive the transmit pulser stage [7]. Power regulation is implemented using the LM2733 boost converter and LM27761 charge-pump converter [9, 10], with the TPS22918 load switch providing controlled power sequencing [11].

5.2.3 6-Layer Printed Circuit Board Layout

The PCB layout can be seen in Fig. 12. This includes 6 layers: Top (analog signal), 2nd (analog ground), 3rd (analog power), 4th (digital power), 5th (digital ground), Bottom (digital signal). Separation in this manner is necessary to maintain signal integrity. The fast switching from digital signaling can induce EMI if the traces are not adequately separated.

Furthermore, having separate grounds for the digital and analog signals prevents noisy, high-speed digital switching currents from contaminating sensitive, low-amplitude analog signals. This way, the ADC can read clean signals.

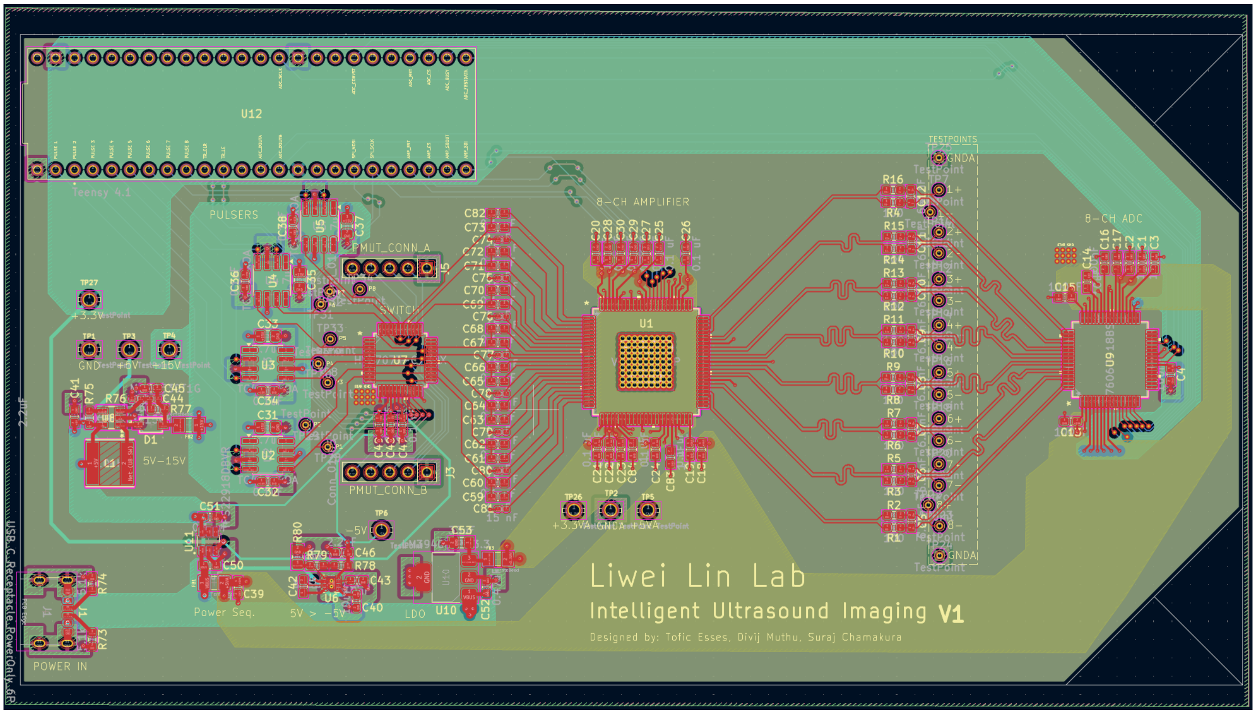

Figure 12: PCB Layout

5.2.4 Electronics Physical Realization

After the PCB was manufactured, the components were assembled onto it, represented by the physical realization seen in Fig. 13.

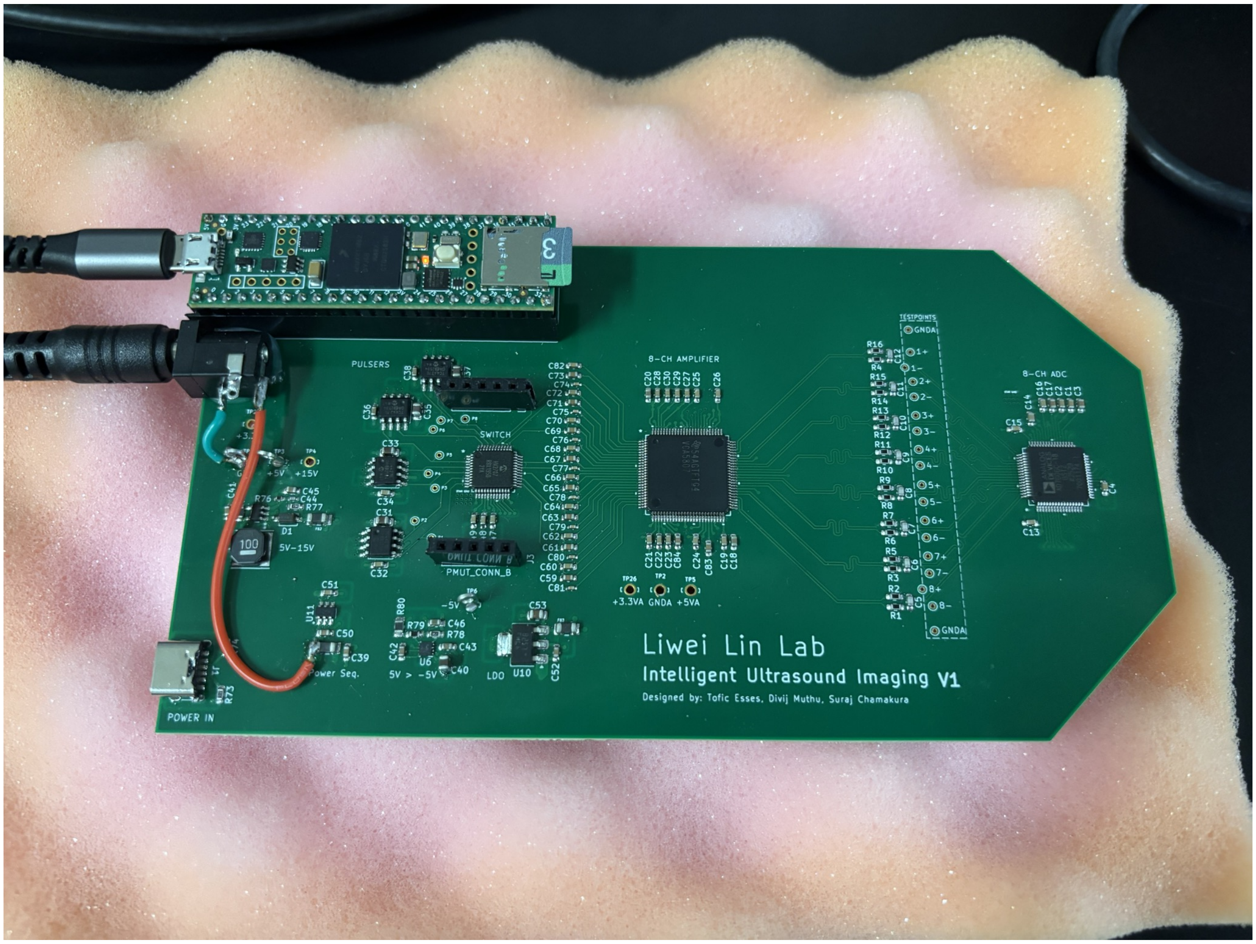

Figure 13: PCB Physical Realization

5.3 Teensy Microcontroller Firmware

The Teensy microcontroller contains the firmware code that is written to schedule the transmitting and receiving sequence for each PMUT. The received signal from each channel of the ADC is then packaged and sent to the computer for further processing. The computer is where the signal processing and reconstruction algorithm is applied.

6 Physics-Based Forward Model via K-Wave Simulation

To ensure the sensing matrix is defined by ultrasound physics the forward model is developed by using K-Wave MATLAB toolbox [12] which is a popular open-source package for acoustic simulation in time-domain based on k-space pseudospectral method [13]. The objective is to utilize this simulated sensing matrix H in the 3D image reconstruction coupled with compressed sensing algorithms. The core idea is to assume PMUTs as point source functions and develop Green's functions which are impulse responses in terms of acoustic pressure time-series at each voxel for the total volume under simulation. The final forward model when simulated with a specific set of frequency and delay combinations relates each voxel of the 3D reconstruction volume to pressure time-series measured at each receiver element as the sensing matrix used for image reconstruction.

6.0.1 Simulation Domain and Grid Configuration

The physics based simulation model is developed assuming the simulation medium to be a homogeneous medium of air with c = 343 m/s, density = 1.21 kg/m^3. The forward model assumes linear acoustic propagation in a homogeneous, weakly attenuating, and weakly dispersive medium, which permits superposition and time-invariant impulse-response modeling. The PMUTs are pulsed close to its resonance frequency of 50.4 kHz corresponding to a wavelength of approximately 6.8 mm. For the K wave simulation grid, the spatial resolution is chosen as 9 points per wavelength to maintain numerical accuracy while not increasing the computational cost significantly. The boundaries of the simulated volume are chosen as 10 grid cells of absorbing cells to suppress spurious reflections from the walls [14] .

The 8 PMUT array is configured as an irregular layout on a rectangular grid derived from profile optimization spanning approximately 96 mm x 39 mm centered at the coordinate origin for maximum incoherence in the H matrix. The simulated domain volume of 7.5 cm x 5.7 cm x 22 cm is chosen for K-Wave simulations enclosing all the PMUTs in the volume together with the absorbing boundary cells. Separately, the image reconstruction is performed on a coarser voxel grid over a 5 cm × 5 cm × 20 cm field of view (FOV) with lower end of the depth chosen as PMUT origin in z-axis direction so only the volume above the PMUT array is reconstructed. This reconstruction grid is discretized into 17 × 13 × 101 voxel array (22,321 total voxels) with spacings of 3.1 mm(x), 4.2 mm(y) and 1.7 mm(z) - the depth spacing being approximately wavelength/4, satisfying the Nyquist criterion for depth resolution. Green's functions are first developed on a finer K-Wave grid which is then mapped to the reconstruction voxel centers by spatial interpolation before assembling H. The duration of simulation is set dynamically to capture the full round-trip propagation, accounting for maximum PMUT-to-voxel distance of 200 mm, encoded frequency delays of up to 300 µs and pulse burst lengths of up to 300 µs.

For K-Wave physics simulation, rather than running a separate simulation for all 64 transmitter-receiver, we exploit the acoustic reciprocity principle for a homogeneous medium. This reduces the required simulations to 8 total base simulations, as the Green's function from element i to any voxel can be assumed to be identical to Green's function from the voxel back to element i [15].

For a given specific transmit waveforms with specific frequency and delay, the round-trip response for any transmitter-voxel-receiver path is calculated via convolution of transmit waveform with corresponding transmitter-to-voxel and voxel-to-receiver Green's functions.

The resulting H matrix from the above calculations represents the time-domain measurement at all receivers for a hypothetical unit reflector at each of the reconstruction voxels each representing one column. This H matrix is then utilized in the FISTA sparse reconstruction algorithm described in the following sections.

7 Resonant Frequency Verification with Microphone

Using a B&K Type 2669 microphone with a preamplifier, the emitted pulses of the PMUTs were recorded with a Pico Technology Type 4224A oscilloscope. As seen in Fig.14, this enables the characterization of each individual PMUT to determine their unique resonant frequencies resulting from manufacturing differences.

The microphone was placed 48 mm directly above each PMUT. The PMUT was driven by the oscilloscope, which swept through 40–70 kHz in steps of 2 kHz. The emitted signal was a 10-cycle square wave with an amplitude of 15 V_{pp}. While the signal was emitted for each frequency, the microphone swept through a 200 x 200 x 200 mm volume to characterize the resulting sound lobe across space.

To determine the peak driving frequency, the maximum amplitude and rise time of the measured signal were calculated for each PMUT. Once the individual profiles are determined, a superposition of the sound profiles can be summed together, and their resulting pattern is seen in Figure 14b. This is the synthetic aperture, and by changing the frequencies and delays of each PMUT, we can change the aperture each time to minimize the coherence between each acquisition and maximize the amount of information available for the reconstruction software to image the scene.

Table 1: Measured resonant frequencies of individual PMUT elements.

PMUT ID #	**Resonant Frequency [kHz]**
2	52
3	50
4	54
5	52
10	58

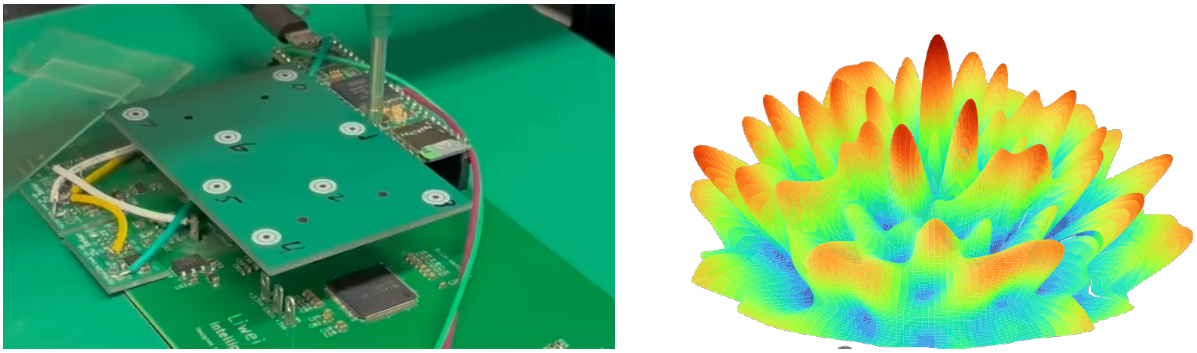
(a) Microphone Test Set Up (b) Superposition of Characterized Sound Pulses

Figure 14: Measured Emitted Sound Pulses From PMUTs.

These differences show that each PMUT affects the resulting maximum pressure emitted. A higher pressure output means there is more energy and therefore a larger amplitude in the received echo. A higher signal-to-noise ratio is required for efficient reconstruction of the measured target.

This also indicates that if the PMUTs are driven at uniform or random frequency profiles with unique delays, some profile masks may be less effective due to variations in resonant frequency. Additionally, there is inherent mechanical leakage from one PMUT being driven at its resonant frequency to another, as they are mounted on the same PCB without any vibration isolation. This means the intended profile emission from the Teensy command will be slightly different in reality.

8 PMUT Array and Transmission Profile Optimization

The ultrasound system must generate measurements that are sufficiently diverse across acquisitions for compressed sensing with a limited number of sensors. In this part of the project, we explored how a small number of PMUT elements can be strategically configured and driven to maximize unique information in the received echoes. We approached this as a hardware-software co-design problem. Both the physical layout of the PMUT array, and the transmission profiles are optimized to improve the reconstruction quality.

Recent research has demonstrated that compressed-sensing-based 3D ultrasound imaging can be achieved by leveraging measurement diversity [4]. Inspired by these approaches, we explore how a small number of PMUT elements can be strategically configured and driven to maximize information in the received echoes.

8.1 PMUT Array Layout

The PMUT array consists of eight sensors mounted on a single printed circuit board (PCB). The usable aperture for sensor placement is defined by the PCB dimensions ($75\,\text{mm} \times 50\,\text{mm}$). To ensure manufacturability and avoid interference between sensors, the optimization enforces two geometric constraints:

- A minimum spacing of 3 mm between any pair of sensors
- A border margin of 5 mm from each edge of the PCB

All candidate PMUT positions must satisfy these physical constraints, ensuring that every layout generated by the optimization algorithm corresponds to a configuration that can be physically fabricated on the PCB.

An example of an optimized PMUT array layout is shown in Fig. 15. The resulting configuration satisfies all geometric constraints while distributing sensors across the aperture to maximize spatial diversity.

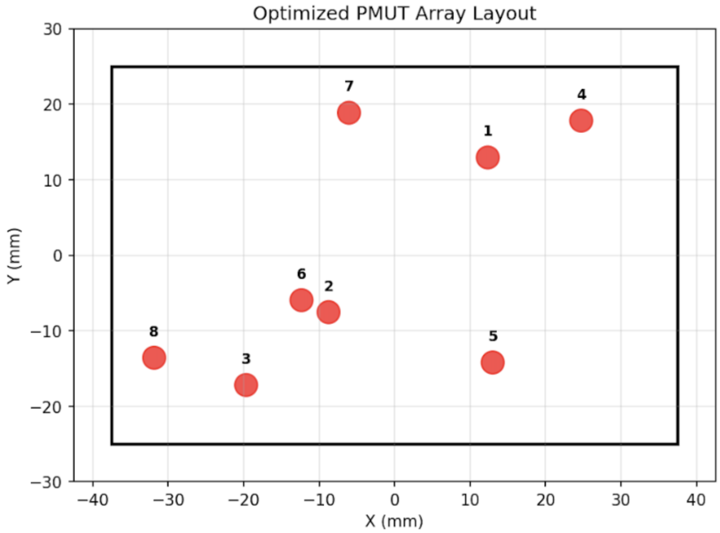

Figure 15: Optimized PMUT array layout on a 75 mm × 50 mm PCB. The black rectangle indicates the allowable placement region. Red markers denote the positions of the eight PMUT elements selected by the optimization algorithm.

8.2 Forward Model and Sensing Matrix Construction

Candidate configurations are evaluated using a simplified ray tracing forward model. The model simulates the echo generated by point scatterers within the imaging volume. Configurations are first evaluated on a coarse $12 \times 12 \times 12$ voxel grid. Then they are validated on a high-resolution $40 \times 40 \times 40$ grid to balance computational cost and accuracy. The simulation responses are assembled into a sensing matrix, which is used to evaluate the sensing performance of each candidate PMUT configuration.

8.3 Optimization Algorithm

The sensing quality of a configuration is evaluated using a mutual coherence metric. This metric measures the similarity between columns of the sensing matrix. Lower metric indicates that echoes from different locations produce more distinguishable measurements. The metric is computed using the Gram matrix $G = A^H A$. Minimizing these correlations reduces the redundancy in each measurement and therefore improves reconstruction quality. The matrix operations are accelerated using GPU computation via CuPy when available.

In addition to geometric diversity, measurement diversity is also increased by optimizing transmission frequency and delay. The algorithm searches frequencies within range of 40–60 kHz in 2 kHz increments and delays within range of 0–200 μs in 10 μs increments. L-BFGS-B optimization method with five random initializations is used to reduce sensitivity to local minima. To ensure compatibility with the system's Teensy controller, optimized parameters are projected onto a discrete grid. This ensures that optimized transmission profiles can be directly implemented on the hardware.

9 Simulation Environment Development to Test Reconstruction Algorithm

Prior to the capture of real data, a simulated $1m^3$ environment was developed to test the ability of the initial reconstruction algorithm to work on various targets. This digital twin of the real world allows quick and effective testing not only various targets, but also the parameters used in the algorithm to tune the results to become more clear. A drawback of this simulation is that the waveforms sent to the targets and simulated echoes are perfect, without any induced noise from the actual hardware system or environment.

Table 2: Tuned Parameters of Simulated Targets for FISTA Reconstruction.

Tuned Parameter	Cross	Cal	TI	AD
Acquisitions	21	24	33	24
Lamba L1	0.16	0.05	0.05	1.9
Iterations	206	183	183	274
Attenuation	1.02	0.82	0.82	0.49
Support Size	123	247	247	172
L1 Continuation	1.0	1.8	1.8	1.7

The parameters in Table 2 were found from a sweep through a range of values, whose resultant reconstruction was scored based on their accuracy and precision of the reconstructed voxels compared to the ground truth. The most critical parameters of number of acquisitions used and number of FISTA iterations taken, give hope for a possibility that there is a generalizable set of parameters for the solver to use for any target. This is contrasted with the wide range of optimized L1 values which are a penalty regularization term in the FISTA solver. It is clear in Fig.16 and Fig.17 that further tuning of the FISTA solver is required for a perfect reconstruction of any target scene.

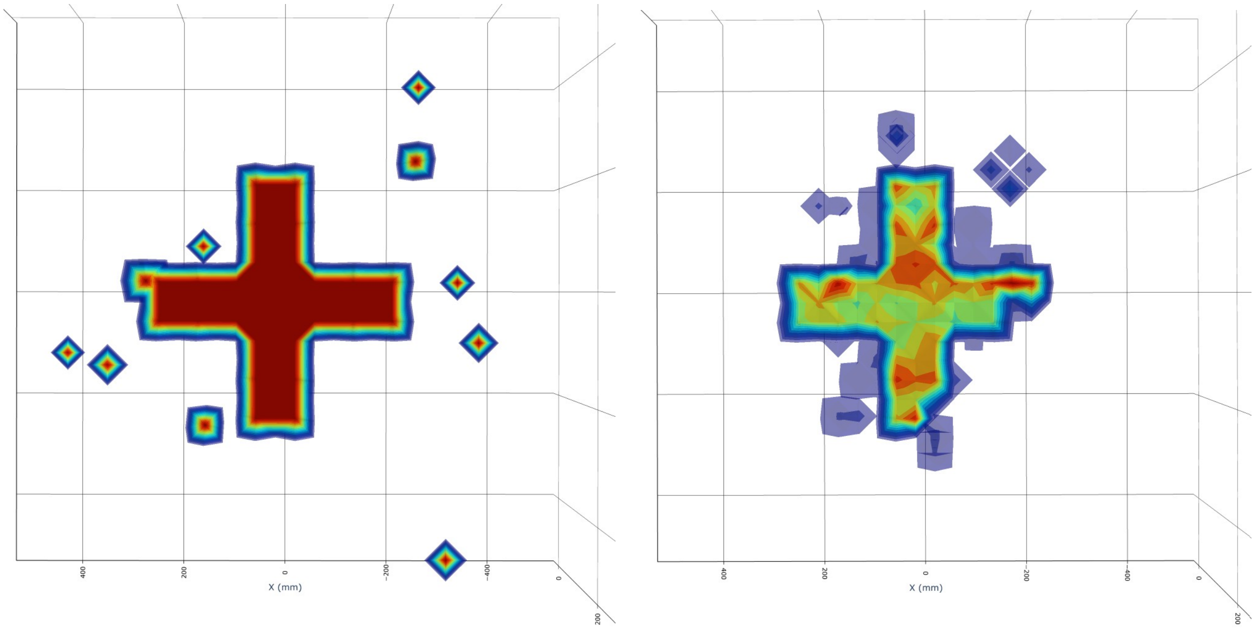

(a) Simulated Cross Ground Truth (b) Best Simulated Cross Reconstruction FISTA

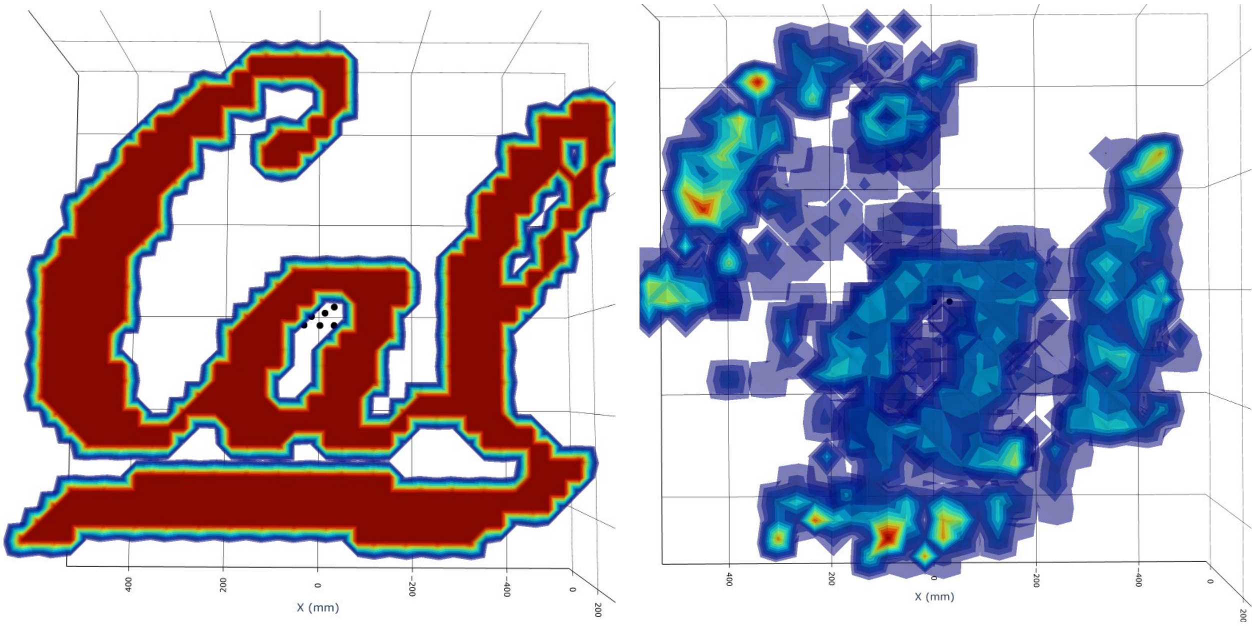

(c) Simulated Cal Ground Truth (d) Best Simulated Cal Reconstruction FISTA

Figure 16: Current Best Reconstructions of Different Targets with Varying FISTA Parameters: Cross and Cal.

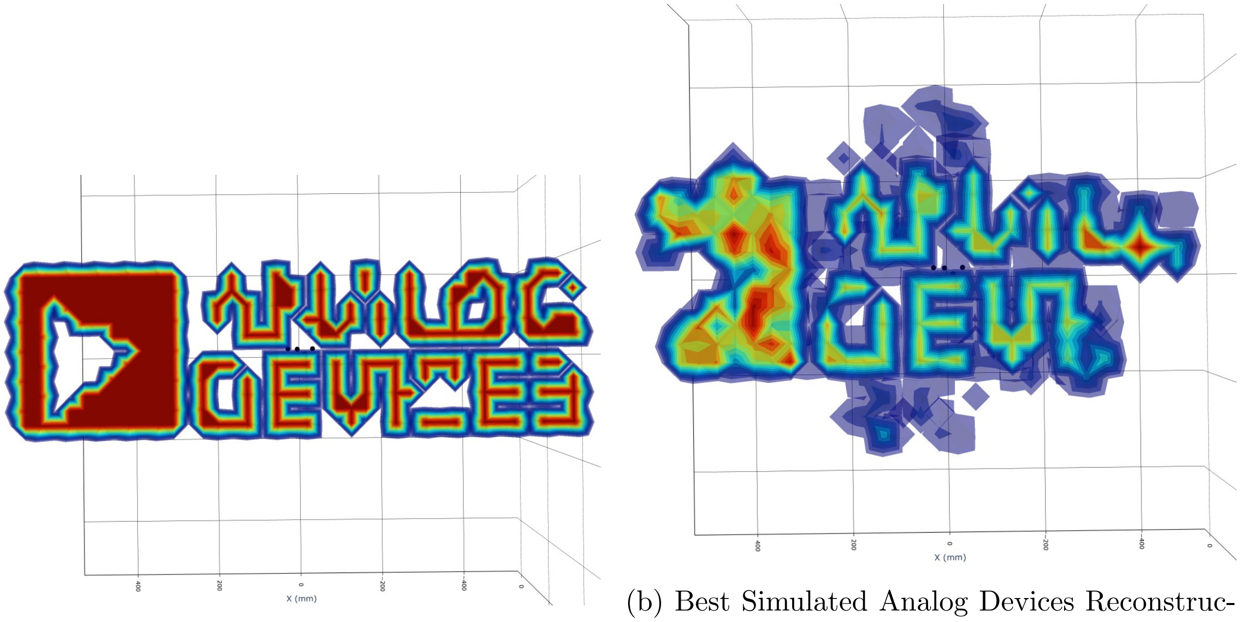

(a) Simulated Analog Devices Ground Truth

(b) Best Simulated Analog Devices Reconstruction FISTA

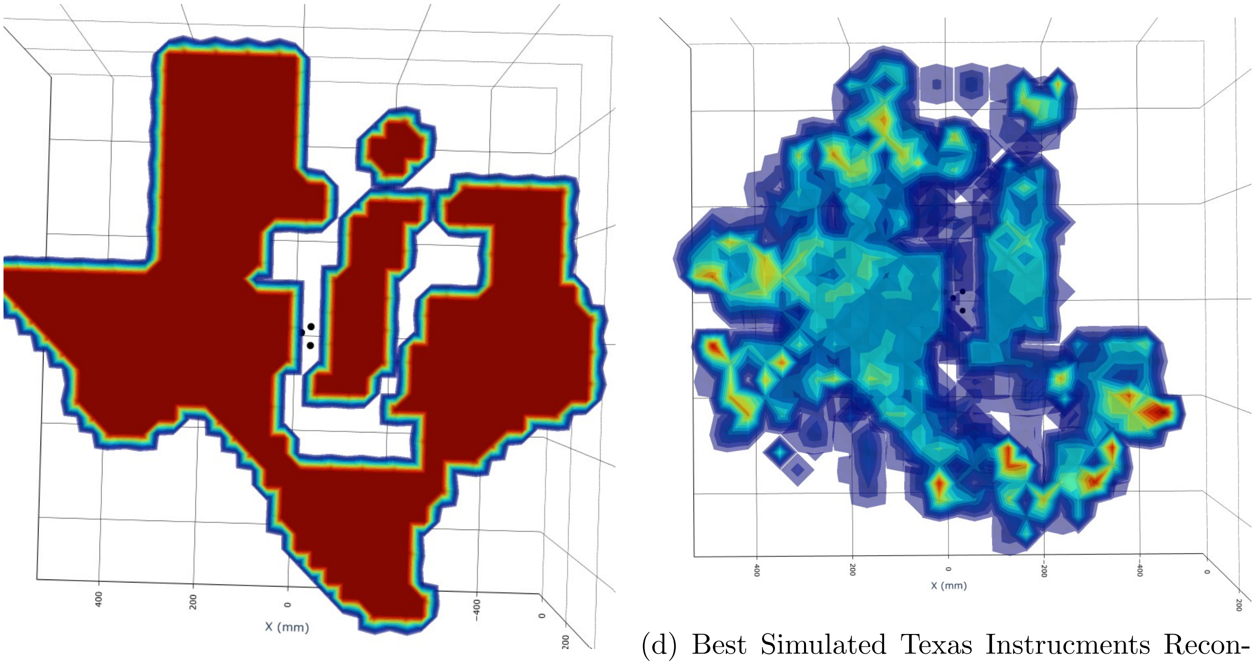

(c) Simulated Texas Instruments Ground Truth

(d) Best Simulated Texas Instrucments Reconstruction FISTA

Figure 17: Current Best Reconstructions of Different Targets with Varying FISTA Parameters: Analog Devices and Texas Instruments.

With the above results, a question arose on if FISTA was the best sparse solver that existed for our solution, and if others would produce better results.

10 Comparison of Various Sparse Reconstruction Algorithms on Performance

Several sparse reconstruction algorithms were investigated to determine which gave the highest accuracy, precision, and fast reconstruction of the target image. FISTA is a fast iterative shrinkage-thresholding algorithm developed for image processing. It is classically used for compressed sensing [16] Numerous proposals have been made in recent years on improvements to the original FISTA algorithm, [17] focused on improving the thresholding of the L1 regularization, or the gradient descent search.

ADMM is the Alternating Direction Method of Multipliers [18], which differs from FISTA in its ability to handle large matrix inversions and multiple constraints at the same time. This increased matrix size is a penalty in solve time that can be addressed with a GPU with sufficient RAM.

SBL is a Sparse Bayesian Learning designed to reward initial results in a greedy feedback, so even if the signal has low amplitude, it will be reinforced with subsequent iterations. The benefit of this is that for voxels further away from the sensor, which suffer from high attenuation, they will not be prematurely zeroed by an aggressive regularization parameter like seen in FISTA and ADMM. However, because the SBL is not aggressively selective, this means that more data has to be stored for each operation. For example, on real data sets, a T100 GPU on Google Colab was required to handle a 25 GB H matrix inversion.

11 Reconstruction Results of Three State of the Art Solvers

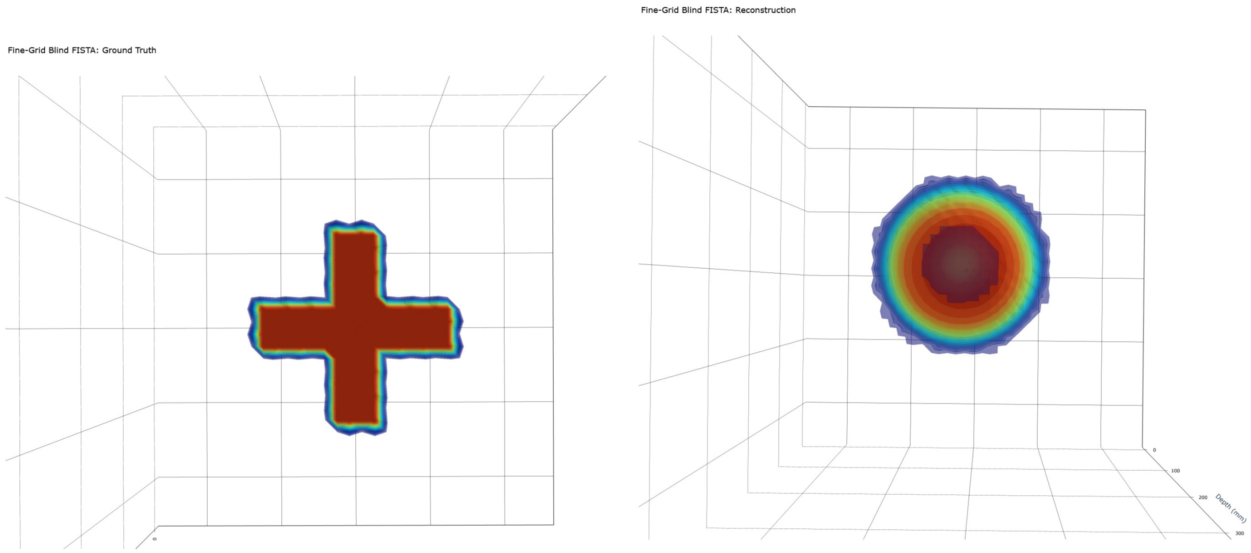

(a) Cross Ground Truth

(b) Best Reconstruction FISTA

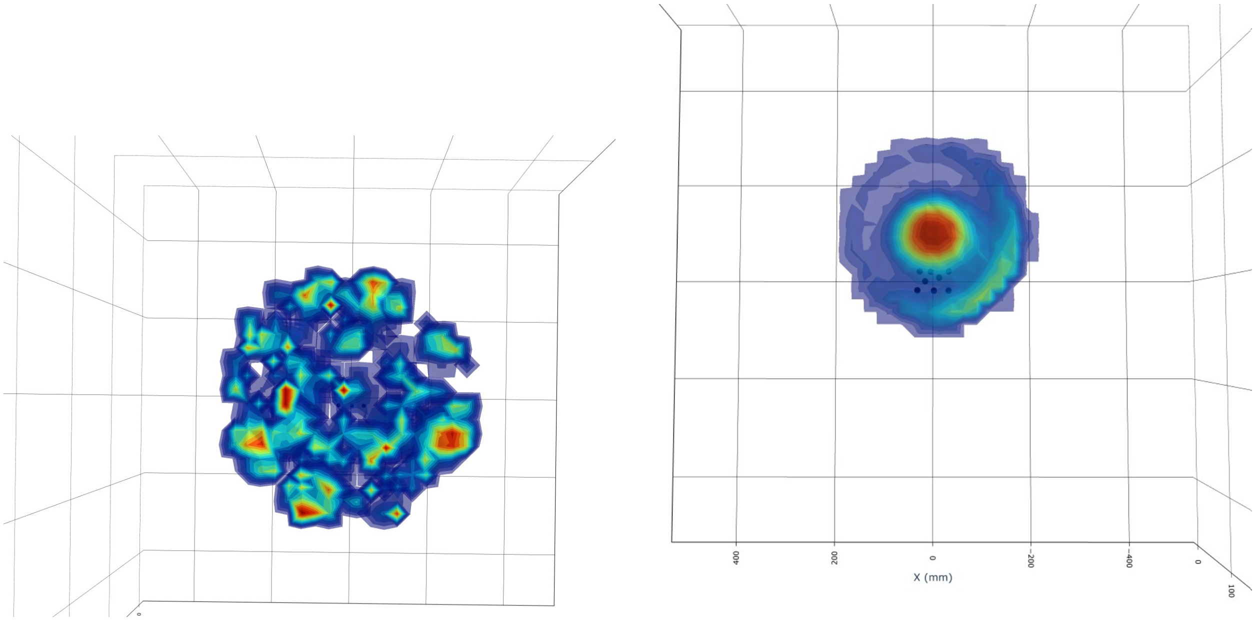

(c) Best Reconstruction SBL

(d) Best Reconstruction ADMM

Figure 18: Current Best Reconstructions of Different Solvers.

12 Limitations of the Algorithms and Current Codebase

The three different reconstruction algorithms all converged to similar results no matter their initial start point and amount of tuning allowed to parameters. Both one at a time parameter sweep, and Design of Experiments sweep were performed. These tests determined the sensitivity each parameter had on the quality of reconstruction, or if there was coupling between parameters. To diagnose the issue, an oracle model was developed, that would feed the reconstruction algorithm only the voxels within the ground truth of the target. This allows the validity of the forward model to be tested, and establishes an upper ceiling of performance for the specific parameters given. The Oracle represents the best case solution, equivalent to a perfect waveform in the simulations from Figure 16. If the Oracle has a poor performance, that means even with an idealized model, the parameters given to the FISTA, ADMM, or SBL algorithm will result in a bad reconstruction. However, if the Oracle perfectly reconstructs the target, then that means those parameters could allow the solver to successfully reconstruct the target.

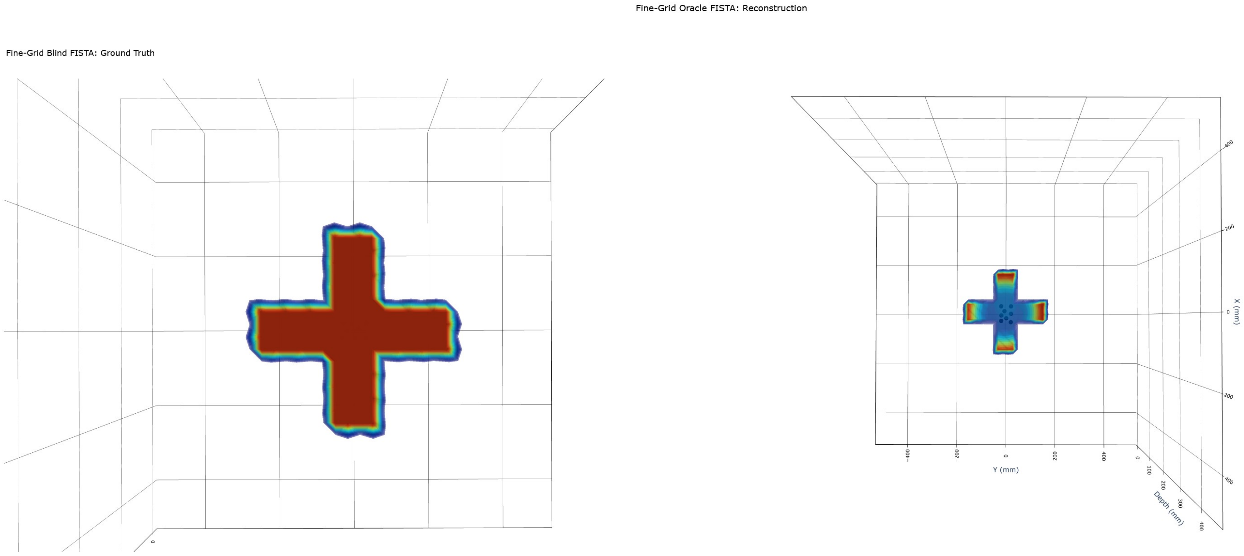

(a) Cross Ground Truth (b) Best Oracle FISTA

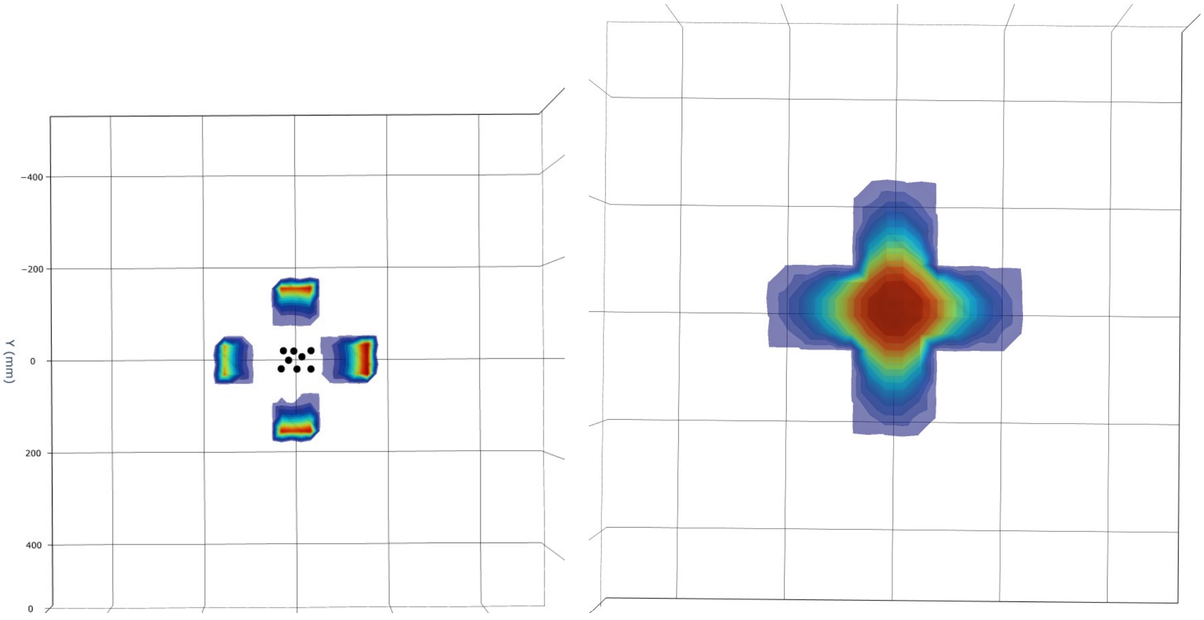

(c) Best Oracle SBL (d) Best Oracle ADMM

Figure 19: Oracle reconstructions for FISTA, SBL, and ADMM on the Cross target.

It is clear that when properly tuned, and given the right selection of voxels, both the FISTA and ADMM solver can properly reconstruct the target scene. Poorly tuned parameters cause the cross shape to devolve into unrecognizable shapes or sub-components. This shows that there is a missing component in the reconstruction pipeline as seen in the poor reconstructions of Fig.18. The solver is highlighting too many voxels as real sources of echos, resulting in an errant solution. At this state, it is clear SBL is not the right algorithm as it has both high compute cost and the worst oracle performance.

13 Conclusion

The miniaturization of the hardware components from a laboratory bench-top in Figure 2 to a single integrated PCB component is a significant leap towards a commercializable product. This PCB also simplified the ease of sending and capturing signals for the reconstruction. The reduction of all the numerous wires and interconnects to a power and data cable allowed the unit to be easily portable and set up in multiple environments. Fitting within the palm of your hand, the PCB can be easily mounted anywhere, and encased in a ruggedized container for multiple environments. There is also high potential for scaling the PCB to have additional sensors and data channels without significant redesign cost.

With the present hardware configuration, there is a limit on the amount of information available in the transmission and received echos by the PMUTs. The encoded information becomes diffuse and obscured within the signal of the overall system. Individual voxels are not able to be distinguished from their neighbors, resulting in a spherical sound cloud shape that is only stretched and distorted by the parameters, rather than revealing the target object within the signal. Compressed Sensing for Ultrasound Imaging using PMUTs, as presented in this paper with the current hardware configuration and software architecture, is not able to reconstruct any target image with real world data. The answer to successful target reconstruction is likely in the signal processing method of the measured echos. Investigations of further signal processing need to be performed on the measured waveforms, to better improve the signal to noise ratio. This will allow distinct voxels to be extracted from the various measured echoes, and allow the reconstruction software to succeed. In parallel, the simulation of pulse and echos needs to better account for real world physics, so that reconstruction strategies and PMUT array layouts are accounting for real world errors.

Pursuing multiple reconstruction algorithms, it allowed a head to head comparison for solver speed, accuracy, and tuning parameters for a generalizable scene. Under an ideal physics simulation from Figure 19, it is clear that FISTA and ADMM are avenues still worth pursuing. SBL was too costly in every metric that mattered, and still did not deliver a realistic result. At this point, ADMM is the most promising as it is already used for machine learning processes. An option for this project to take is implementing machine learning on the measured pulses from a given target, to enhance the reconstruction result from a single pulse, rather than 100 acquisitions, improving the reconstruction speed and accuracy even further. This machine learning would require significant investment on multiple training targets and environments.

14 Future Work

Once signal to noise ratio improves, and the solver produces realistic reconstructions, optimizing the throughput and latency of the algorithm will be the next step to enable real-time acquisition and reconstruction. This can be achieved by developing custom hardware on a Field Programmable Gate Array (FPGA) and much later an Application Specific Integrated Circuit (ASIC) containing all the analog and digital signal processing as well as the compute. With a solved reconstruction algorithm, and accelerated hardware, the compressed sensing ultrasound will allow imaging with fewer sensors. It opens up a new field of competition to existing ultrasound beamforming and optical imaging.

At this size, the consumer product of the PMUT arrays can be configurable to multiple different environments. The initial prototype draws constant power and transfers data via USB-C, however it is possible to integrate an on-board battery power supply with wireless transmission, to enable remote sensing in cordless locations. While the current set up is for in-air applications, swapping the 50kHz range PMUTs for ones in the MHz range allows for underwater and biomedical imaging applications. The reconstruction solver works better when there is an initial idea of what is in the static scene it is capturing. Pursuing this line of thinking, the device could be optimized for fault detection, for example healthy versus unhealthy organs. There is still significant work to require the solver to complete for real-time imaging of optical sensors. However, one field of application that it will have an entirely new niche and be unrivaled in is non-line of sight imaging.

15 Acknowledgements

We would like to thank Nikita Lukhanin for his unwavering support, guidance, mentorship, and contribution to this project.

We would also like to thank Professor Liwei Lin for providing advice and the necessary resources and lab equipment for making this project possible.

Additionally, Mostafa Sedky has been key in providing us with an accurate physics-based forward model via K-Wave Simulation.

And lastly, the contributions from Divij Muthu and Ziv Behar in all aspects of the project have also been key in advancing this project forward.

References

[1] A. Basarab, D. Kouamé, and J.-Y. Tourneret, “Compressive sensing in medical ultrasound,” *HAL Open Science*, 2013. Available: https://hal.science/hal-00830731v1/document

[2] J. Mitrović, E. S. Ebbini, and others, “Compressed sensing for reduced hardware footprint in ultrasound imaging,” *Ultrasonics*, vol. 107, 2020. Available: https://www.sciencedirect.com/science/article/abs/pii/S0041624X20301530

[3] S. M. M. Rahman et al., “Application of compressive sensing to ultrasound images,” *Biomedical Signal Processing and Control*, 2019. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC6885152/

[4] P. Kruizinga, P. van der Meulen, A. Fedjajevs, F. Mastik, G. Springeling, N. de Jong, J. G. Bosch, and G. Leus, “Compressive 3D ultrasound imaging using a single sensor,” *Science Advances*, vol. 3, no. 12, 2017, doi: 10.1126/sciadv.1701423.

[5] Analog Devices, *AD7606C-18: 18-Bit, 1 MSPS, 8-Channel Simultaneous Sampling DAS*, Datasheet, 2023. Available: https://www.analog.com

[6] Texas Instruments, *VCA5807: 8-Channel Variable Gain Amplifier with Low-Noise Preamplifier*, Datasheet, 2022. Available: https://www.ti.com

[7] Microchip Technology, *TC4427: 1.5A Dual High-Speed Power MOSFET Driver*, Datasheet, 2021. Available: https://www.microchip.com

[8] Microchip Technology, *HV2707: 8-Channel High Voltage Analog Switch*, Datasheet, 2020.

[9] Texas Instruments, *LM2733: High Frequency Boost Converter*, Datasheet, 2022. Available: https://www.ti.com

[10] Texas Instruments, *LM27761: Switched-Capacitor Voltage Converter with LDO*, Datasheet, 2021. Available: https://www.ti.com

[11] Texas Instruments, *TPS22918: Ultra-Low Leakage Load Switch*, Datasheet, 2022. Available: https://www.ti.com

[12] B. E. Treeby and B. T. Cox, “k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields,” *J. Biomed. Opt.*, vol. 15, no. 2, p. 021314, 2010.

[13] B. E. Treeby, J. Jaros, A. P. Rendell, and B. T. Cox, “Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method,” *J. Acoust. Soc. Am.*, vol. 131, no. 6, pp. 4324–4336, 2012.

[14] J. P. Berenger, “A perfectly matched layer for the absorption of electromagnetic waves,” *J. Comput. Phys.*, vol. 114, no. 2, pp. 185–200, 1994.

[15] D. A. Steinberg, *Principles of Aperture and Array System Design.* New York, NY, USA: Wiley, 1976.

[16] D. L. Donoho, ”Compressed sensing,” in IEEE Transactions on Information Theory, vol. 52, no. 4, pp. 1289-1306, April 2006, doi: 10.1109/TIT.2006.871582.

[17] Wei, Jianxiang et al. “A Faster and More Accurate Iterative Threshold Algorithm for Signal Reconstruction in Compressed Sensing.” Sensors (Basel, Switzerland) vol. 22,11 4218. 1 Jun. 2022, doi:10.3390/s22114218

[18] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. Found. Trends Mach. Learn. 3, 1 (January 2011), 1–122. https://doi.org/10.1561/2200000016