

A systematic evaluation of hardware and acquisition parameters for near-infrared imaging

Vivian C. Yue

Mentors: Anuradha Godavarty and Lihong Wang

Near-infrared (NIR) imaging allows light to penetrate several millimeters to centimeters into tissue due to reduced absorption and scattering in the 650–1000 nm spectrum, which allows for noninvasive assessment of subsurface physiological information. The objective of this project is to optimize near-infrared (NIR) imaging by identifying the imaging hardware, acquisition parameters, and illumination methods that produce the most consistent high quality images. Camera systems were evaluated under white light and multiple NIR wavelengths (660, 690, 830, and 850 nm) by comparing image quality using different image formats, shutter modes, and downsampling methods during both static and continuous imaging. The optimal camera settings and the selection of cameras that provide the best performance for NIR imaging have been identified through this process. Additional work included evaluating optical filters, alternative LED light sources, compact white-light cameras, and time-of-flight sensors, while developing synchronized imaging software capable of controlling multiplexed illumination and image acquisition. Custom illumination boards and integrated control electronics have also been developed and are awaiting experimental validation. Future work will focus on achieving stability in the illumination with replacement light sources, and integrating the optimized hardware and software into a standardized imaging protocol that provides clean and efficient near infrared imaging results.

Toward event-driven feature extraction for neuromorphic brain-machine interfaces

Layla Adeli

Mentors: Azita Emami and Loai Danial

Higher channel counts under lower power budgets and unstable decoding due to electrode degradation have motivated the field of neural interfaces to move towards sparse feature extraction. This project aims to develop a baseline for such an event-driven feature extractor, using an analog front end (AFE) to encode neural activity as pulse-density events. The AFE is used to replace the uniformly-sampled converter and waveform reconstruction that most interfaces require, and is the natural input to later feature extraction and decoding. The work was structured in four stages: a comparison of seven unsupervised spike sorters under simulated electrode aging, the construction of a framework for comparing spiking neuron models on the AFE's actual event stream, an evaluation of those neuron models as feature extractors for motor decoding, and the hardware implementation of the logic that could turn the AFE's pulses into a decoding feature. Treating the input as already encoded by the front end, the final stage was implemented in Verilog, including an asynchronous counter of its events. This work contributes toward low-power, on-chip feature extraction for event-driven neural interfaces.

A wireless, oxygen-generating macroencapsulation device for high-density islet transplantation in type 1 diabetes

Aruesha Srivastava

Mentors: Siddharth Krishnan, Wei Gao, and Kecheng Wang

Pancreatic islet transplantation has the potential to restore endogenous insulin secretion in patients with type 1 diabetes, but long-term graft survival is limited by immune rejection and insufficient oxygen delivery. Macroencapsulation devices can protect transplanted islets from immune attack, yet their clinical scalability remains constrained by hypoxia, particularly at high cell densities and increased implant thicknesses. This project investigates a wireless, oxygen-generating macroencapsulation system designed to improve oxygen availability while maintaining a compact and retrievable implant architecture. A resonant inductive power-transfer system was developed to wirelessly drive an electrochemical oxygen generator. The receiving antenna, circuit board, and oxygen-generating element were integrated into a functional implant-side assembly. An incubator-compatible stacked transmitter and passive resonator configuration was developed to replace the NeuroLux cage used in initial testing. The optimized six-turn receiving antenna produced sufficient

voltage to operate the oxygenator, and the complete integrated system generated visible oxygen bubbles under wireless power. In parallel, human islets were mechanically downsized to approximately 70 μm to reduce oxygen diffusion distance. Preliminary live/dead fluorescence analysis after 24 hours at 1% oxygen showed a lower measured necrotic area in downsized islets than in native islets. A multilayer BioMed Clear Resin-Ecoflex-PDMS coating was also developed to reduce corrosion in culture medium while preserving oxygen generation during short-term testing. These results show the feasibility of combining wireless power delivery, electrochemical oxygen generation, islet-size reduction, and protective encapsulation within a single platform. The ongoing work focuses on biological replication and functional assessment of insulin secretion.

Wireless non-LOS BAN channel characterization and communication at mmWave

Sadik O. Görgü

Mentors: Ali Hajimiri and Zachary Szekeres

Increased data requirements in wireless communications have pushed devices towards higher frequencies of operation. One application where this is particularly important is body-area networks (BANs), which enable communication between separate on-body devices. At mmWave frequencies, on-body wireless communication is strongly affected by the body's shadowing and scattering from the surrounding environment, which can significantly reduce link reliability. In this work, we make use of these scattering effects by first characterizing the wireless channel then using the estimated channel response to establish a reliable link between on-body transmitters and receivers. Channel characterization will be performed at millisecond time scales, completing the measurement and data transmission within the channel coherence time. For this purpose, we are designing RF boards operating at 24 GHz using a 1 GHz-bandwidth FMCW chirp to estimate environmental scattering paths and their corresponding phase delays. A 3×3 phased array with digital beamforming adds flexibility to measurements and reduces channel estimation time. Following channel estimation, data transmission will be performed within the 24.0–24.25 GHz ISM band.

Advancing monolithic integration of vacuum tubes on CMOS

Richard Y. Li

Mentors: Ali Hajimiri and David A. Baum

Although solid-state transistors dominate modern electronics, vacuum tube devices still show superior power density and high frequency performance. Monolithically integrating vacuum tubes onto standard complementary metal-oxide-semiconductor (CMOS) substrates allows the benefits of vacuum tubes and transistors to be combined to increase the performance of radio frequency integrated circuits. This project addresses numerous constraints of current on-chip vacuum tubes, namely their short lifespans and limited gain. By optimizing the geometry of the vacuum tubes, both the lifespan and gain of the vacuum tube devices may be improved; thus, a variety of structures were designed to determine the optimal geometry for vacuum tubes. The measurement of these structures will also be used to improve models of vacuum tubes. Additionally, supporting digital control circuitry was developed to precisely control the temperature of the vacuum tubes to increase lifespan. These developments will further the goal of building a fully integrated vacuum tube system on standard CMOS chips.

Highly scalable atomic force microscope on an unmodified CMOS chip

Katherine C. Robertson

Mentors: Ali Hajimiri, David A. Baum, and Jatin Mathur

Atomic force microscopy (AFM) offers some of the highest-resolution imaging of surface topography among existing microscopy methods. However, existing AFM techniques are frequently limited by throughput, which inhibits their use in, for example, semiconductor manufacturing and quality control. To overcome this limitation, we present massively parallelizable arrays of integrated AFM units. The AFM units are designed to leverage unmodified CMOS processes and maskless post-processing, which enables mass-manufacturability at scale and the co-integration of supporting electronics. Each AFM unit comprises an electrothermally actuated Z-axis stage and cantilevered probe tip. The resonant

frequency of the probe tip is read capacitively and used to determine the probe's interaction with a sample surface. As probe-sample interactions occur, the resonant frequency of the beam fluctuates, which provides a measure of the probe's distance to the sample surface. Several non-contact mode AFM units and test structures were taped out on a 180 nm bulk CMOS process.

Development of thin-film tactile sensors using PVDF and P(VDF-TrFE-CFE)

Chloe Wang

Mentors: Ioannis Kyriassis, Wei Gao, Eric Chang, and Peter Ballentine

Static and dynamic tactile sensors, which mirror the slow-adapting and fast-adapting mechanoreceptors of the human tactile system, are essential for providing rich sensory information to robotic manipulators. Thin-films are especially well-suited for fabricating tactile sensors as they produce compact and lightweight devices with high sensitivity and accuracy. In this project, we develop thin-film tactile sensors based on polyvinylidene fluoride (PVDF) and poly(vinylidene fluoride-trifluoroethylene-chlorofluoroethylene) (P(VDF-TrFE-CFE)). Chromium-gold electrodes are deposited via electron-beam evaporation through shadow masks to fabricate parallel-plate capacitive sensors on each film. The devices are then characterized under static and dynamic compressive loading by measuring both capacitive and piezoelectric charge response. Charge generated by the sensors is measured using an ADC-based charge readout, while capacitance changes are measured using a CDC. Although PVDF is widely used for dynamic tactile sensing through its piezoelectric response, we investigate if capacitive measurements can also enable static tactile sensing within the same device. We further evaluate P(VDF-TrFE-CFE) for a strong static response through its high dielectric constant. Finally, we aim to separately integrate each device into a compact robotic finger and demonstrate static and dynamic signal acquisition.

Flow-DFL: Decision-focused learning for decision-making under uncertainty via energy-based conditional flow matching

Tommaso Colombo

Mentors: Steven H. Low and Han Xu

Decision-making under uncertainty often optimizes for improving prediction accuracy to improve decision quality rather than to making high quality downstream decisions. This project studies that "predict-then-optimize" gap in the setting of virtual electricity trading. The project compares a classical DPDS non-neural virtual bidding benchmark against four neural decision-focused policies. These policies each use the spread predictions generated by the prediction models to produce optimal trading decisions using either an MLP or an energy-based conditional flow matching (ECFM) decision model. Policies are then evaluated using validation total cumulative profit, Sharpe ratio, spread MSE and latency. Across 10 random seeds, the MLP + ECFM pipeline achieved the highest cumulative validation profit out of all the models, while FM-TS Direct achieved the highest average Sharpe ratio. The results suggest that flow-based models generally make better decisions when the prediction quality is the same, but also that the benefits are sensitive to model architecture, the initial seeds and the different regularization choices. Furthermore, the spread and Sharpe regularized model performed better than the non-regularized model in a single seed run. This suggests that optimizing for downstream decision making rather than prediction accuracy could be more important in making more optimal decisions under uncertainty.

Ultrashort pulses from an optoelectronically mode-locked laser on nanophotonic lithium niobate

Richard H. Feng

Mentors: Alireza Marandi, Rithvik Ramesh, and Yiting Wang

Mode-locked lasers (MLLs) are important sources in modern photonics because they produce periodic trains of ultrashort optical pulses. These pulses are useful for optical computing, time multiplexing, and driving nonlinear optical processes, while the corresponding frequencycomb enables applications in spectroscopy, precision timing, and ranging. The comb-line spacing, equal to the pulse repetition rate, provides a bridge between the optical and microwave frequency domains. Recent efforts have focused on developing integrated MLLs to enable compact, electrically driven photonic sources. Using

a hybrid cavity formed by coupling a semiconductor gain chip to a thin-film lithium niobate (TFLN) waveguide, we demonstrate an electrically pumped integrated MLL with active intracavity loss modulation provided by a tunable mirror architecture. We generate pulses as short as ~ 300 fs with a coherent frequency comb spanning ~ 23 nm centered near 1550 nm, with over 1.5 W of on-chip peak power, over 400 mW of off-chip peak power, and over 1.5 mW of off-chip average power. We further demonstrate self-starting and self-sustained operation of the MLL as a coupled optoelectronic oscillator, with the laser generating its own RF drive signal at its repetition rate, eliminating the need for an external RF signal generator. In this configuration, we maintain ultrashort (~ 330 fs) pulses while simultaneously generating a 9 GHz microwave signal. Ongoing work focuses on fully stabilizing the frequency comb to ultrastable references through injection locking and other locking schemes, toward further development of the platform for low-noise microwave synthesis.

Scaling full-3D inverse design to large-area photonic systems

Alexander P. Plekhanov

Mentors: Kiyoul Yang and Kerry J. Vahala

Inverse design uses gradient-based optimization of Maxwell's equations to discover compact nanophotonic device geometries with prescribed optical responses, inventing structures beyond conventional human design intuition. However, the computational cost of full-wave 3D simulation has generally limited designs to footprints below roughly 30×30 μm . This restricts their performance and their use in complex, highly connected photonic systems.

We developed a fabrication-aware inverse-design framework that substantially expands the accessible design area to approximately 110×110 μm on a single accelerator and 200×200 μm using distributed computation at telecommunications wavelengths. This advance enables large multiport photonic systems, rather than isolated components, to be optimized directly.

We apply the framework to SiN and SOI devices, including a wavelength-division multiplexer and a broadband 50:50 splitter. A central target is a nanophotonic neural-network accelerator implementing a 36×36 optical matrix transformation in a single structure. By performing the core linear operations of neural-network inference directly in light, such devices could provide massively parallel, high-throughput computation with reduced data movement and energy consumption, while substantially exceeding the scale of previous inverse-designed optical matrix processors.

Photonic molecule design for quantum frequency conversion

Matthew J. Cantor

Mentors: Kiyoul Yang and Kerry J. Vahala

Scalable quantum networks require efficient interfaces between various quantum processors that emit and absorb photons at different optical frequencies. This project develops an integrated quantum frequency converter linking two visible-band quantum emitters separated by a frequency shift of tens of terahertz. The device utilizes Bragg-scattering four-wave mixing in a silicon nitride photonic molecule composed of two coupled high-Q microring resonators. Two optical pumps mediate the conversion and hybridize across the coupled resonators, providing a tuning degree of freedom that satisfies energy conservation while keeping all four waves aligned to cavity resonances. The signal and idler, however, are localized in a single ring to maintain strong nonlinear overlap for high conversion efficiency at reduced pump powers. The design methodology combines electromagnetic mode simulations with coupled-mode theory, and has identified geometries that satisfy the conditions of Bragg-scattering four-wave mixing while suppressing competing nonlinear processes. The next steps are to fabricate the device and characterize its performance. This work establishes a pathway toward low-noise, chip-scale quantum interconnects capable of linking dissimilar quantum-network nodes.