Development of a Tone-Tracking Algorithm for Maximizing Dynamic Range of Kinetic Inductance Detectors

Jonathan R. Hoh[1], Adrian K. Sinclair[1], and Ryan C. Stephenson[1]

***Abstract*— Here we present the designs and preliminary results of an FPGA-based tone-tracking firmware that always maintains on-resonance probe tones in frequency multiplexed MKID arrays. The design of the firmware relies on the fact that resonators have a distinct phase shift at resonant frequency that remains invariant even as the resonance itself changes.**

Similar to proportional/derivative/integration (PID) digital controllers, the firmware compares phase feedback loops to resonant phase data collected before beginning observations. Using these phase comparisons, the tone-tracking algorithm can individually update each tone in the comb sent to an MKID array so that resonance is always known and maintained. In practice, this eliminates the need for mid-observation VNA sweeps and allows MKID-based cameras to reach their maximum dynamic range potential.

The application of this tone-tracking firmware will drastically change the approach for reading out MKID arrays. We will conclude with a discussion on how other aspects of the readout system can be modified and simplified to accommodate tone-tracking along with future NASA missions that hope to use the design. Finally, we look forward at the possibility of implementing tone-tracking firmware on the cutting-edge Xilinx RFSoC FPGA.

I. Introduction

Over the new millennium, microwave kinetic inductance detectors (MKIDs) have become a popular technology for use in mm and sub-mm astronomy. MKIDs function by cryogenically cooling certain metals at superconducting temperatures in order to create a resonator which changes in both resonant frequency and quality factor when exposed to incoming photons. By creating an array of these MKIDs, each with a unique resonant frequency, the entire system can be read out with a single transmission line using a method known as frequency multiplexing. Depending on the arrangement MKIDs can be used as pixels for both spatial and spectral deep-space cameras.

In order to frequency multiplex an MKID array, a comb of tones comprised of resonant frequencies must be sent down a single transmission line coupled to all the detectors. By reading the transmitted spectral phase and amplitude information, incident power on each of the MKID pixels can be determined. The simplicity of frequency multiplexing an MKID array is one of the major upsides of using the technology versus competing detectors such as Transition Edge Sensors (TES) bolometers.

Recently, the cutting-edge technology of MKIDs and the associated frequency multiplexing techniques have been flown on balloon missions such as OLIMPO and BLAST-TNG with reasonable success. The BLAST-TNG mission in particular managed to capture images of a far-infrared stellar source using MKID detectors before the early demise of the flight. The successful implementations of MKID technology as a means of ultra-high sensitivity detection shows great promise for the future of mm and sub-mm astronomy. Unfortunately, the growing use of MKIDs for astronomy has also illuminated certain problems with their implementation. Current readout techniques for MKIDs greatly reduce their theoretical dynamic range; an issue that needs to be remedied if KIDs are to continue their widespread use.

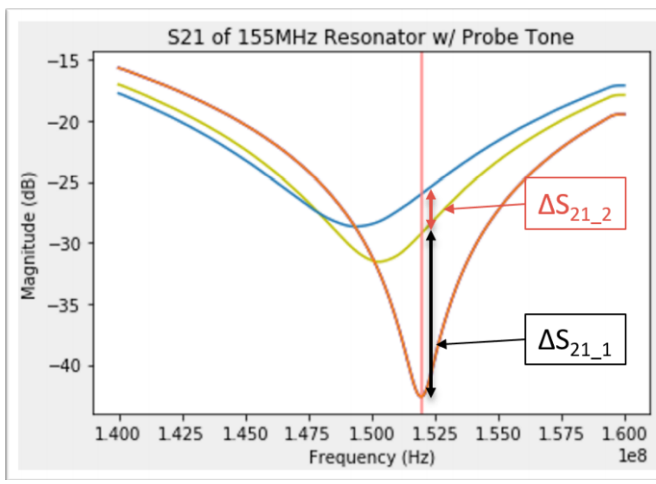

Fig. 1. An exaggerated example depicting the diminishing returns of incident power on a KID. Increasing power returns smaller and smaller changes in amplitude until the difference between two modulated states is dominated by amplifier noise. As of now, the only correction possible is a time-intensive mid-flight VNA sweep.

The current readout technique for MKID arrays as stated earlier is through frequency-domain multiplexing (FDM). By tuning each element of the array to a slightly different resonant frequency, the phase and magnitude of each detector can be individually read out from the same signal line. For small amounts of incident power, this method of monitoring has

[1] The Arizona State University, Tempe, AZ 85282 USA (contact info: jrhoh@asu.edu)

shown to be perfectly acceptable. However, problems begin to arise when there are very large changes to the incident power on the detectors.

As the intensity of light interacting with a KID increases, the change in amplitude seen at port 2 (the output to the readout FPGA) begins to shrink rapidly. Once a KID is saturated with enough power, the change in transmission is small enough that it is lost in system noise, mainly due to amplifiers. Fig. 1 shows an exaggerated example of this difference in visibility. Attempting to read out useful information from a KID which is dominated by amplifier noise is a fruitless effort and the detector is dead weight until the issue is corrected. Currently, the only remedy for this problem is a time-consuming VNA sweep which occurs mid-flight and wastes precious observation time. Losing amplitude information to amplifier noise is not uncommon during flights and on the recent BLAST-TNG mission this problem occurred every few degrees when scanning the sky.

The next NASA mission planning to use MKIDs as a means of deep-space detection is the Terahertz Intensity Mapper (TIM) balloon-borne telescope, led by a group at the University of Illinois in collaboration with a number of other research institutions. TIM aims to survey CII, OI, and OIII in 100 galaxies to demystify many aspects of star formation while simultaneously producing deep tomographic maps of the 3-D structure of the universe. The architecture, spectrometer, and readout system of TIM is modeled from the BLAST-TNG mission that recently flew at McMurdo Base in 2019. Among the lessons learned from BLAST-TNG was the need for a more efficient method to overcome the KID dynamic range issue.

II. MATERIALS AND METHODS

The solution to the restricted dynamic range lies in the development of an FPGA-based firmware that constantly keeps track of the resonance frequencies of the MKIDs within an array. Therefore, rather than producing a static probe tone for each detector that is reset once information is lost, the tone-tracking firmware will constantly update the frequency of the probe tone to assure that it is always on resonance. This tone-tracking firmware can be accomplished with a feedback control system.

The most important aspects of most control systems, known as the "invariant," is a variable which in this case is independent of the incident power on the detectors. Invariant properties allow for a reference during feedback loops which allow the system to change relative to a stable value. Resonators have such a property due to their electrical length at resonance being zero. This means that regardless of incident power on a KID, the phase of a tone at resonance will never change (ignoring linear phase shifts from signal lines connected to the system).

Since the goal of resonance tracking is to move a probe tone in frequency space rather than phase space, we require a linear conversion factor to change our phase feedback into frequency feedback. Luckily as seen in Fig. 2, resonators have a linear regime in the relationship between phase and frequency which can be measured quite simply.

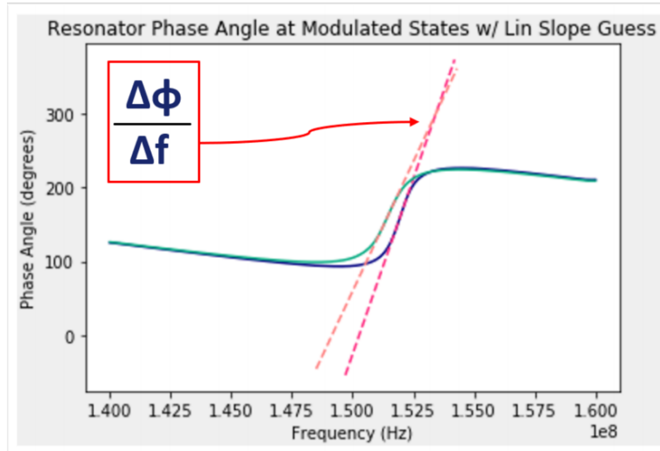

Fig. 2. A plot highlighting the linear regime in the phase/frequency relationship of a resonator in both the rest state and a modulated state. Previous applications of KIDs have confirmed that resonant frequency will not stray from this linear portion. While modulation of the KID causes changes in slope, the tone tracking system is designed in a way where these differences can be ignored.

With the tools of an invariant and a conversion factor in hand, it is now possible to design an algorithm to track the resonant tone of a KID. Before observing, one can determine the original resonant frequency of the KID (f_{res0}), the resonant phase (φ_0) and the linear regime slope ($\Delta\varphi/\Delta f$). Once power falls on a KID, there will be a change in phase of ($\Delta\varphi$). Putting these together, it only takes a simple set of equations to reach our final goal of the new resonant phase (f_{res_new}).

$$(\Delta f/\Delta\varphi) * \Delta\varphi = \Delta f \quad (1)$$

$$f\_res0 + \Delta f = f\_res\_new \quad (2)$$

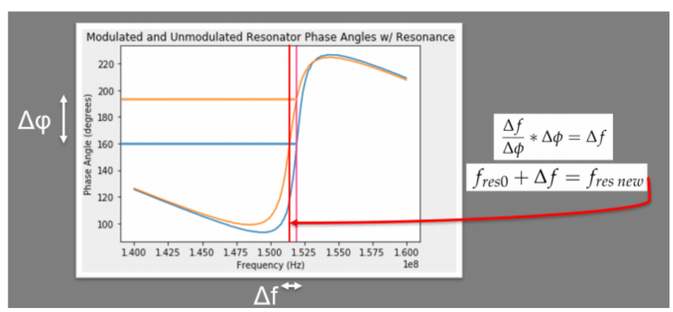

Fig. 3. A graphical representation of the tone tracking algorithm.

Fig. 3 provides a graphical representation of the tone-tracking algorithm and how it uses measurable quantities and known relationships to shift probe tones in response to changes in incident power on the detector. In digital control terms, the tone-tracking algorithm is a combination of a phase-locked loop (PLL) feedback-controlled with a proportional/integrative (PI) controller.

The block diagram shown in Fig. 4 shows in more detail the tone-tracking algorithm for a single detector. Starting with a probe tone at the KID's original resonant frequency, the signal is outputted by the digital to analog converter (DAC) and travels through the transmission line coupled to the KID. The affected tone then re-enters the FPGA by a digitizer (ADC) in quadrature and is multiplied with a complex conjugate of the probe tone to mix down to DC. The phase of the DC signal is calculated and fed through a filter to remove small fluctuations. A differential is taken between the new phase and the resonant phase to find the keystone value of Δφ. This differential is multiplied by the inverse slope of the estimated phase/frequency ratio to convert it to a differential frequency. Since there is lag time due to the length of the tone-tracker as well as uncertainty in the accuracy of the linear regime slope, this frequency difference (Δf) is scaled down considerably as to not "overshoot" the goal. The scaled down (Δf) is then added incrementally over multiple loops through the firmware to the original probe tone and inserted into a CORDIC wave generator to make a new probe tone closer to resonance. Finally, the probe tone exits the DAC and repeats the cycle until resonance is found and a phase differential of zero is achieved.

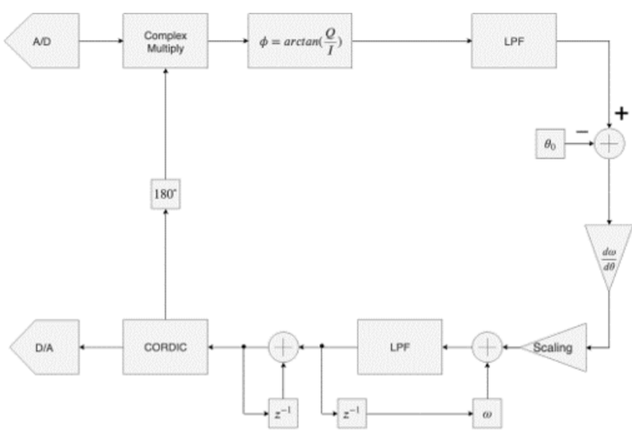

Fig. 4. A block diagram of the single-channel tone-tracking algorithm firmware. The diagram can be interpreted as follows: 1. Probe tone leaves DAC and goes through observing KID 2. Resulting tone returns to the ADC and is mixed to DC 3. Phase of input is calculated and integrated 4. Differential is taken to find Δφ 5. Δφ is scaled by estimated linear slope ratio 6. Differential frequency is added to original probe frequency 7. Probe tone is made with a CORDIC design and enters DAC 8. Return to step one and repeat

III. Simulation and Results

A firmware based on the algorithm explained in Fig. 4 was developed within the Simulink System Generator environment in MATLAB. The System Generator function of MATLAB's Simulink is the steppingstone between the block diagram and usable firmware as it allows for graphical DSP design and converts the designs into a hardware defined language (HDL) compatible with FPGAs. While the structure of the firmware is similar to the block diagram, the digitizer we plan to use operates at twice the clock speed of the FPGA fabric and therefore two samples will be fed for every clock cycle on the FPGA. To account for this, a parallel design was created so that the tone-tracker can handle two samples at once.

With proper handling, Simulink can also act as an extremely effective method of simulating firmware designs. To accomplish this, a sufficiently analogous "digital resonator" was developed within MATLAB and fed into the ports of the tone-tracking design. Simulations were executed starting at the simplest possible model, a simple resonance shift in a noiseless environment. Two shifts were used, one on the scale that is expected to be seen during observation and one significantly larger to stress test the algorithm. In the larger case, the algorithm found resonance on the order of 10 microsecond with a small overshoot. The smaller resonance shift was relocated in less than half the time with no oscillations. The time taken until a residual phase error of zero was reached can be seen in the plot in Fig. 5.

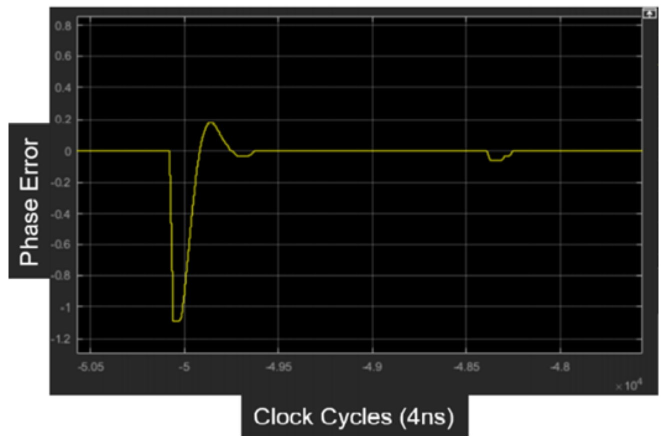

Fig. 5. Noiseless simulation results of tone-tracking algorithm measuring the difference between the real-time phase and resonant phase. Zero means the tone is tracked. The first spike represents a large phase shift and a dampened oscillation is visible. The second bump is a much smaller and more realistic phase shift.

For a more realistic scenario, white noise was added to the system and another set of tests was run, still using the same pair of large and small shifts. Fig. 6 shows these two simulated resonance shifts with an added white noise component. The tone tracking algorithm located the resonance in both cases on millisecond timescales with a negative-exponential-shaped curve.

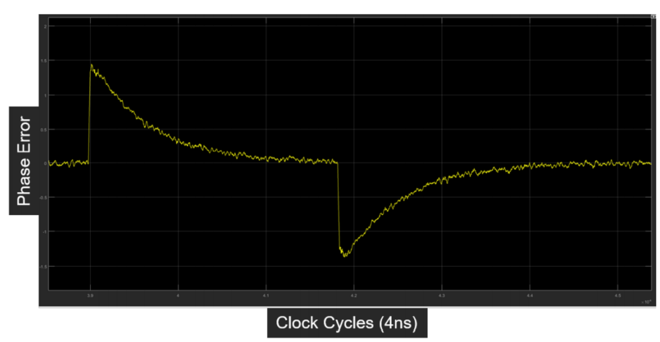

Fig. 6. Noisy simulation showing two simulated resonance shifts both finding the resonance on millisecond timescales.

IV. Discussion

The results of the tone-tracking simulations, while seemingly trivial, are a fantastic milestone in proving the validity of the algorithm for locating resonant frequencies in-situ. The noiseless examples seen in Fig. 5 verify that the fundamental algorithm used is an effective control system for reducing phase error to zero. Tweaks to the weight of the integrative portion in the P.I. controller could improve the speed and accuracy even further.

The noisy simulation in Fig. 6 is much closer to what will be encountered during observation. The results of this test imply that the time required to relocate resonance is largely independent of the modulation magnitude. The negative exponential shape of the phase differential implies that noise will significantly increase the timescale of relocating resonances. However, the shape of the plot also shows that the integrative portion of the controller is effective in eliminating the overshooting seen in the noiseless simulations. In its current state, the noisy tone-tracker was able to relocate resonances on millisecond timescales. Detector sampling rates from BLAST-TNG and other projects are on a similar millisecond timescale, confirming that the current tone-tracking algorithm is just fast enough to be effective in similar observation scenarios.

V. Conclusion

The ability to constantly track the resonant frequency of MKIDs in-situ has been desired for quite a while and the complexity of the problem speaks to why it still does not exist. Significant progress has been made over the past year in developing an FPGA-based tone-tracking system, but there is a great deal more to accomplish before it can be used in future MKID observations.

Since correctly tuning bit growth of the signal through the firmware seems to be posing a considerable hurdle, a Python simulation imitating the algorithm is under development that will focus on where bit-widths can be reduced. By using a software testing environment instead of firmware simulations, we can drastically reduce the iteration time for perfecting bit-propagation.

Furthermore, developing a firmware module which can read and update the tone-tracker in real time with the slope of the linear phase/frequency regime would greatly reduce the amount of time needed to relocate resonance. The current design for the tone-tracker can ignore the differences in these slopes by using very small iterative steps. By knowing the exact slopes, step size towards the new resonance can be maximized without overshooting and therefore reduce the number of clock cycles needed.

With the above steps completed, a single-channel version of a tone-tracking algorithm can finally be compiled and uploaded to an FPGA for tests on room temperature resonators and eventually cooled MKIDs. Once the single-channel design finds success, the firmware will be expanded to its full 1024/2048 channel form for use in MKID arrays. Finally, all designs, software, and firmware modules will be uploaded to GitHub and made open source for others to use and modify so that MKIDs can continue their rising popularity in the astronomy community

VI. References

1. Mauskopf, P. D., "Transition edge sensors and kinetic inductance detectors in astronomical instruments," Publications of the Astronomical Society of the Pacific 130, 082001 (jun 2018).

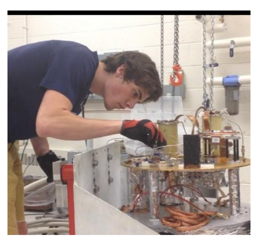

Jonathan R. Hoh is a Ph.D candidate in astrophysics and a Master's student in Microwave engineering at Arizona State University, School of Earth and Space Exploration. His academic interests include the design and miniaturization of Terahertz frequency heterodyne detection systems along with the implementation and readout of microwave kinetic inductance detector.