

Second Generation Readout For Large Format Photon Counting Microwave Kinetic Inductance Detectors

Neelay Fruitwala,^{1, a)} Paschal Strader,² Gustavo Cancelo,³ Ted Zmuda,³ Ken Treptow,³ Neal Wilcer,³ Chris Stoughton,³ Alex B. Walter,⁴ Nicholas Zobrist,¹ Giulia Collura,¹ Isabel Lipartito,¹ John I. Bailey III,¹ and Benjamin A. Mazin¹

¹⁾*Department of Physics; University of California; Santa Barbara; California 93106; USA*

²⁾*Dominican School of Philosophy and Theology; 2301 Vine St; Berkeley; CA 94708*

³⁾*Fermi National Accelerator Laboratory; Batavia; IL 60510; USA*

⁴⁾*Jet Propulsion Laboratory; California Institute of Technology; Pasadena; California 91125; USA*

(Dated: 16 November 2020)

We present the development of a second generation digital readout system for photon counting microwave kinetic inductance detector (MKID) arrays operating in the optical and near-IR wavelength bands. Our system retains much of the core signal processing architecture from the first generation system, but with a significantly higher bandwidth, enabling readout of kilopixel MKID arrays. Each set of readout boards is capable of reading out 1024 MKID pixels multiplexed over 2 GHz of bandwidth; two such units can be placed in parallel to read out a full 2048 pixel microwave feedline over a 4 – 8 GHz band. As in the first generation readout, our system is capable of identifying, analyzing, and recording photon detection events in real time with a time resolution of order a few microseconds. Here, we describe the hardware and firmware, and present an analysis of the noise properties of the system. We also present a novel algorithm for efficiently suppressing IQ mixer sidebands to below -30 dBc.

I. INTRODUCTION

Microwave Kinetic Inductance Detectors (MKIDs) are a low temperature superconducting detector technology used for astronomical observations in the submillimeter through visible bands. In the optical/IR, MKIDs are single photon sensitive, energy resolving, and allow photon arrival time determination on microsecond timescales¹. This makes MKIDs ideally suited to the study of highly time-variable phenomena such as pulsars² and compact binaries³. MKIDs are also suited to serve as IFUs (integral field units) for exoplanet direct imaging, which requires high sensitivity (due to the extremely low flux from companions), and benefits from high time resolution (to resolve atmospheric effects). Three optical/IR MKID instruments have been built and commissioned: ARCONS (Array Camera for Optical to Near-IR Spectrophotometry, a 2 kilopixel IFU at Palomar Observatory)⁴, DARKNESS (DARK-speckle Near-infrared Energy-resolving Superconducting Spectrophotometer, a 10 kilopixel IFU at Palomar)⁵, and MEC (MKID Exoplanet Camera, a 20 kilopixel IFU at Subaru Observatory)⁶.

One of the primary advantages of MKIDs over other superconducting detector technologies is the intrinsic frequency domain multiplexability; each MKID pixel is a microwave frequency resonator with a resonant frequency that can be easily tuned in fabrication. This allows large numbers of resonators, each tuned to a different frequency, to be placed in parallel on a single microwave channel. This greatly simplifies the cryogenic wiring and electronics, shifting the complexity to the room temperature readout system⁴⁻⁷.

A first generation readout for optical MKIDs was developed in 2011 for the ARCONS instrument. This system is

capable of reading out 256 channels in a 512 MHz band on each readout board⁸. Eight boards were used simultaneously to readout the 2048 pixel ARCONS instrument. In order to feasibly scale up to the kilopixel arrays used by the second generation of optical MKID instruments, we have developed a new readout system capable of reading out up to 1024 channels per unit over 2 GHz of bandwidth. This system largely retains the signal processing algorithm chain found in the first generation readout, albeit with major changes to the hardware and firmware to accommodate the higher bandwidth.

II. OVERALL APPROACH AND REQUIREMENTS

An MKID is a cryogenically cooled superconducting resonator consisting of an inductor/capacitor pair. The circuit gets its inductance primarily from the kinetic energy of the Cooper pairs in the supercurrent (hence the term kinetic inductance). When a photon is absorbed by the inductor, some of the Cooper pairs are broken (creating free electrons known as “quasiparticles”), changing the inductance and hence the resonant frequency and quality factor of the circuit. The number of quasiparticles created is proportional to the energy of the incident photon, which gives the MKID intrinsic energy resolution. To read out the detector, the resonator is driven by a microwave probe signal tuned to its resonant frequency and the probe signal is monitored for changes in its relative phase. A photon detection event manifests as a pulse in phase with a steep rise (~ 1 μ s) and an (approximately) exponential decay with ($\tau \approx 15$ μ s). Our devices are designed to operate in the linear regime, where $\delta\phi \propto \delta f \propto \delta E$ ^{1,9}.

The full kilopixel readout generalizes the above approach to an array of resonators placed in parallel on a single microwave feedline. A comb of probe tones, one at each resonator frequency, is used to drive the MKIDs. Each resonator acts as a notch filter centered on its resonant frequency, modifying

^{a)}Electronic mail: neelay@ucsb.edu

the phase and amplitude of its probe tone while leaving other tones comparatively untouched. After transmission through the array, the tones are downconverted and digitized, and each tone is independently monitored for changes in its relative phase.

The second generation of MKID arrays are in the 10-20 kilopixel range, with pixels split up into several microwave feedlines, each containing 2000 pixels inside the 4-8 GHz band. We expect a maximum photon event rate of 2500 counts/pixel/second. Due to data transmission and storage constraints, we require that each event be identified, analyzed, and recorded in real time. Following these constraints, we define the following system requirements:

- Generation of a 2000 tone comb in a 4-8 GHz band, with arbitrary tone frequencies and powers.
- Channelization: the ability to isolate a 200-500 kHz wide channel centered around each tone
- Low noise: the noise floor of each channel must be lower than the device noise and cryogenic amplifier noise
- Sample each channel at a rate of $\sim 1 \mu s$
- Implement a filtering and triggering system that is capable of detecting photon events in real time and accurately determining the associated phase pulse amplitude

III. SYSTEM OVERVIEW

The full readout system consists of several readout units, each of which can read out up to 1024 resonator channels in a 2 GHz band. Two such units can be placed in parallel to read out a full 4 GHz feedline. Because the resonant frequency band lies outside the range of readily available ADCs, we use an IQ modulation scheme to perform up/down conversion between the digital room temperature electronics (-1 to 1 GHz) and the resonators (4 to 6 GHz or 6 to 8 GHz). Each readout unit has three major components: 1) CASPER (Collaboration for Astronomy Signal Processing and Electronics Research) ROACH-2 (Reconfigurable Open Architecture Computing Hardware) board with a Xilinx Virtex-6 FPGA for channelization and pulse detection, 2) Custom ADC/DAC board with dual 2 GSPS Analog Devices AD9625 ADCs and AD9136 DACs and Xilinx Virtex-7 FPGA for control, and 3) RF/IF board for IQ modulation.

The basic outline of the readout procedure (illustrated in figure 1) is as follows:

1. Tone generation in [-1 GHz, 1 GHz] (IF band) using dual 2 GSPS DACs.
2. Upconversion to RF band via IQ mixing tones with local oscillator (LO) (figure 2).
3. Tones filtered through resonators, amplified by cryogenic HEMT (high electron mobility transistor).

4. Output is downconverted by IQ mixer, then digitized by dual 2 GSPS ADCs.
5. ADC stream is sent to Virtex-6 FPGA for channelization, filtering, and photon triggering.
6. Photon events (time and pixel tagged pulse height) are streamed over ethernet to data server.

IV. HARDWARE

A. ADC/DAC Board

Due to the complexity of modern, high sample rate ADCs and DACs, we have developed a dedicated circuit board for signal routing and control of these components. The board features dual 12-bit 2 GSPS ADCs (Analog Devices AD9625), and dual 16-bit 2 GSPS DACs (Analog Devices AD9136). Dual ADC/DACs are required for complex sampling; one unit is required for each of I and Q. The board also has an onboard Xilinx Virtex-7 FPGA (XC7VX330T-2FFG1761C) for control of the ADCs and DACs, routing of the ADC signal to the ROACH-2, and communication with the ROACH-2 and RF/IF boards. A LMK04821 synthesizer is used to generate all clocks required by the ADCs, DACs, and FPGAs on both the ROACH-2 and ADC/DAC boards, synchronized to an external 10 MHz reference signal.

The DAC output is generated using lookup table (LUT) that contains the sum of all of the resonator probe tones in the time domain. The LUT is generated in software and stored in on-board DDR3 RAM. Two DACs are used to generate a complex output signal, one for I and one for Q. A table size of 262144 values (for each of I and Q) imposes a tone frequency quantization of 7.629 kHz. The LUT approach allows us to generate the DAC output for an arbitrary list of resonator frequencies, and independently specify the drive power for each resonator (required to optimize photon pulse SNR). To generate the LUT, tones are first calculated and summed together in software according to their IF band frequencies and relative powers. Tone phases are randomized to maximize DAC dynamic range utilization. The full LUT is then scaled to fit the DAC dynamic range, and the programmable output attenuators are adjusted to set the tones to the desired power. Programmable attenuators in the RF output chain ensure full utilization of the DAC dynamic range over a wide range of tone powers.

The IQ ADC/DAC inputs/outputs are connected to the RF/IF board via SMP connectors. There are also several SPI interfaces to enable the Virtex-7 to program the attenuators and local oscillator on the RF/IF board.

The raw ADC output is sent to the Virtex-6 FPGA on the ROACH-2 over a ZDOK connector 16 samples (8 each from I and Q) per clock in a 192-bit parallel bus clocked at 250 MHz.

B. ROACH-2 Board

The ROACH-2 was developed by the CASPER collaboration at UC Berkeley^{10,11}. It was designed around the Xilinx

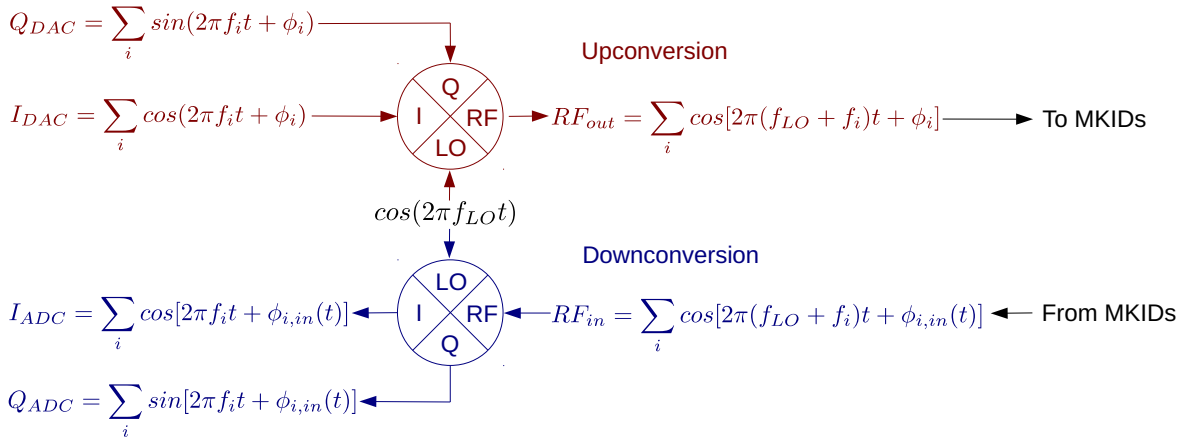
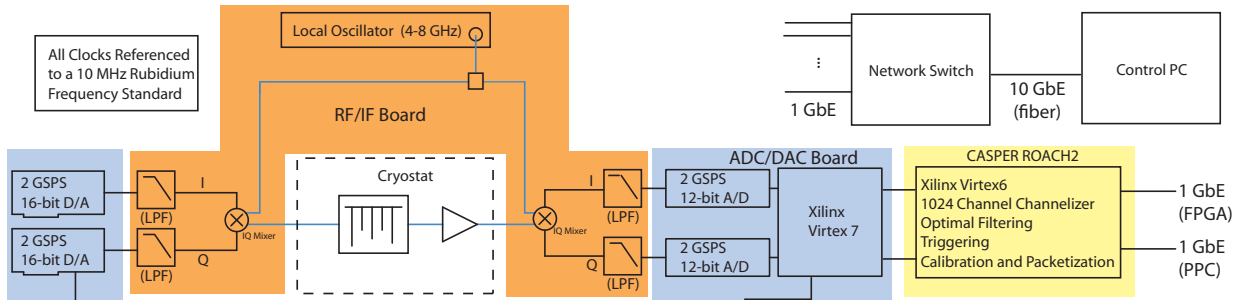


FIG. 2. Schematic of RF up/downconversion scheme. The frequency comb is written explicitly as a sum over the probe tones, where $f_i \in [-1 \text{ GHz}, 1 \text{ GHz}]$ is the IF band tone frequency, and $f_{LO} \in [4 \text{ GHz}, 8 \text{ GHz}]$ is the LO frequency. ϕ_i is the probe tone phase in the DAC LUT. $\phi_{i,in}(t) = \phi_i + \phi_{i,cable} + \delta\phi_{i,res}(t)$ is a time dependent phase term that includes the DAC LUT phase, cable delay, and phase change induced by the resonator. The same LO is used for both upconversion and downconversion to suppress phase/frequency errors in the LO signal.

Virtex-6 FPGA (XC6VSX475T-1FFG1759C). It has a PowerPC CPU, DDR3 and QDR memory modules, and various communications interfaces such as ethernet. All of the core digital signal processing for the readout system (channelization, filtering, and pulse detection) is implemented on the Virtex-6 FPGA. We chose the ROACH-2 platform because CASPER provides an extensive dedicated toolset and programming interface for the FPGA. The CASPER toolflow contains blocks for performing complex signal processing tasks (such as the polyphase filter bank), as well as interfaces to much of the hardware on the ROACH-2. CASPER has also released a series of python libraries for programming and communication with the FPGA over ethernet.

C. RF/IF Board

1. IQ Upconversion and Downconversion

The RF/IF board implements IQ modulation for analog up/downconversion between the IF band (-1 – 1 GHz) and

the RF band (4 – 6 GHz or 6 – 8 GHz). For the upconversion process, the complex IF band frequency comb is mixed with a local oscillator (LO) tone at the center of the RF band (i.e. a 4 – 6 GHz resonator band would use a 5 GHz LO), which generates a real-valued output that is shifted into the resonator frequency band. Downconversion is the reverse process, where the real valued RF output from the cryostat is mixed with the same LO to generate a complex-valued IF band signal which can be digitized by the ADCs. This process is illustrated schematically in figure 2. The LO is generated on the board by a programmable phase locked loop (PLL) IC (Texas Instruments TRF3765). The PLL is referenced to an external 10 MHz signal provided by the ADC/DAC board. A frequency doubler is used on the output of the TRF3765 to shift the frequency range into the resonator band (600 MHz – 9.6 GHz). Analog devices HMC525LC4 I/Q mixers are used for upconversion and downconversion.




FIG. 3. Cartridge containing an assembled readout unit. The ROACH-2 board is connected to the ADC/DAC board by two ZDOK connectors. The RF/IF board is mounted on the ADC/DAC board using SMP blind-mate connectors for signals and GPIO pins for programming. Another readout unit is mounted to the underside of this cartridge. Figure and Caption reproduced with permission from Paschal Strader, "Digital Readout for Microwave Kinetic Inductance Detectors and Applications in High Time Resolution Astronomy", Ph.D. dissertation (University of California at Santa Barbara, 2016).

2. RF Output Chain

Two Peregrine Semiconductor PE43705 programmable step attenuators provide 0 to 63.5 dB (steps of 0.25 dB) of attenuation on the output frequency comb, which enables full utilization of DAC dynamic range over a range of resonator readout powers. The maximum output power of the full RF chain is -12.5 dBm.

3. RF Input Chain

In order to ensure full utilization of the ADC dynamic range over a wide range of resonator drive powers, the RF/IF board has an amplifier/attenuator chain before the IQ mixer. Four Hittite HMC3587 provide 60 dB gain, and two programmable step attenuators (63.5 dB total in steps of 0.25 dB) are used to optimize the input power at the ADC and ensure that the amplifiers are kept away from saturation (see figure 12 for schematic). The optimal signal power at the end of the amplifier chain (IQ mixer input) should be approximately -7 dBm; this implies an input power range of -64 dBm to -0.5 dBm¹². In practice, we find that in order to keep the noise from the room temperature chain sufficiently low, the total programmable attenuation should not exceed ≈ 30 dB, providing an actual input power range of -64 dBm to -34 dBm (see Appendix A for full calculation).

D. Timing

The ROACH-2 firmware has a timing block that keeps track of the absolute (UTC) time with microsecond precision. In or-

der to ensure accurate absolute timing, this block can be synchronized to an external reference using a pulse-per-second (PPS) input. In our current implementation, the PPS input and the 10 MHz reference used to generate the FPGA clocks are generated by a spectracom timing module that is synchronized to a GPS satellite reference.

V. FIRMWARE AND SOFTWARE

A. ADC/DAC (Virtex-7 FPGA)

The Virtex-7 firmware is responsible for controlling the ADCs, DACs, and programmable components on the RF/IF board, and for routing the ADC and DAC data streams. Communication with the Virtex-7 is conducted over a UART interface to the ROACH-2 Virtex-6 FPGA, which can receive and forward commands from the control computer over ethernet. The low-level firmware and communications interfaces are controlled using a Xilinx MicroBlaze soft processor. The MicroBlaze runs a C script which initializes all major peripherals, then receives and executes commands sent by the ROACH-2 over the UART interface. These commands include programming the RF attenuators and LO, receiving the DAC LUT and writing it to DDR3, and resetting the ADCs and DACs. The firmware was developed primarily at Fermilab, and testing and integration with the ROACH-2 were performed at both Fermilab and UCSB.

B. ROACH-2 (Virtex-6 FPGA)

The Virtex-6 firmware is responsible for the bulk of the digital signal processing performed by the system: conversion from the raw 2 GHz I/Q stream from the ADCs to time and phase (or energy) tagged photons. This process can be divided into four stages: channelization, phase conversion, optimal filtering, and photon event triggering. The firmware was developed at UCSB using the MATLAB Simulink and ISE based CASPER ROACH toolflow. All of the firmware blocks are clocked at 250 MHz using a clock supplied by the ADC/DAC board over the ZDOK connectors.

1. Channelization

Channelization is the process of converting the raw 2 GHz I/Q stream from the ADC into a time-multiplexed series of frequency bins, each centered around a resonator probe tone. We use the same two stage approach as the first generation readout, as described in (Ref. 8). The first stage is a dual 2048-sample (more specifically, 2048 branch, 4 tap) polyphase filter bank (PFB), which consists of a finite impulse response (FIR) filter followed by a Fast Fourier Transform (FFT). The configuration is the exact same as in (Ref. 8), except for the total bandwidth/number of points (2048 samples over 2 GHz, instead of 512 samples over 512 MHz). As

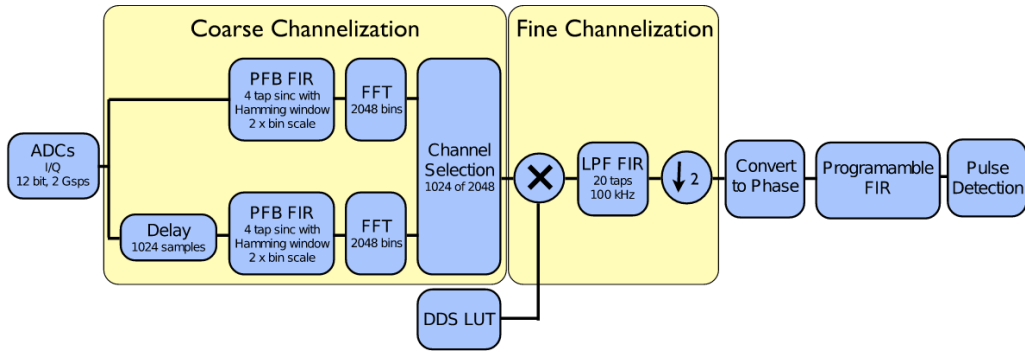


FIG. 4. ROACH-2 firmware block diagram. Data is streamed over the ZDOK over a 16-sample bus (8 12-bit I/Q pairs) clocked at 250 MHz. The data is then copied and processed by two parallel PFBs; one copy is delayed by 128 clocks (half of the complete FFT cycle) to maintain a uniform sampling interval of each bin at 2 MHz⁸. After the PFB stage, up to 1024 resonators are selected from the bins. Following channel selection, DDC is performed on each channel to center the resonator tone within the frequency bin. Channels are then low pass filtered (100 kHz on each of I and Q) and converted to phase. The phase is filtered using a channel-specific “optimal” FIR. Phase pulses meeting the trigger conditions are recorded and streamed to the data server over 1 GBit ethernet. Figure adapted with permission from Paschal Strader, “Digital Readout for Microwave Kinetic Inductance Detectors and Applications in High Time Resolution Astronomy”, Ph.D. dissertation (University of California at Santa Barbara, 2016).

in (Ref. 8), oversized 2 MHz bins are used, so two parallel PFBs are required to Nyquist sample each bin. Oversized bins ensure that there exists a bin with at least 500 kHz of bandwidth around every resonator, no matter its resonant frequency. The output of the PFB+FFT is time multiplexed at $(2 \times FFT) \times (8 \text{ bins/clock}) = 16 \text{ bins/clock}$. The FPGA is clocked at 250 MHz, so the sample rate of each bin is $(250 \text{ clocks/s}) \times (16 \text{ bins/clock}) / (2048 \text{ bins}) \approx 2 \text{ MHz}$.

After the FFT, the bins are sorted into 1024 resonator channels; for each resonator frequency, the bin with the closest center frequency is selected. The allowed spacing between resonators can be as low as 200 kHz, so there may be as many as 10 resonators per bin. These channels are split into 4×256 -channel time-multiplexed “streams”, with each resonator getting 2 samples per clock (one from each FFT). This maintains the 2 MHz channel sample rate: $(250 \text{ MHz}) \times (2 \text{ samples}) / (256 \text{ channels}) \approx 2 \text{ MHz}$. Digital down conversion (DDC) is then performed on each channel to center it at 0 Hz: each channel is multiplied by a (complex) sinusoid with $f = f_{bin} - f_{res}$. These sinusoids are generated using a lookup table which is stored on the ROACH-2 QDR memory.

2. Low Pass Filtering and Phase Conversion

After the channelization stage, each channel is low pass filtered to set the desired channel bandwidth. This is done by separately convolving I and Q with a 20-tap Hanning-windowed sinc function. The following considerations should be taken into account when setting the channel bandwidth: 1) desired time resolution/response time of the detector (1 μ s); 2) photon event (phase pulse) SNR; 3) desired minimum resonator spacing. Exoplanet imaging requires framerates of at most a few kHz ($\sim 100 \mu$ s), so time resolution was not a major consideration. We found that a channel bandwidth of 200 kHz (± 100 kHz around the tone center; which gives 100

kHz bandwidth after phase conversion) provides good performance for considerations (2) and (3). This bandwidth preserves the frequency content of a typical pulse (figure 5) while attenuating much of the high frequency line noise present in many resonator channels (figure 7). It also allows us to impose a minimum channel spacing of 200 kHz, reducing the number of frequency collisions to $\approx 6\%$ of identified resonators. After the low pass filter, the channels are downsampled to 1 MHz, and the I/Q values are converted to phase using a CORDIC algorithm that implements $\arctan(Q/I)$.

3. Optimal Filtering

After phase conversion, each channel is filtered using a 50-point FIR (finite impulse response) filter. The filter coefficients are programmable and are unique to each channel. This customizability is a key feature, as it allows us to tailor each filter to that channel’s noise PSD and photon pulse shape. To generate these filters, we use the formalism described in (Ref. 13) and extended in (Ref. 14). This “optimal filter” formalism uses the noise covariance and a template of the photon pulse to produce a minimum variance linear estimator of the photon pulse amplitude. The number of filter coefficients is well matched to the characteristic pulse timescale ($\tau \approx 15 \mu$ s), as indicated by the filtered pulse being fairly symmetric (figure 6).

4. Photon Event Triggering and Streaming

The total phase data rate for a single readout unit is approximately 18 Gbps (1024 channels \times 1 MHz per channel \times 18 bits). This data rate is difficult to stream over ethernet (a 20 kpix array would require 360 Gbps of bandwidth) and infeasible to save for a full night of observing. To reduce the




FIG. 5. Top: representative phase pulse template, made by averaging together phase collected on a single resonator channel, then fitting to a triple exponential function. Bottom: power spectrum of the above pulse. Note that there is very little power above 100 kHz.




FIG. 6. Example photon signal on a single resonator channel. (blue) is data collected after the 100 kHz LPF and phase conversion steps; (orange) is the same data stream after applying a 50 coefficient optimal FIR filter. Photon wavelength is 1050 nm. The sampling interval is $1.024 \mu\text{s}$

data rate, we use a firmware trigger to find and record photon pulses. The trigger consists of two conditions:

1. Local minimum detection - pulse must be preceded by 9/10 points w/ negative derivative, and immediately proceeded by two points w/ positive derivative.
2. Threshold - pulse amplitude must exceed a threshold value. This threshold is programmable and can be specified individually for each channel. Typically, we use some multiple of the standard deviation of the phase value ($3.5 - 4\sigma$).

If the trigger conditions are met, the pulse amplitude, arrival time, and channel ID are recorded in a 64-bit data packet. Multiple such data packets are packed into UDP frames (up to 100 packets per frame) and streamed to the data server over ethernet.

VI. PERFORMANCE

A. Phase Noise

A key system performance metric is the phase noise in each resonator channel (i.e. the noise PSD of the phase timestream used for photon detection). The phase noise has three major sources: 1) noise intrinsic to the device, 2) noise contributed by the cryogenic amplification chain, 3) noise contributed by the readout system¹⁵. Ideally, phase noise is dominated by sources (1) and (2); contributions from the readout system should be negligible. To characterize this, we measure the readout system phase noise in loopback (RF input connected directly to RF output), and compare this to the expected best case performance of the device and cryogenic amplifier.

The noise contributed by the device comes primarily in the form of a $1/f$ contribution due to two-level system (TLS) noise^{16,17}. We ignore this contribution, and instead use the cryogenic amplifier noise as a basis for comparison. This is because the TLS noise varies significantly between devices and is difficult to estimate, and using cryogenic amplifier noise alone provides a more stringent performance criterion than using the combination of TLS noise and cryogenic amplifier noise. For the cryogenic amplifier we assume a 2.3 K HEMT (high electron mobility transistor), which is used in both MEC⁶ and DARKNESS⁵. We assume that it has 40 dB gain and contributes thermal white noise at the device stage with a 2.3 K noise temperature (eqn. A1). We report all phase noise measurements in dBc/Hz, where dBc is the noise power relative to the probe (carrier) tone in a single quadrature.

The noise contributed by the digital readout electronics is dominated by room temperature thermal noise and amplifier noise in the input RF chain. The exact phase noise contribution from the room temperature stage depends on the programmable attenuator settings; we calculate that for typical values the contribution of room temperature noise floor should be at least 6 dB lower than the HEMT noise contribution (see appendix A for calculation). The measured phase noise floor




FIG. 7. Phase noise spectrum of a single resonator channel. Measurement was taken in loopback with a full 1024 tone comb in 4 – 6 GHz band. Aside from spurious signals, the noise floor is approximately -93 dBc/Hz.

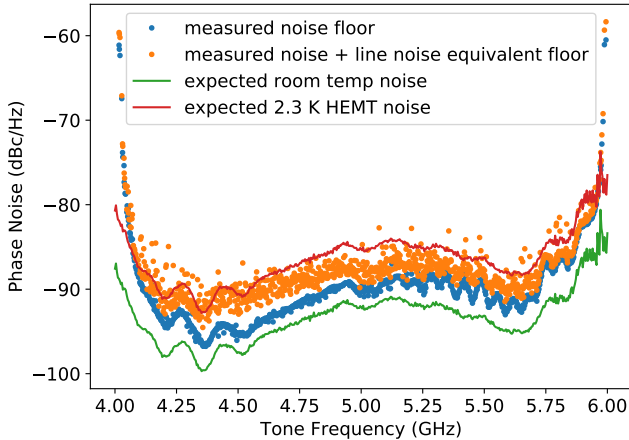


FIG. 8. Loopback phase noise measurement for 1024 tone comb on a single readout unit. The measured phase noise value for each channel was computed by fitting the white noise component of the phase noise spectrum (example spectrum in figure 7). The equivalent phase noise floor after accounting for line noise is also plotted (calculation in appendix B). The expected phase noise contributions from the HEMT and room temperature chain were estimated using the measured tone power at the RF input (see appendix A for calculation).

(figure 8) is ≈ 3 dB higher than expected, likely due to unaccounted for attenuation in the RF input chain.

To verify system performance, we performed a loopback measurement of the phase noise of each tone in a full 1024 tone frequency comb. The tones were given frequencies in 4 – 6 GHz range and uniform output power in the DAC LUT. The output power of each RF tone ranged from approximately -65 dBm to -70 dBm, where the variation results from the nonuniform frequency response of the RF output chain (which we calibrate out when reading out actual devices). The exact readout power (power at the MKID) that this corresponds to

depends on the array properties and experimental setup, but the typical RF attenuation before the device is ≈ 35 dB, which corresponds to a readout power of approximately -100 dBm to -105 dBm. Results are plotted in figure 8, along with expected contributions from the room temperature amplifier chain and HEMTs.

The measured phase noise has significant line noise (figure 7). The line noise occurs at multiples of 7.629 kHz, which is the quantization frequency of the DAC comb. So, we suspect that the lines are generated as intermodulation products in the input RF chain. In order to compare the line noise to the HEMT and room temperature noise, we examine the effect each has on the variance of the measured photon pulse amplitude. We do this by removing the spectral lines from the noise PSD, then scaling this “flattened” PSD such that its corresponding pulse height estimator variance is equal to that of the original PSD (see appendix B for full calculation).

There is a significant apparent difference between our phase noise measurements (≈ -93 dBc/Hz), and the measured phase noise for the first generation system (-106 dBc/Hz, including HEMT noise⁸). This is because the first generation system measurements were performed at -85 dBm⁸ (significantly higher than used in practice with ARCONS), and our measurements were performed at an expected readout power of -100 dBm to -105 dBm (which is close to the -106 dBm readout power¹⁵ of the PtSi devices used in DARKNESS⁵ and MEC⁶). So, we expect that any thermal noise source (which includes both HEMT and room temperature amplifier noise) will contribute up to 20 dB higher phase noise (phase noise is inversely proportional to tone power; see eqns. A1, A2, A5) for readout tone powers characteristic of our system as compared to the first generation system; this is reflected in our measurements in figures 7 and 8.

It is difficult to determine the effect of readout system noise on the ultimate resolving power of the MKIDs, as this depends on detector sensitivity (i.e. pulse height as a function of photon energy), quality factor, and device noise, which vary significantly between pixels^{1,6}. However, the majority of channels are below the expected HEMT noise floor, so these would obey the HEMT noise limits measured in (Ref. 15). We also note that the measurements in figures 7 and 8 are close to worst-case performance, with RF input attenuators $A_1 = 14.25$ dB and $A_2 = 14.5$ dB (see figure 12 for block diagram of RF input chain). A more typical use case would have $A_{1,2} < 10$ dB, which will reduce the phase noise floor by > 2.7 dB (eqn. A6), making almost every channel HEMT noise limited.

VII. SIDEBAND SUPPRESSION

Imperfections in IQ mixer performance lead to the presence of “sidebands”, or reflections across the LO frequency when the mixer is used to upconvert or downconvert RF tones. For example, a 6.5 GHz tone generated using a 6 GHz LO will have a sideband tone at 5.5 GHz (which may also show up at -500 MHz in the IF band after downconversion). Ideally, these sidebands are > 30 dB down from the original tone. In our system, we have observed sidebands with powers as high

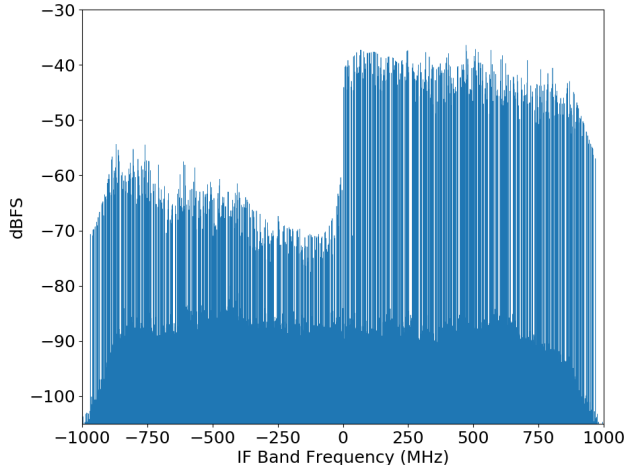


FIG. 9. Example sideband spectrum (taken using the ADC on loopback). Tones on the right half ($f > 0$) are actual tones generated by the DAC, while tones on the left half are undesirable sideband reflections of these generated tones. In the worst case, sideband power is only 10-15 dB down from the tone power.

as 10 – 15 dB below the original tone. A sideband is problematic if it happens to be within 100 kHz of another tone, as it will leak into that channel and cause its phase timestream to oscillate at $f = f_{\text{tone}} - f_{\text{sideband}}$, degrading the performance of that pixel.

The input to the I/Q mixer can be adjusted to compensate for mixer errors and suppress sidebands. This can be done for both upconversion and downconversion:

- Upconversion: adjust the relative phase/amplitude of the I and Q inputs to the mixer. This effectively introduces sidebands in the IF input to cancel them out in the output.
- Downconversion: introduce real RF sidebands to cancel out sidebands produced in the downconverted I and Q.

Performing these adjustments separately for upconversion and downconversion is infeasible, as injecting RF sidebands after upconversion would require major design changes. So, we instead adjust the output I and Q generated by the DAC to minimize total sideband power remaining after both upconversion and downconversion.

For our system, we experimentally determined that the optimal phase and amplitude adjustments are sensitive to tone power and frequency and are not always smooth functions of these. So, a global fitting approach would be suboptimal, and any adjustments must be made on a per-tone basis in the DAC LUT. We have developed an algorithm to perform this optimization for an arbitrary list of tone frequencies and powers. Our algorithm minimizes total sideband power remaining after both RF upconversion and downconversion (and everything else that might be in the RF chain, including resonators, etc). The optimization problem for N tones can be written as:

$$\max_{\Delta\phi_1, r_{IQ,1}, \dots, \Delta\phi_N, r_{IQ,N}} \sum_i P_i - P_{\text{sideband},i}(\Delta\phi_i, r_{IQ,i}) \quad (1)$$




FIG. 10. Plot of sideband suppression of a single tone as a function of ϕ and r_{IQ} . Close to the optimal value, the sideband suppression decays approximately exponentially, which is why an exponential fit is used in the search algorithm.

where i indexes the tones, P_i is the tone power, and $P_{\text{sideband},i}(\Delta\phi_i, r_{IQ,i})$ is the sideband power of tone i as a function of the phase offset $\Delta\phi_{IQ,i}$ (from 90°) between I_i and Q_i and $r_{IQ,i} = |I_i/Q_i|$ (I_i and Q_i are the DAC outputs for tone i). Sideband power is measured in the input IF band by the ADCs, in units dBFS (dB relative to ADC full scale). To simplify the problem, we have assumed that the tones are independent (i.e. adjusting the phase or amplitude balance of one tone will not affect the sideband power of any other tone); so the problem can be thought of as a series of N independent maximizations over the two dimensional domain given by $\{\Delta\phi_{IQ} \times r_{IQ}\}$.

The following procedure is used to evaluate the objective function for a specific set of $\{(\Delta\phi_{IQ,1}, r_{IQ,1}), (\Delta\phi_{IQ,2}, r_{IQ,2}), \dots, (\Delta\phi_{IQ,N}, r_{IQ,N})\}$:

1. Load DAC LUT containing these offsets
2. Take a 8,388,608-point snapshot (set by the amount of available memory) of the ADC input
3. Take an FFT of the snapshot to find sideband powers

This evaluation process takes approximately one minute, so a brute force optimization with the required range and precision (phase: $[-20^\circ, 20^\circ]$, 1° increments; amplitude ratio: $[0.8, 1.2]$ in increments of 0.2) would take approximately 13 hours. Since the optimization must be repeated for each new configuration of tones, the brute force approach is infeasible and a more efficient algorithm is needed. We have developed the following approach:

1. Sample the objective function at n randomly chosen sets of $\{(\Delta\phi_{IQ,1}, r_{IQ,1}), (\Delta\phi_{IQ,2}, r_{IQ,2}), \dots, (\Delta\phi_{IQ,N}, r_{IQ,N})\}$ (in addition to $(\Delta\phi_{IQ,i}, r_{IQ,i}) = (0, 0)$ for all i)
2. For each tone, fit $P_i - P_{\text{sideband},i}(\Delta\phi_i, r_{IQ,i})$ to an exponential decay function (see figure 10).
3. Randomly sample a new set of $\{(\Delta\phi_{IQ,1}, r_{IQ,1}), (\Delta\phi_{IQ,2}, r_{IQ,2}), \dots, (\Delta\phi_{IQ,N}, r_{IQ,N})\}$, either close to the




FIG. 11. Test of the sideband suppression algorithm on loopback using a 868 tone MEC array frequency comb. Top/middle: comparison of each tone’s power with its corresponding sideband power, before (top) and after (middle) running the sideband suppression algorithm. Bottom: Magnitude of sideband suppression across the frequency comb.

fit centers or uniformly over the domain (choose this with some probability ϵ).

4. Repeat (2) and (3) until all sidebands are sufficiently suppressed (typically $P_{tone} - P_{sideband} > 30$ dB), or the number of iterations has reached some threshold value.

Using this approach, we are able to suppress sidebands by at least 30 dB for almost every tone, with a convergence time of 30 - 45 minutes. Results from a loopback test using a MEC array frequency comb with 868 tones are shown in figure 11. It is difficult to provide a general estimate of the improvements to energy resolution resulting from suppressing sidebands to ≥ 30 dBc, as this depends on highly variable device characteristics (including sensitivity, quality factor, device noise), as well as the location of the interfering sideband in the spectrum of a nearby resonator channel. We feel that the 30 dBc benchmark is appropriate, as it has an equivalent RMS phase > 3 dB below that of a typical resonator channel (appendix C), and is the same benchmark used for

the first generation system⁸.

VIII. CONCLUSION

We have developed a digital readout system which, for the first time, makes it feasible to read out kilopixel-scale photon counting optical/IR MKID arrays. Our system is actively being used by two MKID cameras; MEC at Subaru Observatory, and DARKNESS at Palomar Observatory. All of our array calibration and control software, as well as the Virtex-6 firmware, is open source and can be found at: <https://github.com/MazinLab/MKIDReadout>.

IX. ACKNOWLEDGEMENTS

The authors would like to thank Jonas Zmuidzinas for valuable feedback regarding the phase noise calculations. NF is

supported by a grant from the Heising-Simons Foundation. NZ was supported throughout this work by a NASA Space Technology Research Fellowship. IL was supported by the National Science Foundation Graduate Research Fellowship under grant number 1650114. This work was funded in part by an NSF ATI grant AST-1308556, and by US Department of Energy, Office of Science, Office of High Energy Physics contract No. DE-AC02-07CH11359 w/ Fermi Research Alliance, LLC. We would also like to thank Xilinx for donating the FPGAs used for this research.

X. DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Appendix A: Phase Noise Calculations

The single-sided phase noise spectral density $S_{\delta\phi}$ due to Johnson noise from the HEMT is given by¹⁸:

$$S_{\delta\phi,HEMT} = \frac{k_B}{P_{device}} (T_{device} + T_{HEMT}) \quad (A1)$$

To facilitate comparison with loopback data taken at room temperature, we can instead use the tone power at the readout system RF input; $P_{RF} = G_{HEMT}P_{device}$, where G_{HEMT} is the HEMT amplifier gain. So,

$$S_{\delta\phi,HEMT} = \frac{k_B G_{HEMT}}{P_{RF}} (T_{device} + T_{HEMT}) \quad (A2)$$

The noise from the readout system is dominated by the room temperature amplifier chain immediately following the RF input of the RF/IF board. This chain consists of four HMC3587, each with a noise temperature of 438 K and two programmable attenuators (figure 12). The total single sided noise spectral density at the end of this chain (IQ mixer input) is given by:

$$\begin{aligned} S_{mixer} &= G \left\{ \frac{G}{2} \left(\frac{G}{A_2} \left[\frac{G}{A_1} (S_{RT} + S_{amp}) + S_{RT} + S_{amp} \right] + S_{RT} \right. \right. \\ &\quad \left. \left. + S_{amp} \right) + S_{RT} + S_{amp} \right\} \\ &= \left(\frac{G^4}{2A_1A_2} + \frac{G^3}{2A_2} + \frac{G^2}{2} + G \right) (S_{RT} + S_{amp}) \end{aligned} \quad (A3)$$

where $S_{RT,amp} = 2k_b T_{RT,amp}$ is the Johnson noise PSD at room temperature and from the amplifier, respectively, G is the room temperature amplifier gain, $A_{1,2}$ are the programmable attenuator values, and $T_{RT,amp}$ are noise temperatures at room temperature (290 K) and of the amplifiers (438 K).

The tone power at the end of the amplifier chain is given by:

$$P_{mixer} = \frac{G^4}{5A_1A_2} P_{RF} \quad (A4)$$

To obtain the single quadrature phase noise spectral density, we have:

$$\begin{aligned} S_{\delta\phi,RT} &= \frac{S_{mixer}}{2P_{mixer}} \\ &= \frac{k_B}{P_{RF}} \left(\frac{5}{2} + \frac{5A_1}{2G} + \frac{5A_1A_2}{2G^2} + \frac{5A_1A_2}{G^3} \right) (T_{RT} + T_{amp}) \end{aligned} \quad (A5)$$

Comparing the HEMT and room temperature terms:

$$\begin{aligned} \frac{S_{\delta\phi,RT}}{S_{\delta\phi,HEMT}} &= \left(\frac{5}{2} + \frac{5A_1}{2G} + \frac{5A_1A_2}{2G^2} \right. \\ &\quad \left. + \frac{5A_1A_2}{G^3} \right) \frac{T_{RT} + T_{amp}}{G_{HEMT}(T_{device} + T_{HEMT})} \end{aligned} \quad (A6)$$

In normal operation we require $A_{1,2} \leq G$. So we can set $A_{1,2} = G$ to determine an upper bound on $S_{\delta\phi,RT}$:

$$\frac{S_{\delta\phi,RT}}{S_{\delta\phi,HEMT}} \leq \left(\frac{15}{2} + \frac{5}{G} \right) \frac{T_{RT} + T_{amp}}{G_{HEMT}(T_{device} + T_{HEMT})} \quad (A7)$$

For $T_{RT} = 290$ K, $T_{amp} = 438$ K, $T_{device} = 0.1$ K, $T_{HEMT} = 2.3$ K, $G_{HEMT} = 40$ dB, and $G = 15$ dB, $\frac{S_{\delta\phi,RT}}{S_{\delta\phi,HEMT}} \leq -6.3$ dB.




FIG. 12. Schematic of room temperature amplifier/attenuator chain. An RF power splitter is used to send the full 4 – 8 GHz output comb to both the 4 – 6 GHz and 6 – 8 GHz readout units. The amplifier gain G is 15 dB and $0 \leq A_{1,2} \leq 31.75$ dB. In practice, we require $A_{1,2} \leq 15$ dB.

Appendix B: Correcting for Line Noise

According to the optimal filtering formalism, the pulse amplitude estimator variance is given by¹⁹:

$$\sigma^2 = \left(\sum_k \frac{|\hat{s}(f_k)|^2}{J(f_k)} \Delta f \right)^{-1} \quad (\text{B1})$$

where k indexes frequency, Δf is the frequency spacing, $\hat{s}(f_k)$ is the DFT (discrete Fourier transform) of the pulse template, and $J(f_k)$ is the noise PSD.

Consider some $J(f_k)$ with significant spectral lines. We can remove these lines – denote this $J'(f_k)$ – then compute a new pulse height variance $\sigma^2[J'(f_k)]$, where we’ve written σ^2 as a function of $J(f_k)$. Define:

$$C = \frac{\sigma^2[J(f_k)]}{\sigma^2[J'(f_k)]} \quad (\text{B2})$$

Since σ^2 scales proportionally with $J(f_k)$,

$$\sigma^2[J(f_k)] = C \sigma^2[J'(f_k)] = \sigma^2[CJ'(f_k)] \quad (\text{B3})$$

$CJ'(f_k)$ is a “clean” spectrum with the same σ^2 as $J(f_k)$, so we define the noise floor of $CJ'(f_k)$ as the “true” channel noise floor after accounting for spectral lines. Since $J'(f_k)$ has the same noise floor as $J(f_k)$, we can perform this correction by scaling the fitted noise floor of $J(f_k)$ by C .

To compute $J'(f_k)$, we set all spectral lines that are 3 dB above the fitted noise floor to the local spectral floor. This correction is performed over 500 Hz to 300 kHz band; it does not account for the $1/f$ “knee” present in some channels. The IQ low pass filter rolls off transmission significantly past 300 kHz, so it is not necessary to remove lines in this region.

In order to keep this calculation consistent with the optimal filtering scheme used in firmware, we calculate σ^2 directly from the optimal filter and pulse template. In the frequency domain, the optimal filter is given by:

$$\hat{\phi}(f_k) = \frac{\hat{s}^*(f_k)}{J(f_k)} \quad (\text{B4})$$

Substituting eqn. B4 into B1, we have:

$$\sigma^2 = \left(\sum_k \hat{\phi}(f_k) \hat{s}(f_k) \Delta f \right)^{-1} \quad (\text{B5})$$

We can convert this to the time domain using the convolution theorem (up to a normalization constant, since we are only interested in the ratio $\sigma^2[J(f_k)]/\sigma^2[J'(f_k)]$):

$$\begin{aligned} \sigma^2 &= \left(\sum_k \hat{\phi}(f_k) \hat{s}(f_k) e^{if_k t} \Delta f \right)^{-1} \Big|_{t=0} \\ &= iFFT[\hat{\phi}(f_k) \hat{s}(f_k)](t) \Big|_{t=0} \\ &= [\phi * s](t) \Big|_{t=0} \end{aligned} \quad (\text{B6})$$

where $*$ denotes the convolution operator. We use equation B6 to compute all σ^2 , using a 50 coefficient optimal filter and representative pulse template (figure 5). To avoid potential periodicity artifacts inherent to DFT, optimal filters are computed in the time domain¹⁴.

Appendix C: Noise Contributed by Sideband Tones

Consider a sideband reflection falling within 100 kHz (inside the channel bandwidth) of some resonator tone. The power contributed by the sideband reflection to the resonator tone phase power spectrum (in units dBc) is given by:

$$P_{SB,\delta\phi} = \frac{1}{2} \frac{P_{SB}}{P_{tone}} \quad (\text{C1})$$

where P_{tone} resonator tone power and P_{SB} is the sideband reflection power. The corresponding RMS amplitude (in radians) is given by:

$$A_{SB,\delta\phi} = \sqrt{P_{SB,\delta\phi}} = \sqrt{\frac{P_{SB}}{2P_{tone}}} \quad (\text{C2})$$

For $\frac{P_{SB}}{P_{tone}} = -30$ dB, $A_{SB,\delta\phi} = 1.3^\circ$. This is below the RMS phase noise for a typical resonator channel; the channel in figure 7 has $A_{\delta\phi,RMS} \approx 2^\circ$. Depending on the location of the sideband reflection relative to the in-band tone, this can likely be further reduced by optimal filtering.

¹P. Szypryt, S. R. Meeker, G. Coiffard, N. Fruitwala, B. Bumble, G. Ulbricht, A. B. Walter, M. Daal, C. Bockstiegel, G. Collura, N. Zobrist, I. Lipartito, and B. A. Mazin, *Opt. Express* **25**, 25894 (2017).

²M. J. Strader, A. M. Archibald, S. R. Meeker, P. Szypryt, A. B. Walter, J. C. van Eyken, G. Ulbricht, C. Stoughton, B. Bumble, D. L. Kaplan, and B. A. Mazin, “Search for optical pulsations in PSR J0337+1715,” *Monthly Notices of the Royal Astronomical Society* **459**, 427–430 (2016), <https://academic.oup.com/mnras/article-pdf/459/1/427/8115535/stw663.pdf>.

- ³P. Szypryt, G. E. Duggan, B. A. Mazin, S. R. Meeker, M. J. Strader, J. C. van Eyken, D. Marsden, K. O'Brien, A. B. Walter, G. Ulbricht, T. A. Prince, C. Stoughton, and B. Bumble, "Direct detection of SDSS J0926+3624 orbital expansion with ARCONS," *Monthly Notices of the Royal Astronomical Society* **439**, 2765–2770 (2014), <https://academic.oup.com/mnras/article-pdf/439/3/2765/3850573/stu137.pdf>.
- ⁴B. A. Mazin, S. R. Meeker, M. J. Strader, P. Szypryt, D. Marsden, J. C. van Eyken, G. E. Duggan, A. B. Walter, G. Ulbricht, M. Johnson, B. Bumble, K. O'Brien, and C. Stoughton, *Publ. Astron. Soc. Pac.* **125**, 1348–1361 (2013).
- ⁵S. R. Meeker, B. A. Mazin, A. B. Walter, P. Strader, N. Fruitwala, C. Bockstiegel, P. Szypryt, G. Ulbricht, G. Coiffard, B. Bumble, G. Cancelo, T. Zmuda, K. Treptow, N. Wilcer, G. Collura, R. Dodkins, I. Lipartito, N. Zobrist, M. Bottom, J. C. Shelton, D. Mawet, J. C. van Eyken, G. Vasisht, and E. Serabyn, *Publ. Astron. Soc. Pac.* **130**, 065001 (2018).
- ⁶A. Walter, N. Fruitwala, S. Steiger, J. I. Bailey, N. Zobrist, N. Swimmer, I. Lipartito, J. P. Smith, S. R. Meeker, C. Bockstiegel, G. Coiffard, R. Dodkins, P. Szypryt, K. K. Davis, M. Daal, B. Bumble, G. Collura, O. Guyon, J. Lozi, S. Vievard, N. Jovanovic, F. Martinache, T. Currie, and B. Mazin, "The MKID Exoplanet Camera for Subaru SExAO," *Publications of the Astronomical Society of the Pacific* (Accepted).
- ⁷B. Mazin, P. Day, J. Zmuidzin, and H. LeDuc, "Multiplexable kinetic inductance detectors," in *AIP Conference Proceedings*, Vol. 605 (American Institute of Physics, 2002) pp. 309–312.
- ⁸S. McHugh, B. A. Mazin, B. Serfass, S. Meeker, K. O'Brien, R. Duan, R. Raffanti, and D. Werthimer, *Review of Scientific Instruments* **83**, 044702 (2012).
- ⁹P. Day, H. Leduc, B. Mazin, A. Vayonakis, and J. Zmuidzin, "A broadband superconducting detector suitable for use in large arrays," *Nature* **425**, 817–21 (2003).
- ¹⁰D. Werthimer, "The casper collaboration for high-performance open source digital radio astronomy instrumentation," in *2011 XXXth URSI general assembly and scientific symposium* (IEEE, 2011) pp. 1–4.
- ¹¹J. Hickish, Z. Abdurashidova, Z. Ali, K. D. Buch, S. C. Chaudhari, H. Chen, M. Dexter, R. S. Domagalski, J. Ford, G. Foster, *et al.*, "A decade of developing radio-astronomy instrumentation using casper open-source technology," *Journal of Astronomical Instrumentation* **5**, 1641001 (2016).
- ¹²There is a fixed 3 dB of attenuation in addition to the programmable attenuators, so the power at the end of the chain is given by $P_{out} = P_{in} + 60 \text{ dB} - A_{prog} - 3 \text{ dB}$.
- ¹³S. H. Moseley, R. L. Kelley, R. J. Schoelkopf, A. E. Szymkowiak, D. McCammon, and J. Zhang, "Advances toward high spectral resolution quantum x-ray calorimetry," *IEEE Transactions on Nuclear Science* **35**, 59–64 (1988).
- ¹⁴B. K. Alpert, R. D. Horansky, D. A. Bennett, W. B. Doriese, J. W. Fowler, A. S. Hoover, M. W. Rabin, and J. N. Ullom, "Note: Operation of gamma-ray microcalorimeters at elevated count rates using filters with constraints," *Review of Scientific Instruments* **84**, 056107 (2013), <https://doi.org/10.1063/1.4806802>.
- ¹⁵N. Zobrist, B. H. Eom, P. Day, B. A. Mazin, S. R. Meeker, B. Bumble, H. G. LeDuc, G. Coiffard, P. Szypryt, N. Fruitwala, I. Lipartito, and C. Bockstiegel, "Wide-band parametric amplifier readout and resolution of optical microwave kinetic inductance detectors," *Applied Physics Letters* **115**, 042601 (2019), <https://doi.org/10.1063/1.5098469>.
- ¹⁶J. Gao, J. Zmuidzin, B. A. Mazin, H. G. LeDuc, and P. K. Day, "Noise properties of superconducting coplanar waveguide microwave resonators," *Applied Physics Letters* **90**, 102507 (2007).
- ¹⁷J. Gao, M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzin, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc, *Appl. Phys. Lett.* **92**, 152505 (2008).
- ¹⁸J. Zmuidzin, "Superconducting microresonators: Physics and applications," *Annual Review of Condensed Matter Physics* **3**, 169–214 (2012).
- ¹⁹S. R. Golwala, *Exclusion limits on the WIMP nucleon elastic scattering cross-section from the Cryogenic Dark Matter Search*, Ph.D. thesis, UC, Berkeley (2000).
- ²⁰M. Strader, *Digital Readout for Microwave Kinetic Inductance Detectors and Applications in High Time Resolution Astronomy*, Ph.D. thesis, University of California Santa Barbara (2016).