

A Prototype System for SQUID Multiplexing of Large- Format Transition-Edge Sensor Arrays

C.D. Reintsema, J. Beyer, S.W. Nam, S. Deiker, G.C.
Hilton, K.D. Irwin, J.M. Martinis, J.N. Ullom, L.R. Vale

Mail Stop 814.03

National Institute of Standards and Technology

325 Broadway

Boulder, CO 80305

M. MacIntosh

UK Astronomy Technology Centre

Royal Observatory, Blackford Hill

Edinburgh, Scotland EH9 3HJ

Review of Scientific Instruments, in press

Submitted March, 2003

A Prototype System for SQUID Multiplexing of Large-Format Transition-Edge Sensor Arrays*

Carl D. Reintsema, Jörn Beyer, Sae Woo Nam,

Steve Deiker, Gene C. Hilton, Kent Irwin, John Martinis, Joel Ullom,

Leila R. Vale

*Electromagnetic Technology Division, National Institute of Standards & Technology, Boulder,
Colorado, 80305*

Mike MacIntosh

*UK Astronomy Technology Centre, Royal Observatory, Blackford Hill, Edinburgh, Scotland,
EH9 3HJ*

ABSTRACT

We discuss the implementation of a time-division SQUID multiplexing system for the instrumentation of large-format transition-edge sensors (TES) arrays. We cover the design and integration of cryogenic SQUID multiplexers and amplifiers, signal management and wiring, analog interface electronics, a digital feedback system, serial-data streaming and management, and system configuration and control. We present data verifying performance of the digital-feedback system. System noise and bandwidth measurements demonstrate the feasibility of adapting this technology for a broad base of applications, including x-ray materials analysis and imaging arrays for future astronomy missions such as Constellation-X (x-ray) and the SCUBA-2 instrument (submillimeter) for the James Clerk Maxwell Telescope.

* Contribution of the National Institute of Standards and Technology, not subject to copyright.

1. INTRODUCTION

The development of detectors based on superconducting transition-edge sensors has led to a variety of innovative approaches for the detection of photons in the wavelength range of millimeter-wave through γ -ray. The low noise, low power, and low impedance of superconducting quantum-interference devices (SQUIDs) make them ideal amplifiers for first-stage signal conditioning. Independent SQUID-coupled TES elements are adequate for the readout of small numbers of detectors. However, the readout of large-scale, two-dimensional arrays of cryogenic detectors presents several challenges. Both wire count and power dissipation scale with pixel count in a non-multiplexed, conventional layout. Multiplexing can be used to reduce these and other problems associated with scaling.

To multiplex an array of detectors, each element is coupled to a dedicated first-stage SQUID amplifier. The first-stage amplifiers are configured in a column format. The signals from all SQUIDs in the column are summed and routed to a second-stage amplifier. Only one first-stage SQUID is on at a time, and so the signal from the corresponding detector is presented to the second-stage amplifier. Since the other first-stage SQUIDs remain in a superconducting state, they contribute no signal or noise and dissipate no power. By multiplexing the first-stage amplifiers, only a single output channel is required per column of detectors. The address lines for a row of first-stage SQUIDs are wired in series, so only one set of address lines is required per row of the array. Custom room-temperature digital electronics are used to process signals from each column, control the timing of the row multiplexing, and apply a switched feedback signal to a common feedback line for each column.

Applications for this technology range from materials analysis¹ to astronomical imaging^{2,3}. The FIBRE instrument, a Fabry-Pérot interferometer operating in wavelength bands centered at 350 μm and 450 μm and using an array of eight SQUID-multiplexed TESs, has been demonstrated in an astronomical application.⁴

2. SYSTEM DESIGN

An overall block diagram of the system we have designed, implemented, and characterized is shown in Fig. 1. The discussion of the system design will begin at the first-stage amplifiers (left side of the figure), proceed down the electronic signal chain, and end with the host computer.

2.1 Cryogenic multiplexer and amplifier chain

The first element in the signal path after the detector is the SQUID multiplexer (SMUX). A schematic of the multiplexing circuit is shown in Fig. 2. Each column represents a separate multiplexer with M first-stage input SQUIDS. The N SQUIDS in each row are wired in series and addressed (biased) with time-sequential boxcar-modulated current drives. Second- and third-stage amplifiers and flux offset lines are shown as well (one set for each column). The first- and second-stage amplifiers are integrated on a single chip that is situated on the detector mount at base temperature (typically <100 mK). The third-stage series-array SQUID amplifier is located on the 4 K base plate of the cryostat. A superconducting wiring harness carries signals between these thermal stages.

A 1×32 SQUID multiplexer was designed and fabricated at NIST. A detailed description of the design and performance of this multiplexer⁵ and other similar preliminary designs,⁶ can be found in the literature. A photograph of a portion of the integrated circuit is shown in Fig. 3. The chip has 32 first-stage SQUID amplifiers. The input coils to the SQUID terminate in pads at the bottom of the chip to which the TESs from each corresponding pixel in a column of the array are wire bonded. The first-stage SQUIDS are biased (via contacts at the top) each in parallel with a $\sim 1 \Omega$ resistor and in series with a transformer coil. In a multiplexing configuration the first-stage SQUIDS are sequentially addressed (biased) and sampled. The transformer coils (the smaller coils visible across the top of the circuit) couple the SQUID signal into a summing loop common to all 32 first-stage SQUIDS. The summing coil couples to the input of a single second-stage SQUID.

The integrated second-stage SQUID amplifier averts the need to couple the first stage output directly to a SQUID series-array amplifier at 4 K. While direct transformer coupling could be implemented, a superconducting loop to 4 K would require careful magnetic shielding efforts to prevent coupling of stray fields into the measurement circuit, and would reduce the efficiency of the summing coil due to parasitic inductance.

The terminals of the second-stage SQUID are connected, via superconducting leads, to the input coil of a SQUID series-array amplifier in series with a bias resistor (which breaks the superconducting loop), both of which are situated on the 4 K cold stage. The series-array amplifiers have been developed⁷ and refined⁸ at NIST over the last decade. The series-array amplifier currently in use consists of 100 SQUIDs arrayed in series and sharing a common input coil.

2.2 Analog interface electronics

An electromagnetically shielded housing, or tower (see Fig. 4), on the cryostat contains several custom electronic cards that interface with the cryogenic electronics. The four varieties of tower cards are: power, preamplifier, bias, and feed-through. A single power card distributes analog and digital power from low-noise linear supplies or battery sources over a back plane to the other cards in the tower. In addition to filtering and regulation, this card also has a field programmable gate array (FPGA) used to decode optical serial communication from the host computer. The preamplifier card provides a low-noise ($1.1 \text{ nV/Hz}^{1/2}$), gain 100, 5 MHz bandwidth, room temperature amplifier for the voltage signal from the series-array SQUID amplifier. Each channel on the preamplifier card also has a programmable bias source to drive the series arrays. The series-array bias-current level is set using the optical addressing. Bias cards provide adjustable voltages to set biases and flux offsets for the SQUID amplifiers and also to bias detectors. Feed-through cards are used to pass signals into and out of the cryostat. All of the analog interface cards for the tower are designed for 8 channels.

2.3 Digital-feedback electronics

The core of the room-temperature electronics is a digital-feedback (DFB) card developed at NIST. The card digitizes a multiplexed signal from a column of the array and maintains a flux-locked loop through application of a feedback signal to a common feedback coil for the first-stage SQUID amplifiers of a single column. The card has three primary components: a 12-bit analog-to-digital converter (ADC) to sample the input signal; a FPGA that applies the switched-feedback algorithm in a multiplexed mode; and a 14-bit digital-to-analog converter (DAC) that provides the feedback signal to the first-stage SQUID amplifiers. Both the ADC and DAC are clocked at 50 MHz.

Significant effort has been focused on developing firmware for handling the multiplexed feedback. The kernel implements a discrete proportional/integral feedback algorithm that can be expressed (in a simplified form) as

$$y_n = P_{user} x_n + I_{user} \sum_{i=0}^{n-1} x_i, \quad (1)$$

where y_n is the feedback signal, x_n is the error signal, and P_{user} and I_{user} are software-configurable gain constants.

The error signal, x_n , is the difference between the ADC-sampled input, accumulated and averaged over a user-defined number of samples ($NSAMP$), and the ADC lock point. The feedback signal is the output from a multiply-and-accumulate cell (26-bits) that implements the function expressed in equation 1, truncated to 14-bit word size to drive the DAC output.

There are 3 clock signals: the master clock (CLK), the line clock ($LSYNC$), and the frame clock (FRM). Figure 5 shows the detailed timing for multiplexing. The master clock is generated by an oscillator on a clock card and distributed electrically to the other digital-feedback system cards. The master clock rate is 50 MHz (clock period 20 ns). The $LSYNC$ clock is also generated and distributed by the clock card, using a

simple divide-by- n counter with $n_{min}=32$. The maximum line (row) sampling rate is thus $50 \text{ MHz}/32 \sim 1.56 \text{ MHz}$. The line period $LSYNC$ is software configurable. The frame clock is derived from the $LSYNC$ clock in a similar manner by dividing down by the number of rows to multiplex $NMUX$ (a frame is one cycle through all of the rows). This firmware timing logic is embedded in a row-control module that keeps track of the row indexing for multiplexing. The FRM clock is synchronous with the reset of the row counter, i.e. a pulse is generated for only the first row of a frame. The FRM bit is encoded in the data stream to allow software demultiplexing of archived data.

The multiplexing firmware is designed to use a pipelined architecture for sampling and determination of error signals and computation of feedback signals, in combination with memory allocation, storage, and indexing for row-specific parameters. Memory is allocated for the following parameters for each row: P_{user} , I_{user} , ADC offset (lock point), DAC offset (feedback flux offset), error signal, and feedback signal. For the purposes of the following discussion, let f represent the frame index and n represent the row index within a frame. Following a line-synchronizing trigger pulse for the n th row, the algorithm latches the DAC output to the feedback value from the last computation cycle (previous frame) for the indexed row $y_{f-1,n}$, waits a user-defined number of master clock cycles for settling ($SETTLE$), and then samples (at the master clock rate CLK) and averages the ADC input a user-defined number of times ($NSAMP$) before subtracting the ADC offset to get the error signal $x_{f,n}$. In parallel, the algorithm computes the feedback signal for the previous row $y_{f,n-1}$ based on the pipelined error signal $x_{f,n-1}$, and writes this value to memory.

Data streaming to the host computer utilizes a fiber-optic link to serially stream the error and feedback signals in real time. The 14-bit feedback and 16-bit error signals are packaged with overflow and data-framing bits and streamed at 50 MHz. Each DFB card has a direct fiber link to one of 16 input channels on a NIST-developed PCI card. The PCI card handles parallel decoding, memory management, and data-archiving. The PCI card also has a fiber link to the clock card to receive the $LSYNC$ signal for data-synchronization purposes. Multiple PCI cards can be configured on the PCI bus in order to handle column counts higher than 16 for larger-format arrays. Limitations of PCI bus rate and memory allocation will ultimately set the maximum data-transfer rate for large-format arrays.

In order to multiplex, the first-stage SQUID amplifiers on the multiplexer chip must be row-addressed sequentially and synchronized with the DFB, a task that is handled by an eight-channel address driver card (designated RA8). This card has eight DACs, one dedicated for each of eight output channels, and an FPGA. The firmware controls the DAC outputs (which are connected to the row bias lines), switching each one between an ON and OFF state, coincident with a specific row address. With this approach one RA8 card is needed for each eight rows of the detector array. Each channel on each card is programmatically assigned a row index and the same row-control logic used by the DFB controls the ON/OFF state of the DACs. Configurable parameters of the address line driver include the *LEVEL* of the ON/OFF states, the duration of the on state *WIDTH*, and an adjustable *DELAY* between the *LSYNC* pulse edge and the latching of the ON level to the DAC (see Fig. 5). A cryogenic CMOS circuit is also being developed to implement the address-line driver at 4 K.⁹

Our discussion thus far has focused on the operation of a single channel of digital feedback. Implementing a two-dimensional array requires a DFB card for each column of the array. One crate can accommodate up to 20 cards, one of which must be a clock card. As shown in Fig. 6, an 8x8 array can be fully instrumented and multiplexed using a single crate with ten cards (1 clock, 8 DFB, 1 RA8) and a host computer with one PCI card. For larger arrays that require card counts in excess of twenty, multiple crates (with synchronized timing), and a host PC with multiple PCI cards are required. The system design accommodates this capability.

Control and configuration of the digital and analog electronics is accomplished via optical serial communication between various components and the host computer. At the transmitter end, an electrical-to-optical adapter is connected to the RS232 serial port of the host computer. At the receiving end, an optical receiver in combination with a FPGA decodes the optical signal on the various boards. A variety of registers have been defined to be software-programmable over this optical link, allowing the user to set counts, rates, gains, biases, lock points, offsets, and other system parameters.

3. SYSTEM INTEGRATION

3.1 Cryogenic platform & detector mount

The cryogenic system is comprised of a conventional cryostat with a two-stage adiabatic demagnetization refrigerator (ADR). A base temperature of 50-60 mK is routinely achieved in this system. The detector mount is machined from a block of oxygen-free high-conductivity (OFHC) copper, and attaches to a cold finger protruding from the inner salt pill of the ADR. The mount is designed to multiplex pixels in a one-dimensional format. As seen in Fig. 7a, wire bonds are used to connect the 1x32 SMUX chip to the detector chips and the circuit board. Due to the sensitivity of SQUID amplifiers and TESs to magnetic flux, substantial measures are taken to shield the amplifiers and sensors from background fields. The SMUX is backed by a thin film of niobium that provides a superconducting ground plane and magnetic shield. Further, the detector mount is placed inside a double-walled cylinder composed of one shield machined from niobium and a second fabricated from an alloy of high-magnetic-permeability. These shields are thermally anchored to the 4 K base plate. Future designs for this enclosure will assure a light-tight feed-through to prevent stray 4 K radiation from illuminating the TES. The signals to/from the SMUX are routed in and out of the detector enclosure on the circuit board. A niobium-titanium twisted-pair wiring harness carries signals between the circuit board on the detector stage and the 4 K stage. The harness is thermally grounded on a 1 K intermediate stage to minimize heat load on the detector stage. In order to further minimize stray inductance and heat load due to wiring while maximizing lead count, we intend to replace the twisted pairs with niobium-on-Kapton flexible-substrate wiring currently under development at NIST.

Signals on the 4 K stage are routed as seen in Fig. 7b. Termination and breakout for up to eight copper-nickel flex strips (see next section) are visible on the semi-circular circuit board. This arrangement provides 64 signal pairs. Five of the flex terminations (40 pairs) are routed as microstriplines to connectors for configurable jumper connection to the detector mount. The remaining three terminations are routed to the SQUID series-array module. For each series array there are two bias leads, two feedback leads, and two input leads. The input leads are wired in a loop (see Fig. 2) with a bias resistor and the second-stage SQUID on the SMUX chip. As mentioned above, the series-array amplifiers must be well shielded

magnetically. We have designed a magnetically shielded enclosure similar to that described earlier that accommodates up to eight series arrays, individually mounted on miniature circuit boards. This assembly is visible in Fig. 7b as the rectangular metal box in the upper left, and in more detail in Fig. 7c.

3.2 Flex wiring

Flexible, multi-conductor, copper-nickel wiring harnesses provide a high density, low thermal conductivity, high bandwidth, controlled-impedance transmission path between the room-temperature electronics and the 4 K stage. Each flexible harness is fabricated from two sheets of copper-nickel alloy laminated to opposite sides of a Kapton carrier. The composite structure is patterned using conventional lithographic and etch processes. Each lead pair is composed of a narrow signal line over a wider, isolated ground plane. The asymmetric geometry was chosen primarily to minimize electromagnetic cross talk between adjacent pairs. The flex strips are eight pairs wide, matching the bus width of the analog interface tower cards. Eight strips of flex are epoxy-potted in a single vacuum feed-through flange, providing 64 signal pairs per assembly. We currently produce flex in lengths of 36 inches and 40 inches. The cold end of an integrated flex assembly is pictured in Fig. 8.

3.3 Electronics Configuration

A rack-mounted 3U crate hosts the clock card, the digital-feedback cards, and the address cards. Analog and digital power are supplied to this crate from low-noise linear power supplies. The crate back plane filters and distributes power as well as distributing timing signals from the clock card. Coaxial cables with SMB terminations connect the analog tower cards to the digital electronics.

In a typical multiplexing configuration, a pre-amp card biases and monitors up to 8 channels of series-array amplifiers. The amplified series-array voltage output signal from each column is routed to the input of a digital-feedback card. A bias card adjusts the bias of the second-stage amplifiers. Additional bias cards can apply a flux offset in the series array and/or second-stage feedback coils. The host computer sets all control currents and voltages over the serial fiber link. Each address-driver card biases a row of first-stage SQUIDs

through one of the eight channels on a tower feed-through card. Finally, the output from the DFB cards is wired, via a feed-through card, to the first-stage feedback coil of the corresponding column SMUX.

3.4 Software

The system runs on a LINUX platform utilizing software written and developed in the PYTHON programming language. A variety of low-level programs can be used to set analog biases on tower cards, change the timing of the line sync on the clock card, and set the bias levels and relative timing of the address cards. Higher-level programs control the digital-feedback functions. These programs typically include global-variable configuration for items such as number of rows, number of samples, settling time, and local registers for row-specific variables such as P and I parameters, lock point, and feedback open/closed state. There are also programs that address the PCI card for streaming data. Currently we stream blocks of data for analysis offline. Real-time analysis and automated amplifier-tuning software are currently under development.

4. EXPERIMENTAL RESULTS

The performance and behavior of the digital-feedback system have been analyzed in detail. Room-temperature tests using resistor networks show predictable gain and bandwidth behavior in agreement with calculations of the feedback algorithm. Cryogenic tests at 4 K verify correct multiplexing behavior, including synchronization with the address driver cards. In these tests, emulated input signals are injected into the SMUX input coils with function generators and the signals are recovered from demultiplexed data. Cryogenic tests at 60 mK demonstrate correct noise scaling in the demultiplexed signals.

4.1 Closed loop bandwidth & linearity

Consideration of the P-I algorithm shows that the closed loop bandwidth can be expressed as

$$f_{3dB} = \frac{1}{2\pi t_{clk}} \left(\frac{I_{user}}{512} \right) \left(\frac{NSAMP}{16} \right) \left(\frac{1}{LSYNC \times NMUX} \right) \left(\frac{g_{input}}{g_{fb}} \right), \quad (2)$$

where $I_{user}/512$ is the normalized integration constant, $t_{clk}=20$ ns, $NSAMP$ is the number of samples, $LSYNC$ is the line period, $NMUX$ is the number of channels (rows) to multiplex, and the gain terms are for the input gain and feedback gain. The input gain for the measurements, $g_{input}=2.5$, is determined by a low-noise amplifier on the DFB card. The feedback gain is set by the resistance ratio of the room-temperature feedback network. Table 1 lists the measurements of the closed-loop bandwidth.

As can be seen from the data, the model is in good agreement with experiment (with the exception of the final row). The data with $NMUX=1$ were acquired with a non-multiplexing version of the firmware that uses the same kernel for the P-I computation and hence also obeys equation 2. The discrepancy in the final row is a result of the omission of dead time from the model. Dead time is the delay between when the error signal is measured and when the feedback signal is applied. In a digital system such as this it includes both propagation delays due to wiring and a digital delay of one frame period due to the feedback algorithm. Dead time causes peaking in the closed-loop gain as f_{3dB} approaches the sampling frequency. This criterion¹⁰ can be stated explicitly as $f_{3dB} > 0.08 / [NMUX * LSYNC * t_{clk}]$. This effect manifests itself as an apparent increase in bandwidth, in agreement with our observations.

Measurements of the linearity of a SQUID multiplexer and series array on a 4 K dip probe are used to ascertain whether the feedback algorithm or other systematics dominate the system nonlinearity. The results of these measurements are shown in Fig. 9, which shows the total harmonic distortion (THD) and the $2f$ component versus the ADC lock point. A dynamic signal analyzer monitors the DFB output while an ultra-low-distortion excitation source is applied to the first-stage SQUID amplifier input ($f_s=3.55$ kHz, amplitude modulation of $3\phi_0$ at the first-stage SQUID). The lock point is plotted as direct ADC units (14 bit effective full scale). Shown for comparison is the behavior that theory predicts considering the inherent nonlinearity of a SQUID amplifier in a flux-locked loop. The theory curve is for the $2f$ component of the THD and is based on the Taylor-series expansion of a sinusoidal $V(\phi)$ relation about the operating point.¹⁰

The THD is minimized for ADC lock point ~ 7000 , which coincides with the inflection point in the composite SQUID amplifier $V(\phi)$ curve. As the lock point is moved either way from this point the harmonic distortion increases dramatically. The two curves demonstrate that the majority of the THD is from the $2f$ component, except at the inflection point where the $2f$ component goes to zero, in agreement with theory. This measurement shows that the system nonlinearity is dominated by the SQUID amplifier chain and not by the digital-feedback electronics or algorithm (at the given excitation frequency).

4.2 Noise scaling

The dominant sources of noise in the system are SQUID-amplifier noise, room-temperature amplifier noise, and digital-electronics noise. Since this is a system paper, we neglect discussion of detector noise. Detailed discussions of noise in TES detectors can be found elsewhere¹¹. The system design goal is to be detector-noise limited over the signal bandwidth of interest with sufficient noise margin to allow many channels to be multiplexed. In our measurements, which are presented below, we show that the SQUID-amplifier noise dominates and that we observe the expected noise-scaling behavior with system-parameter variations.

Since we operate our system in a multiplexing mode, aliasing of high-frequency noise into the signal bandwidth is unavoidable. In addition, the decay of transients associated with switching events can further limit the system performance. The bandwidth of the amplifier chain must be sufficient to allow transients to decay on a time scale short compared to the sampling interval. The result is that the bandwidth of the amplifier chain is the dominant system pole affecting the aliasing. However, the results underscore that the interpretation of noise scaling behavior is nontrivial.

The maximum line-sampling frequency is limited by the fiber-optic data link between the digital-feedback board and the PCI card. During each line sample, the feedback and error signals from the previous line are sent to the computer. Both the feedback and error signals require 16 clock cycles. The firmware sends the feedback and error signals serially on one fiber-optic link, requiring 32 cycles, each taking 20 ns, for a total

of 640 ns minimum line time (a maximum line rate of 1.5625 MHz). This sample time is not limited by the ~ 3 MHz bandwidth of the multistage SQUID plus room-temperature amplifier chain⁵. Applications that require faster sampling rates can be accommodated with reasonable system modifications. The line-sampling frequency could be increased by a factor of two (to 3.125 MHz) by modifying the firmware to either drop the transmission of the error signal, or to stream the error and feedback signal in parallel on the two fiber-optic links. The line-sampling frequency could be further increased by modifying the hardware to increase the clock rate of the master timing chip, or by running the fiber-optic data link on a faster clock. The bandwidth of the amplifier chain would need to be increased to shorten the transient decay (i.e., shorten the *SETTLE* time). Work is underway to replace the room-temperature analog preamplifier with a higher bandwidth circuit. The cryogenic SQUID MUX circuit itself will also be modified to increase its bandwidth.⁵

Equation 3 shows the expected scaling from a system-noise source. Terms attributable to aliasing and averaging are on the left. The expression is rearranged with respect to system-dependent terms on the right.

$$S_{system}^{1/2} \propto \underbrace{\sqrt{\frac{BW_{noise}}{f_s}}}_{\text{aliasing}} \underbrace{\sqrt{\frac{1}{NSAMP}}}_{\text{averaging}} \propto \underbrace{\sqrt{\frac{LSYNC \times NMUX}{NSAMP}}}_{\text{configuration (1/duty cycle)}^{1/2}} \quad (3)$$

In this equation, f_s represents the sampling frequency (frame rate). This equation is accurate as it applies to the scaling of white noise. However, since detector and amplifier noise have a more structured spectrum when referred to the ADC input (i.e., the detector noise is rolled off by the L/R filter on the SQUID input and the SQUID amplifier noise is rolled off by the pole of the amplifier chain), the interpretation of noise spectra from sources farther up the electronic sampling chain is more complicated.

The noise contribution from the low-distortion differential ADC driver used in our design has been determined to be ~ 60 nV/Hz^{1/2} (referred to the ADC input). The noise of the room-temperature preamplifier stage is 1.1 nV/Hz^{1/2}, or ~ 110 nV/Hz^{1/2} when referred to the ADC input. The corresponding referred input noise of the SQUID amplifier is 2 μ V/Hz^{1/2} at 4 K. This value is dominated by the Johnson noise of the damping shunt resistors across the first stage SQUID amplifier input coil. This scales with temperature to a 280 nV/Hz^{1/2} at 80 mK¹². Hence, at operational temperatures, the system is SQUID-noise limited over the measurement bandwidth. The system noise floor in the measurement band is slightly higher than the SQUID amplifier noise floor due to aliasing effects.

We have made a series of measurements of input SQUID noise (at ~ 80 mK) that confirm the correct noise scaling behavior. Figure 10a shows the measured flux-noise-spectral-density as a function of sampling frequency. The three sets of data are for varying *NMUX*, *NSAMP*, and *LSYNC* independently. Note that for the variable *NSAMP* data set, all points have the same f_s since *NSAMP* does not affect f_s (see Fig. 2). For the variable *NMUX* and *LSYNC* data sets the noise scales as $1/f_s^{1/2}$, consistent with equation 3. The measurement is dominated by first stage SQUID noise over these sampling frequencies and changing f_s affects aliasing of the first stage SQUID noise (white) into the measurement bandwidth.

As shown in Fig. 10b, the $1/NSAMP^{1/2}$ scaling predicted by equation 3 is incorrect. This is because averaging low numbers of samples affects frequency components well above both the sample frequency and the amplifier chain roll-off. In this frequency regime the noise floor is dominated by lower-amplitude system-noise sources. While the noise aliased into the measurement band is changing in the expected manner, the effect on the total noise is suppressed (i.e., total noise is dominated by the SQUID noise.) A power-law fit to the data points yields $S_V^{1/2} \propto 1/NSAMP^{0.1}$. Applying this scaling law to the data sets for varying *LSYNC* (taken for *NSAMP*=4) and *NMUX* (taken for *NSAMP*=16) accounts directly for the observed offset in Fig. 10a.

4.3 Multiplexing

We have run the system configured as a 1 x 8 multiplexer at 4 K. Input signals were emulated using function-generator signals applied across the first-stage SQUID input coils. Sine waves at frequencies of 300 Hz to 2.4 kHz and amplitude 13 μ A ($2\phi_0$ amplitude modulation at first-stage SQUID) were applied and sampled in a multiplexed mode with the digital-feedback system. Figure 11a shows the real-time data stream sampled from the multiplexer. Figure 11b shows the individual channel signals after demultiplexing in software. The following parameters were used for this test; $LSYNC=64$, $NSAMP=4$, $NMUX=8$, $SETTLE=56$. These settings correspond to a line rate of ~ 781 kHz and a frame rate of ~ 98 kHz.

Multiplexing tests of a more limited scope were performed in the ADR at 165 mK. In these tests, two nominally identical TES detectors were coupled to adjacent first-stage SQUID inputs of the multiplexer. The detector platform was regulated above the TES transition temperature at an operating point of 165 mK presenting a normal-state detector resistance of 15 m Ω to the SQUID inputs. The two channels were then sampled in a multiplexed mode using digital feedback. The measured voltage-noise spectra were recovered from a Fourier transform of the demultiplexed time-series data and exhibit white noise levels consistent with that of a resistor of the expected value and temperature.¹²

The results presented here underscore the feasibility of scaling this technology to large-format arrays of cryogenic detectors. We are currently developing detectors and SMUX for a number of different large-format arrays. These include a 32-channel x-ray microcalorimeter array for materials analysis¹³, a 32 \times 32 x-ray microcalorimeter array for the Constellation-X observatory³, a 288-pixel first-light far-infrared bolometer instrument for an upcoming airborne observatory (SOFIA) called SAFIRE¹⁴, and SCUBA-2^{2,15}, a 12,800-pixel submillimeter bolometer instrument to be deployed at the James Clerk Maxwell Telescope in ~ 2005 . The SCUBA-2 camera will image at 450 μ m and 850 μ m using four 40 x 40 element sub-arrays in the focal plane for each observation wavelength. The reported measurement bandwidth and flux noise is adequate for the far infrared and submillimeter imaging applications. The measured bandwidth and noise levels need further improvement for x-ray applications.

REFERENCES

- ¹ D. A. Wollman, K. D. Irwin, G. C. Hilton, L. L. Dulcie, D. E. Newbury, and J. M. Martinis, *Journal of Microscopy* **188**, 196-223 (1997).
- ² W. D. Duncan, W. Holland, D. Audley, D. Kelly, T. Peacock, P. Hastings, M. MacIntosh, K. D. Irwin, S. W. Nam, G. C. Hilton, S. Deiker, A. Walton, A. Gundlach, W. Parkes, C. Dunare, P. A. R. Ade, and I. Robson, in *Proceedings of the 9th International Workshop on Low Temperature Detectors; Vol. 605*, edited by F. S. Porter, D. McCammon, M. Galeazzi, and C. K. Stahle (AIP, Madison, Wisconsin, 2001).
- ³ C. K. Stahle, S. Bandler, T. W. Barbee, J. W. Beeman, R. P. Brekosky, B. Cabrera, M. Cunningham, S. Deiker, E. Figueroa-Feliciano, F. M. Finkbeiner, M. A. Frank, K. C. Gendreau, E. Haller, G. C. Hilton, and K. D. Irwin, *Proceedings of the SPIE* **82**, 3765 (1999).
- ⁴ D. J. Benford, T. A. Ames, J. A. Chervenak, E. N. Grossman, K. D. Irwin, S. A. De Kotwara, B. Maffei, S. H. Moseley, F. Pajot, T. G. Phillips, C. D. Reintsema, C. Rioux, R. A. Shafer, J. G. Staguhn, C. Vastel, and G. M. Voellmer, in *Proceedings of the 9th International Workshop on Low Temperature Detectors; Vol. 605*, edited by F. S. Porter, D. McCammon, M. Galeazzi, and C. K. Stahle (AIP, 2001), p. 589-592.
- ⁵ P. A. J. de Korte, J. Beyer, S. Deiker, G. C. Hilton, K. D. Irwin, M. MacIntosh, S. W. Nam, C. D. Reintsema, and L. R. Vale, *Review of Scientific Instruments* **in press** (2002).
- ⁶ J. A. Chervenak, K. D. Irwin, E. N. Grossman, J. M. Martinis, and C. D. Reintsema, *Applied Physics Letters* **74**, 4043-4045 (1999).
- ⁷ R. P. Welty and J. M. Martinis, *IEEE Transactions on Applied Superconductivity* **3**, 2605-2608 (1993).
- ⁸ M. E. Huber, P. A. Neil, R. G. Benson, D. A. Burns, A. M. Corey, C. S. Flynn, Y. Kitaygorodskaya, O. Massihzadeh, J. M. Martinis, and G. C. Hilton, *IEEE Transactions on Applied Superconductivity* **11**, 4048-4053 (2001).
- ⁹ S. H. Moseley, (2002).
- ¹⁰ D. Drung, in *SQUID sensors: fundamentals, fabrication, and applications; Vol. E329*, edited by H. Weinstock (1996), p. 63.

- ¹¹ K. D. Irwin, *Applied Physics Letters* **66**, 1998-2000 (1995).
- ¹² J. Beyer, P. A. J. de Korte, C. D. Reintsema, S. W. Nam, M. MacIntosh, G. C. Hilton, L. R. Vale, and K. D. Irwin, *IEEE Transactions on Applied Superconductivity* **in press** (2002).
- ¹³ K. D. Irwin, J. A. Beall, S. Deiker, G. C. Hilton, L. King, S. W. Nam, D. E. Newbury, C. D. Reintsema, J. A. Small, and L. R. Vale, *Microscopy and Microanalysis* **8**, 78-79 (2002).
- ¹⁴ R. A. Shafer, S. H. Moseley, P. A. R. Ade, D. J. Benford, G. Bjoraker, E. Dwek, D. A. Neufeld, F. Pajot, T. G. Phillips, and G. J. Stacey, *Proceedings of the SPIE* **4014**, 98-108 (2000).
- ¹⁵ W. Holland, W. D. Duncan, B. D. Kelly, K. D. Irwin, A. J. Walton, P. A. R. Ade, and I. E. Robson, *Proceedings of the SPIE* **4855**, 1-18 (2002).

I_{user}	$NSAMP$	$LSYNC$	$NMUX$	$f_{3dB} MODEL$	$f_{3dB} EXPERIMENT$
				[kHz]	[kHz]
1	4	16	1	12.0	13.0
3	4	16	1	36.1	39.0
8	2	64	2	6.0	6.6
8	4	64	2	12.0	13.5
8	6	64	2	18.0	24.0
8	4	64	4	6.0	6.6
24	4	32	2	72.1	120.0

Table 1 – The measured and predicted closed-loop bandwidth of the digital-feedback system for various multiplexing configurations.




Figure 1 – System block diagram.




Figure 2 – Circuit schematic showing concept for two-dimensional multiplexing. Shading represents the 1x32 SMUX chip on the detector stage at <100mK, other components are at 4 K.




Figure 3 – A photograph of a portion of the 1 x 32 channel SQUID multiplexer chip.




Figure 4 – A photograph of the tower enclosure atop the cryostat with analog interface electronics sufficient to operate an 8 x 8 array of detectors. (Cabling removed for clarity.)




Figure 5 – A detailed timing diagram of the algorithm used for digital multiplexing.




Figure 6 – A photograph of a partially filled crate of digital feedback electronics. Configuration shown is adequate for multiplexing of a 8 x 8 array of detectors; [8] digital feedback cards, [1] clock card, and [1] eight channel address line driver.




Figure 7 - Photographs of the cryostat and components: a) a close-up photograph of the detector mount showing a 1 x 32 SQUID multiplexer wire bonded to two prototype 8 x 8 detector arrays, b) the 4 K working space in the cryostat, c) a disassembled series array module.




Figure 8 - A photograph of the Cu-Ni flex harness at the 4 K end (cryostat thermal shields removed).




Figure 9 - Closed loop harmonic distortion versus ADC lockpoint. A comparison to theory for a SQUID in a flux-locked loop (solid line) shows good agreement.






Figure 10 - Noise performance of the digital multiplexing system. Displayed in units of magnetic flux/Hz^{1/2} as referred to the 1st stage SQUID. Measurements made with an open (first stage) input: a) noise versus sampling frequency, b) versus number of samples averaged during a single sampling event.

a)



b)




Figure 11 – Data from the multiplexing system under operation with variable frequency sine waves applied to first stage SQUID inputs: a) real-time stream, b) data after demultiplexing in software.