

ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer

B. A. MAZIN, S. R. MEEKER, M. J. STRADER, P. SZYPRYT, D. MARSDEN, J. C. VAN EYKEN,
G. E. DUGGAN, A. B. WALTER, G. ULBRICHT, AND M. JOHNSON

Department of Physics, University of California, Santa Barbara, CA 93106; bmazin@physics.ucsb.edu

B. BUMBLE

NASA Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91125

K. O'BRIEN

Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH, UK

AND

C. STOUGHTON

Fermilab Center for Particle Astrophysics, Batavia, IL 60510

Received 2013 June 19; accepted 2013 September 12; published 2013 October 16

ABSTRACT. We present the design, construction, and commissioning results of ARCONS, the Array Camera for Optical to Near-IR Spectrophotometry. ARCONS is the first ground-based instrument in the optical through near-IR wavelength range based on microwave kinetic inductance detectors (MKIDs). MKIDs are revolutionary cryogenic detectors, capable of detecting single photons and measuring their energy without filters or gratings, similar to an X-ray microcalorimeter. MKIDs are nearly ideal, noiseless photon detectors, as they do not suffer from read noise or dark current and have nearly perfect cosmic ray rejection. ARCONS is an integral field spectrograph (IFS) containing a lens-coupled 2024 pixel MKID array yielding a $20'' \times 20''$ field of view and has been deployed on the Palomar 200 inch and Lick 120 inch telescopes for 24 nights of observing. We present initial results showing that ARCONS and its MKID arrays are now a fully operational and powerful tool for astronomical observations.

Online material: color figures

1. INTRODUCTION

Detectors and instrumentation have always been the biggest drivers of sensitivity improvements in astrophysics. While the collecting area of ground-based telescopes has only improved by a factor of 4 in the last 50 years, the per-pixel sensitivity of astronomical detectors has grown by at least a factor of 20, and the improvements in quality and size have been even more impressive. These improvements have led to a golden age in astrophysics, but we are rapidly reaching a plateau in the per-pixel performance of traditional semiconductor-based (CCD [Smith 2011] and HgCdTe [Hall et al. 2012]) detectors, with improvements only coming in the total size and pixel count of the final mosaic (Abell et al. 2009; Flaugher et al. 2012).

Maintaining the rapid pace of advance in per-pixel performance of detectors for astrophysics requires breaking away from semiconductors-based detectors and developing technologies with potentially greater performance, for instance, by reducing read noise and dark current, having a wider bandwidth, or having inherent spectral resolution. The most promising

avenue is to use superconducting detectors (Moseley et al. 1988; Irwin et al. 1996; Li et al. 2001; Day et al. 2003) These detectors significantly increase performance by operating below 4 K, thereby reducing the contribution from thermal noise, allowing them to measure the energy and arrival time of a single photon with no false counts.

The Array Camera for Optical to Near-IR Spectrophotometry (ARCONS) was built around a highly multiplexible type of low temperature detector (LTD) known as microwave kinetic inductance detectors (MKIDs) (Day et al. 2003; Mazin et al. 2012), optimized for optical and near-infrared astronomy. The goal of ARCONS is to demonstrate the viability of both the MKID technology as well as do science that would be difficult or impossible with conventional instruments.

For certain types of observations, the MKID arrays in ARCONS gives significant advantages over a conventional lenslet, fiber fed, or image slicer integral field spectrograph (IFS) such as:

1. Simple optical design that enables very high throughput.

2. Time resolution up to 6 orders of magnitude better than a CCD.

3. Extremely broad intrinsic bandwidth (0.1–5 μm) with good quantum efficiency.

4. No read noise or dark current and nearly perfect cosmic ray rejection.

5. No observing time is lost to readout of the array.

6. Simpler scaling to much larger arrays than conventional IFSs. For example, HARMONI (Tecza et al. 2012) for the E-ELT will use eight $4\text{ k} \times 4\text{ k}$ detectors but only has 32 768 spaxels. In an MKID, each pixel is a spaxel, albeit at significantly lower spectral resolution than an IFS like HARMONI.

7. Time domain information allows *ex post facto* use of calibration stars for monitoring atmospheric transparency, setting dynamic apertures, and applying tip/tilt corrections. This improves signal-to-noise and calibration accuracy and allows the instrument to make optimal use of the atmospheric conditions.

8. Photon arrival time, spectral resolution, and the large number of pixels allow for monitoring and the removal of sky emission, potentially down to the Poisson limit, even in spectral regimes dominated by time variable OH emission.

The MKIDs used in ARCONS are relatively immature, leaving significant room to improve future instruments. Current optical MKID technology has a spectral resolution $R = \lambda / \Delta\lambda \sim 10$ at 4000 \AA . Improvements in the MKID and readout electronics are on track to increase the spectral resolution towards the theoretical limit of around 100 for a 100 mK operating temperature. The array size, though small compared to an imaging CCD (although not small compared to IFSs), is currently only limited by what we can afford to read out with our room temperature electronics (§ 5). The components required for increasing readout pixel count are faster Analog to Digital Converter (ADCs) and more powerful Field Programmable Gate Array (FPGAs), and both of these components are driven by large commercial markets. Assuming current trends hold, it should be possible to read out a megapixel scale MKID array at a reasonable cost within the next decade.

We describe the MKID arrays in detail in § 2. The optical design is presented in § 3, and the instrument design, including cryogenics, mounting hardware, guide camera, and wavelength calibration system, is presented in § 4. Section 5 describes the cryogenic and room temperature readout electronics, and § 6 contains a brief outline of the data storage and processing. Section 7 contains on-sky performance results.

2. MICROWAVE KINETIC INDUCTANCE DETECTORS

Low temperature detectors (LTDs) with operating temperatures on the order of 100 mK, are currently the preferred technology for astronomical observations over most of the electromagnetic spectrum, notably in the far-infrared through millimeter (0.1–3 mm) (Bintley et al. 2010; Niemack et al. 2008;

Carlstrom et al. 2011), X-ray (Kelley et al. 2009), and gamma-ray (Dorise et al. 2007) wavelength ranges. In the important ultraviolet, optical, and near-infrared (UVOIR) (0.1–5 μm) wavelength range, a variety of detector technologies based on semiconductors, backed by large investments from both consumer and military customers, has resulted in detectors for astronomy with large formats, high quantum efficiency, and low readout noise. These detectors, however, are fundamentally limited by the large band gap of the semiconductor, which restricts the maximum detectable wavelength (1.1 eV for silicon), and thermal noise sources from their relatively high ($\sim 100\text{ K}$) operating temperatures. LTDs allow the use of superconductors with gap parameters over 1000 times lower than semiconductors. This difference allows a leap in capabilities. A superconducting detector can count single photons with no false counts while determining the energy (to several percent or better) and arrival time (to a microsecond) of the photon (the optical analog of an X-ray calorimeter). It can also have much broader wavelength coverage since the photon energy is always much greater than the gap energy. While a CCD is limited to about 0.3–1 μm , the MKIDs described here are in principle sensitive from 0.1 μm in the UV to greater than 5 μm , enabling observations at important infrared wavelengths.

Superconducting UVOIR detectors have been pursued in the past with two technologies, superconducting tunnel junctions (STJs) (Martin et al. 2006; Hijmering et al. 2008) and transition edge sensors (TESs) (Romani et al. 2001; Burney et al. 2006). While both of these technologies result in functional detectors, they are limited to single pixels or small arrays due to the lack of a credible strategy for wiring and multiplexing large numbers of detectors, although recently there have been proposals for larger TES multiplexers (Niemack et al. 2010).

Microwave kinetic inductance detectors, or MKIDs (Day et al. 2003; Mazin et al. 2012), are a newer cryogenic detector technology that has proven important for astrophysics (Schlaerth et al. 2010; Monfardini et al. 2010) due to their sensitivity and the ease with which they can be multiplexed into large arrays. The “microwave” in MKIDs comes from their use of frequency domain multiplexing (Mazin et al. 2006) at microwave frequencies (0.1–20 GHz) that allows thousands of pixels to be read out over a single microwave cable. The optical lumped element (OLE) (Doyle et al. 2008) MKID arrays we have developed have significant advantages over semiconductor detectors. They can count individual photons with no false counts and determine the energy and arrival time of every photon with good quantum efficiency. Their physical pixel size and maximum count rate is well matched with large telescopes. These capabilities enable powerful new astrophysical instruments usable from the ground and space.

MKIDs work on the principle that incident photons change the surface impedance of a superconductor through the kinetic inductance effect (Mattis & Bardeen 1958). The kinetic inductance effect occurs because energy can be stored in the

supercurrent (the flow of Cooper pairs) of a superconductor. Reversing the direction of the supercurrent requires extracting the kinetic energy stored in it, which yields an extra inductance term in addition to the familiar geometric inductance. The magnitude of the change in surface impedance depends on the number of Cooper pairs broken by incident photons and is hence proportional to the amount of energy deposited in the superconductor. This change can be accurately measured by placing a superconducting inductor in a lithographed resonator. A microwave probe signal is tuned to the resonant frequency of the resonator, and any photons that are absorbed in the inductor will imprint their signature as changes in phase and amplitude of this probe signal. Since the quality factor Q of the resonators is high and their microwave transmission off resonance is nearly perfect, multiplexing can be accomplished by tuning each pixel to a different resonant frequency with lithography during device fabrication. A comb of probe signals can be sent into the device, and room temperature electronics can recover the changes in amplitude and phase without significant cross talk (McHugh et al. 2012), as shown in the right panel of Figure 1.

The primary theoretical limitation on the spectral resolution is from the intrinsic quasi-particle creation statistics during the photon absorption event. The energy from the photon can end up in two places, the quasi-particle system and the phonon

system. These systems interact, allowing energy exchange between the two, which reduces the statistical fluctuation from Poisson by the Fano factor F , typically assumed to be 0.2 (Fano 1947). The spectral resolution of the detector, $R = \lambda/\Delta\lambda = E/\Delta E$, can be written as $R = (1/2.355)(\eta h\nu/F\Delta)^{1/2}$, where η is the efficiency of conversion of energy into quasi particles, typically 0.57 (Kozorezov et al. 2007), $h\nu$ is the photon energy, F is the Fano factor, and Δ is the energy gap. The energy gap depends on the superconducting transition temperature (T_c) of the inductor, $\Delta \approx 1.72k_B T_c$, and we typically operate at a base temperature of $T_c/8$. Going to lower T_c , and hence lower operating temperature, improves the theoretical R . Operating at 100 mK yields a theoretical spectral resolution of $R \approx 100$ at 400 nm. Previous research with superconducting tunnel junctions (STJs) with superconducting absorbers has shown that superconducting absorbers can approach the Fano limit (Li et al. 2001; Huber et al. 2004; le Grand et al. 1998).

MKIDs are very versatile, as essentially any resonator with a superconductor as the inductor will function as an MKID. In 2008, we decided to pursue a lumped element resonator design (Doyle et al. 2008). The resonator itself consists of a 60 nm thick (Marsden et al. 2012) substoichiometric titanium nitride (TiN_x) film (Leduc et al. 2010) with the nitrogen content tuned with $x < 1$ such that the superconducting transition temperature

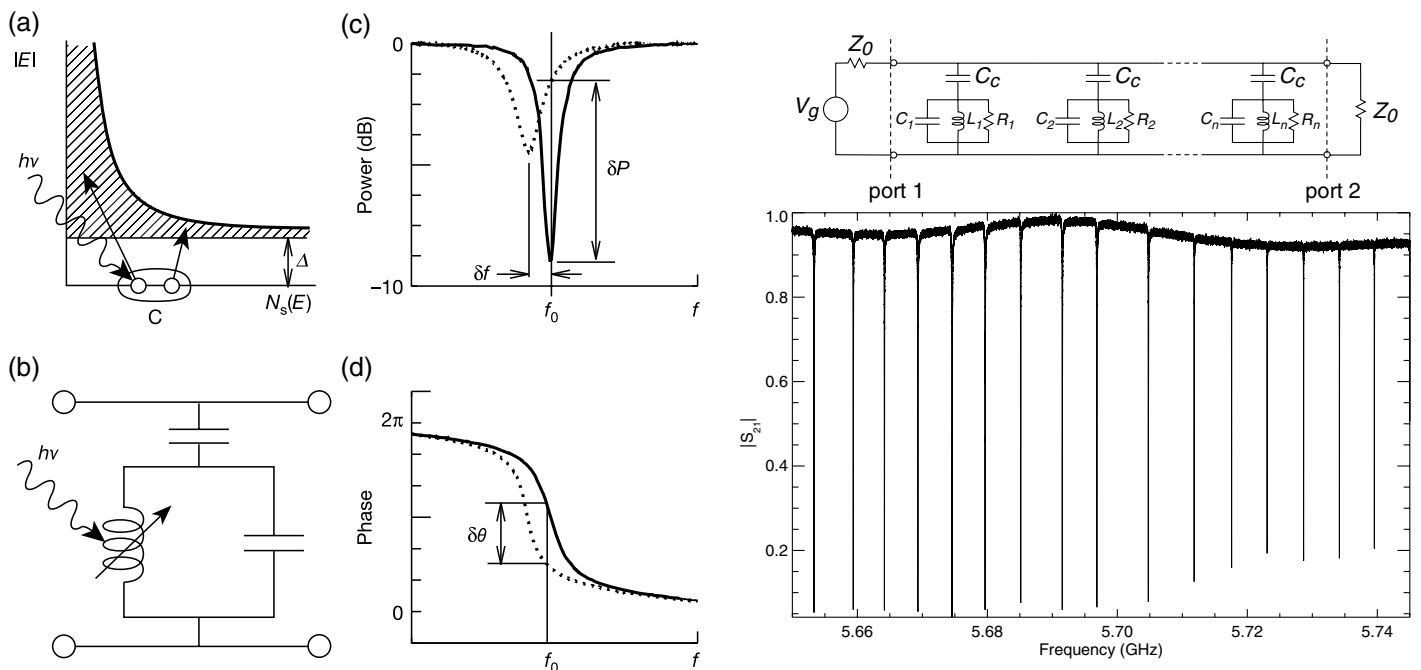


FIG. 1.—*Left*: Basic operation of an MKID, reprinted from Day et al. (2003). (a) An energy level diagram showing the number of states $N_s(E)$ of a superconductor with gap Δ . Photons with energy $h\nu$ are absorbed in a superconducting film, breaking Cooper pairs C , and producing a number of excitations called quasi-particles. (b) To sensitively measure these quasi-particles, the film is placed in a high frequency planar resonant circuit. In the *right panels*, the amplitude (c) and phase (d) response of a microwave excitation signal sent through the resonator as a function of frequency are shown. The change in the surface impedance of the film following a photon absorption event pushes the resonance to lower frequency and changes its amplitude. If the detector (resonator) is excited with a constant on-resonance microwave signal, the energy of the absorbed photon can be determined by measuring the degree of phase and amplitude shift. *Right*: Example of frequency domain multiplexing (FDM) of MKIDs showing many resonators being read out through a single transmission line.

T_c is about 800 mK. Due to the long penetration depth of these films (~ 1000 nm), the surface inductance is a high 25 pH/square, allowing a very compact resonator fitting in a $222 \times 222 \mu\text{m}$ square. A square microlens array with a 92% optical fill factor is used to focus light onto the photosensitive inductor. Microlens arrays with overall sizes over 100 mm are available from commercial vendors. High fill factor multilayer MKIDS, with the inductor made on top of a parallel plate capacitor, are possible for applications where the microlens array is problematic, such as the deep UV. The pixel pitch is easily controlled, with pitches between $75\text{--}500 \mu\text{m}$ relatively easy to achieve. The quasi-particle lifetime in our TiN films is $50\text{--}100 \mu\text{s}$. This sets the pulse decay time, allowing a maximum count rate of approximately $2500 \text{ counts pixel}^{-1} \text{ s}^{-1}$ before problems arise in separating pulses.

After significant development (Mazin et al. 2012), we have arrived at the 2024 (44×46) pixel array shown in Figure 2 and 3. Typical performance for single pixel devices is $R \sim 10$ at 4000 \AA . Over an entire array, we see a median spectral resolution $R \sim 8$ at 4000 \AA (§ 7.2). Typically $>90\%$ of the resonators show up in frequency sweeps, but due to the variations in the TiN gap there are a significant number of collisions (two or more pixels with resonant frequencies closer together than 500 kHz), reducing usable pixels, as discussed in § 7.1. After cutting out collisions and pixels with especially high or low quality factor, we usually have $\sim 70\%$ of our pixels usable. A more uniform film, such as the ones made at the National Institute of Standards and Technology (NIST) with Ti/TiN multilayers (Vissers et al. 2012) or with atomic layer deposition (ALD), could significantly reduce the number of collisions, dramatically improving yield.




FIG. 2.—Photograph of a 2024 pixel UVOIR MKID array (with microlens removed for clarity) in a microwave package. See the electronic edition of the *PASP* for a color version of this figure.

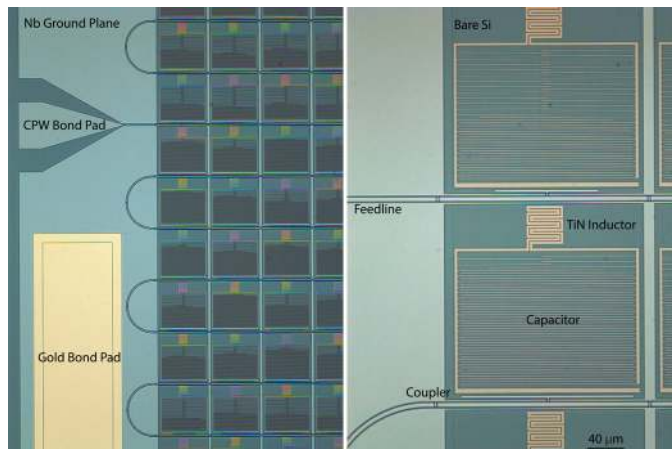


FIG. 3.—*Left*: Microscope image of a portion of the 2024 pixel MKID array used at the Palomar 200 inch telescope. A microlens focuses the light on to the $40 \times 40 \mu\text{m}$ inductor. Various features of interest are labeled. *Right*: Zoomed in view of the array in the *left panel*. See the electronic edition of the *PASP* for a color version of this figure.

Even with the current moderate spectral resolution and yield, these arrays are very useful tools for narrow field-of-view astrophysics.

3. OPTICAL DESIGN

An optical system, shown in Figure 4, has been designed with Zemax to couple ARCONS with the $f/30$ beam at the Palomar 200 inch Coudé focus. A very similar system (using just a different dewar entrance lens) is used to couple to the Lick 3 m Coudé focus. This focus is stationary, which makes interfacing with the cryostat (and its compressor) significantly easier, and the time resolution of the MKIDs means that rotation of the image plane over time can be easily removed.

The optical path is as follows: light bounces off the primary and secondary, then off the Coudé M3 that directs light down the polar axis of the equatorial mount. Once in the Coudé antechamber, a 100 mm diameter pickoff mirror (M4) sends the light towards the dewar entrance window. Before the dewar entrance, two pickoff mirrors form a $\sim 25''$ wide slit and direct the rest of the field of view towards the guide camera's reimaging optics. Immediately in front of the entrance window is a commercial Thorlabs 25 mm filter wheel, which can hold up to six filters




FIG. 4.—Optical layout of ARCONS. Light enters ARCONS at the *left side* of the figure and is absorbed in the MKID array on the *right side* of the figure. See the electronic edition of the *PASP* for a color version of this figure.

TABLE 1
OPTICAL PRESCRIPTION FOR ARCONS

Surface	Type	Radius (mm)	Thickness (mm)	Glass	Diameter (mm)	Comment
1	STANDARD	-46.44035	-7	N-BK7	25	
2	STANDARD	33.77009	-2.5	N-SF5	25	
3	STANDARD	95.94	-70.7		25	
4	STANDARD	Infinity	2.538639		2.1	Lyot stop (not used)
5	STANDARD	Infinity	-10		6	SuperCold Filter
6	STANDARD	Infinity	-37.3		6	
7	STANDARD	-35.09	-6.6	N-BK7	25	
8	STANDARD	35.09	-2.5		25	
9	STANDARD	41.18	-4	N-SF8	17	
10	STANDARD	28.35	-1.5		18	
11	STANDARD	-82	-4	N-BK7	20	
12	STANDARD	82	-2.2		20	
13	STANDARD	23.54	-3.5	N-SF11	20	
14	STANDARD	Infinity	-12		20	
15	STANDARD	Infinity	-1	BK7	13	100 mK IR Blocker
16	STANDARD	Infinity	-4.5		13	
17	STANDARD	Infinity	-1	LITHOSIL-Q	10	
18	USERSURF	0.43	-0.78	LENS ARRAY	10	APO-Q-P222-F0.93mm

for calibration. Observing with ARCONS is typically done at a filter wheel position that does not contain a filter.

The dewar entrance window is an achromatic doublet which collimates the light. All the powered optics have an Edmund Optics VIS-NIR antireflection coating (reflection $<1\%$ per surface between 4000–10000 Å) except the N-SF8 meniscus lens (surfaces 9 and 10 in Table 1), which is uncoated. The light then passes through the dewar's 77 K shield through a baffle and encounters the dewar's 4 K radiation shield. At the 4 K radiation shield an Asahi SuperCold filter¹ is placed to set the red edge of ARCONS's wavelength sensitivity. This was set conservatively at 1.1 μm in order to reduce sky count rates, but in subsequent runs we will attempt to extend the wavelength coverage out through *J* band to 1.35 μm .

There is a pupil image at 4 K suitable for placing a cold 2.1 mm diameter Lyot stop, but this turned out not to be necessary because the glass in the optics at 4 K blocks out enough of the 300 K thermal blackbody radiation to keep the 100 mK stage from warming up. The SuperCold filter does effectively function as a cold Lyot stop with a diameter of 12 mm. There are two achromatic doublets at 4 K to decollimate the beam and set the plate scale, but these achromats are physically separated to eliminate any potential issues with differential thermal contraction between the different glasses. After the final 4 K optic, the light reaches the 100 mK stage. Here the 4 K blackbody radiation is blocked out with a thin piece of borosilicate glass coated in indium tin oxide (ITO), a transparent conductor. The glass absorbs

infrared up to roughly 400 μm , and the ITO reflects a substantial portion of the longer wavelength radiation. Finally, a square microlens with a 0.92 mm focal length and a 92% optical fill factor focuses the light on the sensitive MKID inductor. The optical prescription is shown in tabular form in Table 1.

The resulting spot diagram can be seen in Figure 5. The simulation includes the microlens array in front of the detector. The MKIDs' pixel pitch of 222 μm means that the optical system does not need high magnification, simplifying the system. The system is designed to have 0.45" pixels on the sky. The clear aperture of the optical system was designed to allow an unvignetted field of view of $20'' \times 20''$.




FIG. 5.—Spot diagram of the ARCONS optics at 0.4 (blue crosses), 0.7 (green x's), and 1.0 μm (red squares) plotted on a grid with overall dimensions of $50 \times 50 \mu\text{m}$. The left panel shows the spot on axis, and the right panel shows the spot at the extreme right side of the array. Even at the edges of the array the vast majority of the spot is contained within the $40 \times 40 \mu\text{m}$ inductor. See the electronic edition of the *PASP* for a color version of this figure.

¹ <http://www.asahi-spectra.com/opticalfilters/detail.asp?key=YSC1100>.

This optical system was overdesigned with every precaution possible to block out thermal blackbody radiation, which did not turn out to be required. Future MKID cameras will likely completely eliminate the powered cold optics in favor of much simpler optical systems, including just warm optics to adjust the plate scale, a flat dewar window, a 4 K BK7 flat window, and a 100 mK ITO on BK7 IR blocker.

4. INSTRUMENT ASSEMBLY AND MOUNTING

4.1. Cryogenics

ARCONS is built around an adiabatic demagnetization refrigerator (ADR) made by Janis Research. The system contains a Cryomech PT-405 two stage pulse tube cooler. The first stage of the pulse tube cools the outer radiation shield to 50–70 K, while the second stage cools the ADR to 3.3 K. The pulse tube is powered by an air-cooled compressor, which can be located up to 100 feet from the cryostat. At Palomar, the compressor is placed in a small room on the telescope floor that has a hatch that can be opened to vent the heat outside of the dome. At Lick, the compressor is placed in the battery room outside the dome.

The ADR is a single shot magnetic cooler. During the afternoon, the superconducting ADR magnet is ramped to 40 kG, the heat switch is opened, and then the magnet is slowly ramped down during the course of the night. A PID loop on a Lakeshore 370 controller provides input to a Lakeshore 625 superconducting magnet power supply, allowing us to stabilize the temperature of the ADR to better than 50 μ K. Stabilizing the temperature at 110 mK provides over 12 hr of hold time. There are no obvious drifts in the MKIDs' resonant frequencies as the magnetic field remaining in the superconducting magnet ramps down from 0.75 to 0 kG over the 12 hr hold time.

The limited cooling power of the ADR may be a limitation for some instrument designs. In this case, a dilution refrigerator with approximately 500 times the cooling capacity of an ADR will allow continuous cooling of extremely large arrays at the cost of a more complex cryostat.

4.2. Mounting to the Telescope

At the Palomar 200 inch, an aluminum frame is attached to a steel plate in the Coudé antechamber that is used to accommodate the Coudé spectrograph's slit jaws. ARCONS hangs down from this aluminum frame, as shown in Figure 6.

At the Lick 3 m, ARCONS is also placed in the Coudé antechamber; but since the Hamilton spectrograph is still in use, the mounting plate is not used. Instead, ARCONS is hung from an I beam in the ceiling and then stabilized to the mounting plate with a steel strut.

4.3. Guide Camera and Software

The guide camera for ARCONS is a SBIG STF-8300M CCD camera. Light is reflected off mirrors on either side of the



FIG. 6.—ARCONS installed at the Palomar 200 inch. See the electronic edition of the *PASP* for a color version of this figure.

ARCONS entrance aperture, bounces to an off-axis paraboloid, and then goes through a Nikon 50 mm $f/1.4$ lens mounted to the SBIG camera. The unvignetted field of view of the guide camera is about 1.5' in diameter. The SBIG camera, using 3×3 binning, is read out over USB by a small Microsoft Windows computer located in the electronics rack. In 10 s exposures we are able to guide off of stars in the 18–20th magnitude range. For brighter objects, exposure times down to 0.1 s are used. This computer is accessed by remote desktop in the control room. The guiding software is a modified version of the Palomar Observatory's standard guiding package.

4.4. Wavelength Calibration System

As part of the data taking procedure, each pixel's raw phase response must be mapped to photon energy. We developed a system to uniformly illuminate the ARCONS focal plane with three lasers at the same time before and after science observations, thereby providing a reference to calibrate each detector's response. Roughly 1 minute is sufficient to collect the required wavelength calibration information.

The three lasers operate at 4066, 6710, and 9821 Å. They are integrated into a package that allows for adjustment of the intensity of each laser's light. The light from the three lasers is combined in an integrating sphere with a fiber optic output. The fiber optic carries the combined light to optics and diffusers that simulate the $f/30$ telescope beam. A mirror on a mechanical arm can be rotated into position to block out the sky light and direct the calibration beam into the dewar. The amount of exposure time and frequency of calibrations were controlled through software. During observations, this process occurred approximately once every hour.

The response for a single detector is shown in Figure 7, where the number of counts for a 1 minute integration is plotted




FIG. 7.—An example of wavelength calibrating a pixel. Typical histogram of pulse heights returned by the wavelength calibration system is shown in blue, and a model fit to three Gaussian functions is shown in red. Lower values of phase response indicate higher energy photons. See the electronic edition of the *PASP* for a color version of this figure.

against detector phase in blue, using bins of width one digital number as returned by the firmware. Three Gaussian functions were simultaneously fit to the three peaks produced in response to the laser light, with the width of each Gaussian function being a measure of the noise and, therefore, spectral resolution of the detector. While the widths of the Gaussian functions change with count rate (§ 7.2), the locations of the peaks appear to be stable. The fit is shown in red. The tail out to low energy (less negative phase response) is due to the low energy background

(thermal blackbody radiation, noise triggers, etc.) and is not included in the fit. The location of the three peaks is then fit with a second-order polynomial in phase and used to convert signals on that pixel to photon energy. A given calibration solution is applied to observations near in time. An analysis of the calibration solutions over time showed the polynomial coefficients to be stable over a given night, although more work needs to be done to understand the wavelength calibration over multiple nights. The wavelength calibration is stable enough that it is not noticeably degrading our spectral resolution.

5. READOUT

5.1. Room Temperature Readout

A block diagram for the digital readout for ARCONS is shown in Figure 8 and is described in detail in McHugh et al. (2012). The digital readout for ARCONS utilizes eight reconfigurable open architecture computing hardware (ROACH) boards produced by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) (Parsons et al. 2009). Each ROACH board houses a Virtex5 field programmable gate array (FPGA) and a PowerPC processor, and each board is connected to two digital-to-analog converters (DACs), two analog-to-digital converters (ADCs), and a circuit board used for mixing signals, which is called the intermediate frequency (IF) board. The firmware on the FPGA is used to control the output of the DAC and to process the digitized signal from the ADC.

The two outputs of the DAC make up the in-phase and quadrature (IQ) components of a complex signal used in a quadrature amplitude modulation strategy. The DAC is used to output a frequency comb that is mixed with a local oscillator (LO) frequency to move the comb into the frequency range of the MKID

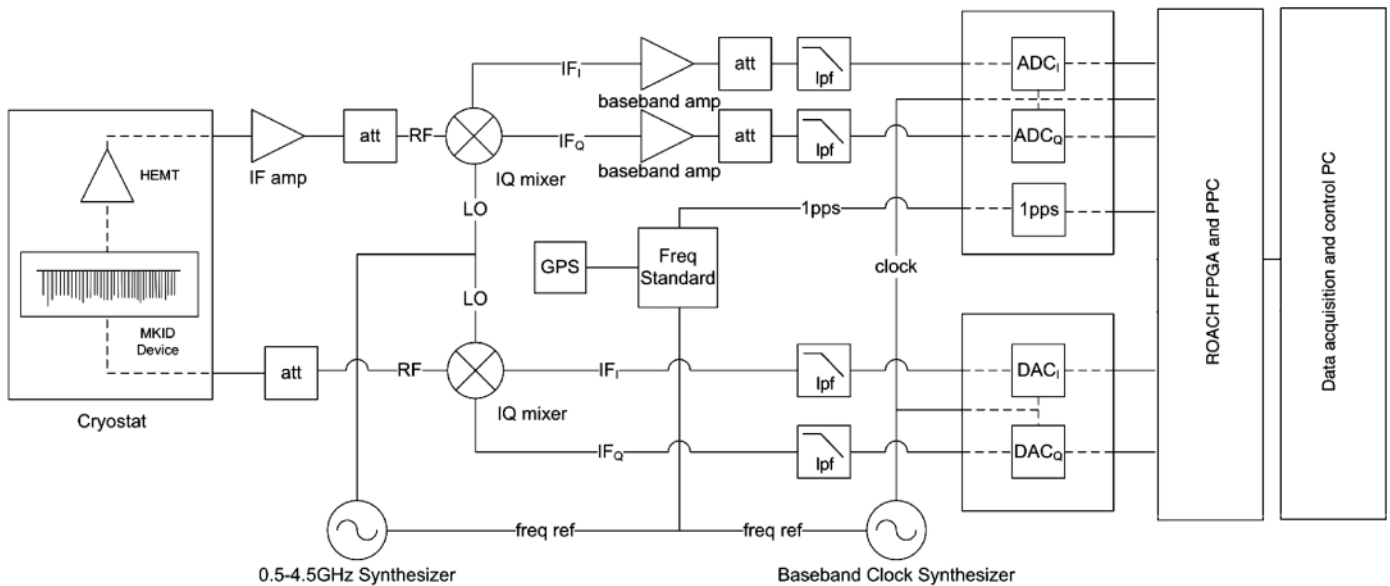


FIG. 8.—Block diagram of the digital readout for ARCONS, reprinted from McHugh et al. (2012).

resonator frequencies. This is then sent through the MKID array. Each frequency tone in the comb corresponds to one MKID resonator frequency. The phase of the tone is set so the baseline phase of the same frequency read in the ADC is approximately zero. The amplitude of the tone is set according to the optimal readout power of the resonator addressed relative to the other resonators. The complete frequency comb sent to the DAC is retrieved by the FPGA firmware from a look-up table (LUT) stored in memory on the ROACH board.

The DACs and ADCs are clocked at 512 MHz and the FPGA is clocked at 256 MHz. The frequency comb output from each DAC can read out up to 256 resonators in 512 MHz of bandwidth, with 2 MHz spacing between tone frequencies. The 2024 pixel array used in ARCONS is divided into two feedlines with 1012 pixels each. Four readout boards are connected to each feedline through power combiners and power splitters.

Once the readout signal has passed through the MKIDs and has been mixed back down to baseband frequencies by the IF board and digitized by the ADC, the FPGA firmware divides the signal into 256 frequency bins (channels) through a “channelization” process so that the contribution from each probe tone in the comb can be treated separately.

The firmware then searches for indications of photon absorption in the phase of each channel’s signal. These are seen as a sudden decrease in the phase followed by an exponential decay back to the baseline, with the characteristic time determined by the quasi-particle recombination time of $\sim 50 \mu\text{s}$. A template of this phase pulse shape was made for each pixel and used to create a customized optimal filter for that pixel. The firmware applies the optimal filter to a channel’s signal, then searches for negative peaks in the filtered phase that pass below a set threshold value. This threshold value is determined by taking a time series of the optimally filtered phase data while the filter wheel is set to block out any incoming light. These data are made into a histogram, and a threshold value is chosen around the 4σ level so false triggers are rare but present. These false events will not impact the final data as the wavelength calibration module in the software pipeline (§ 6) determines the maximum useable wavelength for a pixel at a shorter wavelength than these false triggers. This can be seen on the right side of Figure 7, where the trigger threshold would be set to around -80 digital units, but the wavelength calibration would flag any photons redder than about -120 digital units as unusable false counts.

Each time the peak detection triggers, the firmware creates a 64 bit data packet containing the channel number (so that the photon can be identified with a particular pixel), a time stamp, and values specifying the height of the phase pulse peak and the baseline phase level before the pulse to help remove any $1/f$ phase fluctuations. A dead time of $100 \mu\text{s}$ is added after a trigger to keep pixels with incorrectly set trigger thresholds from filling up the buffers with false pulses. A more elegant way of dealing with incorrect thresholds will be developed for future versions of the readout firmware.

The ROACH boards are synchronized using the 1 PPS output of a Meinberg GPS170PEX GPS board and a Stanford Research FS725 Rubidium 10 MHz frequency standard, which is needed for the time stamps to be recorded with $2 \mu\text{s}$ resolution.

The firmware outputs the data packets for the detected photons to two shared memory blocks on the ROACH. A C program called PulseServer running on the PowerPC located on each ROACH monitors the shared memory blocks. When they are filled to a certain point, PulseServer sends the data over a TCP connection to the data acquisition and control computer. A C program called PacketMaster, running on the data acquisition and control computer, receives the TCP data bundles from all eight ROACH boards. It sifts through each photon packet and organizes them by pixel and second of arrival to write into an Hierarchical Data Format (HDF)² data file. PacketMaster also creates images for the real-time display on the control computer. The resulting HDF data file can later be processed with our software pipeline, detailed in § 6.

5.2. Cryogenic Wiring

ARCONS contains two independent microwave signal paths, which we refer to as feedline 1 and 2. Each feedline reads out half of the array. The signal path begins with a hermetic SMA bulkhead connector that brings the comb of frequencies into the dewar. A laser welded 0.087 inch stainless steel semirigid coaxial cable brings the signals from room temperature to 4 K, with a heat sink clamp cooling only the outer conductor at the 50–70 K stage. At 4 K there is a 30 dB attenuator, followed by an inner and outer DC block. After the DC block a 0.062 inch NbTi/NbTi (outer conductor/center pin) coax brings the signals to the 1 K stage, where the coax is interrupted with a heat-sunk SMA barrel connector. A second NbTi coax brings the signals to 100 mK. After passing through the array, the signal goes through another NbTi coax to 1 K, then a second to 4 K, where it enters the input port of a High Electron Mobility Transistor (HEMT) amplifier (Weinreb et al. 2009). After the HEMT, another stainless coax brings the signals to a hermetic SMA bulkhead connector and outside the dewar.

6. DATA PROCESSING PIPELINE

The ARCONS data reduction pipeline, though currently under development, is outlined here (see Fig. 9 for a schematic). The algorithms, written in Python, will be made publicly available once the pipeline has become stable and reliable. Raw “observation” files consist of a set of tables, one for each pixel, with each row in the tables containing one second’s worth of data packets, where each data packet represents a single photon event. These files form the raw input to the pipeline. The main steps in data reduction are as follows:

² <http://www.hdfgroup.org/>.

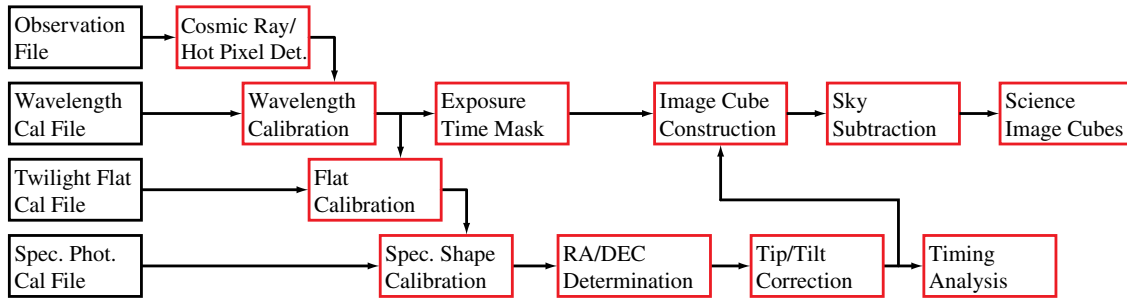


FIG. 9.—General outline of the ARCONS data reduction pipeline. See the electronic edition of the *PASP* for a color version of this figure.

Barycenter Photon Arrival Times: The time of the start of each observation is stored in that observation HDF file’s header. The photon arrival time stamps recorded in the file are relative to this observation start time, so absolute arrival times at the location of the telescope are found by simply adding the recorded time stamps to the header time. These time stamps can then be corrected to solar system barycentric times with the TEMPO2 pulsar timing package (Hobbs et al. 2006) using a custom plugin that allows it to treat photons individually. The barycentered time stamps can then be used to compare the timing between signals recorded from different observatory locations without the effects of the Earth’s motion interfering.

Cosmic Ray Cleaning: Cosmic rays hitting a MKID or the silicon wafer they sit on will heat the wafer, causing phonons to propagate throughout the wafer. These phonons break Cooper pairs in the MKID, causing many pixels to trigger nearly simultaneously. To clean these undesired events from the data, the photon arrival times across all pixels are binned into histograms with overlapping ($\sim 10 \mu\text{s}$) bins, in order to search for multiple photon detection coincidences. Where an interval is found with an unusually high number of coincident photon detections, that interval of $\sim 100 \mu\text{s}$ is flagged as a cosmic-ray contaminated interval for the entire chip. The cosmic-ray event rate is low enough that the integration time lost to this procedure is negligible.

Wavelength Calibration: Wavelength calibration files are created by uniformly exposing the detector array to the combined light of three lasers of known wavelength, as described in § 4.4. For each photon event, the measured phase excursion is assumed to be a monotonically increasing function of photon energy. A histogram of these phase shifts is created for each pixel. In a functional pixel, the histogram shows a superposition of three approximately normal distributions corresponding to the three laser wavelengths. The locations of the maxima of these peaks are fit in order to find the phase shift corresponding to the three wavelengths. A second-order polynomial fit to these three values yields a smooth, continuous, analytic function, mapping phase shift to wavelength for each given pixel.

Flat-field Calibration: Twilight sky exposures are used to normalize the response of each pixel with respect to each other.

The average twilight spectrum is computed by taking the median of the wavelength-calibrated twilight spectrum for each pixel. The ratio of this average spectrum to an individual pixel’s wavelength-calibrated spectrum then gives a weighting as a function of wavelength for that pixel. This weighting is used to normalize the response of different pixels when combining data from multiple pixels, effecting the traditional CCD “flat fielding,” but at every wavelength.

Spectral Shape Calibration: Once the pixel responses are normalized relative to one other, a final absolute calibration of the overall spectral response is performed by comparing the spectrum of an observed reference target against its known spectrum. This provides a single additional whole-array wavelength weighting function. The combination of this absolute weighting and the per-pixel relative flat-fielding weights corresponds to the overall Quantum Efficiency (QE) response of each pixel as a function of wavelength. Applied together, the weights yield a fully calibrated spectrum for each pixel.

Exposure Time Masking: High levels of photon illumination induce a number of the pixels in the array to go “hot” from time to time, registering false counts and reporting artificially high count rates. This is likely due to an interaction of the MKID with free electrons in the silicon wafer, which should be eliminated in new MKID designs. This pixel behavior is detected by comparing the flux distribution in the neighborhood of each pixel to that of the expected point-spread function (PSF) (similar to the algorithm used in the *XMM-Newton* pipeline).³ In each 1 s time step, the ratio of the counts in each pixel to the median counts in a small surrounding box (typically 5×5 pixels) is compared to the same ratio for a two-dimensional Gaussian function with FWHM matching that of the expected PSF (specified manually at the time of the data reduction). If, after accounting for photon shot noise errors, the measured ratio is significantly greater (typically $>3\sigma$), then the source of the counts cannot be astrophysical, and the pixel is flagged as hot

³ *XMM-Newton* Science Analysis System, Users Guide to the *XMM-Newton* Science Analysis System, Issue 9.0, 2012 (ESA: *XMM-Newton* SOC). See also <http://xmm.esac.esa.int/sas/>.

for the given time step. Some pixels are also known to show anomalously *low* count rates (“cold” pixels) or no counts at all (“dead” pixels). In a similar way, therefore, if the count rate for a pixel is significantly lower than the median for the surrounding box, the pixel can be flagged as “cold”. These intervals are then concatenated and tracked to maintain consistency in the total effective exposure time attributable to each pixel. Since the number of bad pixels is relatively small compared to the total number of pixels, and the time intervals over which such unwanted behavior occur are small, the time masks are most efficiently stored as per-pixel lists of bad intervals, also in HDF file format. This process is applicable to any array observation, including the laser wavelength calibrations and the twilight flats.

R.A./Decl. Determination: An astrometric solution is found for the data at regular intervals (typically 30 s). The centroid location of the target source in pixel coordinates is found using the Python language “PyGuide” centroiding algorithm,⁴ which calculates the point of minimum radial asymmetry for a source in a time integrated image. The known sky coordinates of the target can then be assigned to that pixel location. Since the instrument location at the Coudé focus is fixed with respect to the telescope motion, the field of view also rotates at the sidereal rate. The hour angle of the target therefore relates directly to the orientation of the field. In combination with a precalibrated fixed rotational offset for the instrument, a right ascension and declination can be determined for every pixel at each time step. In addition, atmospheric refraction spreads light of different frequencies onto different parts of the array. This phenomenon is well understood (Auer & Standish 2000) and can be easily removed, allowing coordinates to be assigned for every photon.

Tip/Tilt Correction: For sources with sufficient count rates it should be possible to perform the above centroiding procedure on very short timescales, comparable with the timescale of atmospheric seeing variation (~ 10 ms). Alternatively, for non-point-source objects, image cross-correlation can be performed to determine offset corrections to the data on the same timescales. Hence it should be possible in principle to apply *ex post facto* first-order tip-tilt seeing corrections, purely in software, to improve the final recovered PSF.

After these steps have been performed, processed photon lists will form the output data product: calibrated lists of photon events, with each photon having an absolute time stamp, R.A. and decl. location, wavelength, and any associated flags for each photon, along with associated time-mask lists.

Further analysis depends on the science goals and the desired end product. For imaging and spectroscopic analysis, image cubes are constructed by creating images integrated over different wavelength ranges and over the desired time intervals. Background subtraction is performed as in traditional CCD data

reduction (for example, by subtracting a constant or slowly varying two-dimensional polynomial fit to the image after masking bright sources). Spectra are then obtained by integrating over the desired pixels in the images.

Alternatively, science involving timing analysis (e.g., pulsar and eclipsing binary observations) requires a nondiscrete approach in the time axis. We integrate first over selected pixels or a selected area in sky coordinates and maintain the individual photon timing information, avoiding binning into larger time intervals. In the case of pulsar or eclipsing binary observations, for example, this allows for folding of the data on the known period.

Comparing the data from ARCONS with conventional multicolor imaging or spectroscopy is not trivial. For example, dealing with a spectral resolution that varies slightly from pixel to pixel and with count rate can significantly complicate an analysis. More experience with the data will be required to develop best practices for calibration, analysis, and comparison to existing instruments.

7. ON-SKY PERFORMANCE

Even with the current engineering-grade MKID arrays, the extra mirror bounces and path length of the Coudé focus, and nonoptimal optical design and antireflection coatings, ARCONS is a scientifically useful tool for narrow field-of-view astrophysics. It has now been proven in the field over the course of 24 observing nights at 5 m class telescopes. In this section, we explore in detail the performance of the instrument.

7.1. Array Yield

Typically over 90% of the resonators show up in frequency sweeps, but due to the variations in the TiN gap, there are a significant number of collisions, reducing usable pixels, as shown in Figure 10. After cutting out collisions and pixels with especially high or low quality factor, usually about 70% of our pixels are usable. A more uniform film, such as the ones made at NIST with Ti/TiN multilayers (Vissers et al. 2012) or with atomic layer deposition (ALD), could significantly reduce the number of collisions, dramatically improving yield.

7.2. Spectral Resolution

The MKIDs used in ARCONS have an average spectral resolution of $R \sim 10$ at 4000 Å, declining linearly with increasing wavelength, as expected. The median spectral resolution is degraded to $R \sim 8$ with our two channel analog electronics (see Fig. 11) with the full ARCONS digital readout due to limitations in the number of coefficients (taps) in the programmable optical filter in our digital electronics. This issue is being addressed by upgrading our electronics from ROACH to ROACH2, which will allow for significantly more taps.

The achieved spectral resolution is drastically below the theoretical limit for a 100 mK operating temperature of $R = 100$ at

⁴ See <http://www.astro.washington.edu/users/rowen/PyGuide/Manual.html>.



FIG. 10.—Results of beam mapping the ARCONS array. Pixels with good locations are shown in white. Vignetting is apparent at the bottom right side of the array. The overall pixel yield in this engineering-grade array is $\sim 70\%$.




FIG. 11.—Histogram of spectral resolution of one feedline from the array. The black line shows the energy resolution from calibration files taken at Palomar with a high count rate of ~ 1000 counts s^{-1} . The green curve shows the projected energy resolution we expect to recover at the nominal sky count rate of 100 counts s^{-1} by scaling the black data by the expected degradation in spectral resolution with count rate, the red curve in Fig. 13. The other feedline is very similar. See the electronic edition of the *PASP* for a color version of this figure.

4000 Å (Mazin et al. 2012). We believe the current MKIDs are limited by positional effects within the MKID inductor due to local variations in the TiN superconducting gap. Fixing this should roughly double the spectral resolution. After this issue is addressed, lower noise microwave amplifiers (Ho Eom et al. 2012) and new MKID designs should continue to boost the resolution towards the theoretical limit.

As shown in Figure 12 as a red line (*labeled on the right y-axis*), there is also a pronounced degradation of the spectral resolution with increasing count rate. This is caused by two effects. First, the baseline value that the firmware tracks to reduce $1/f$ noise gets noisy from averaging many counts. Second, pulses begin overlapping, degrading resolution. Improvements in the firmware baseline determination algorithm should improve this degradation significantly. At Palomar, the ARCONS sky background rate is 50–100 photons s^{-1} , so only objects significantly brighter than sky will suffer from severe degradation of the spectral resolution.

7.3. Sensitivity

The MKIDs in ARCONS absorb light directly in the TiN film that comprises the inductor. This TiN has an intrinsic absorption of roughly 70% at 4000 Å and 30% at 1 μm . The total system throughput is reduced by the cumulative effect of all the intervening optical materials and surfaces: the atmosphere, the telescope M1–M4 mirrors (the Coudé M3 at Palomar was degraded due to >10 years since realuminization), the collimating lens, the Asahi SuperCold filter at 4 K, the four optical elements

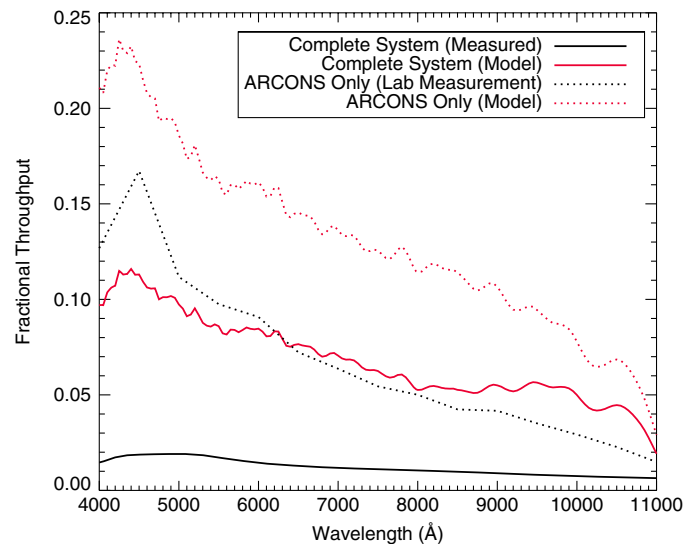


FIG. 12.—Linearity and spectral resolution vs. count rate per pixel of the MKID array in ARCONS measured in the lab. The measured count rate is the count rate actually recorded by the MKID, and the scaled input count rate is the count rate we expect to measure based on the measured flux in a calibrated photodiode. See the electronic edition of the *PASP* for a color version of this figure.

at 4 K, the 100 mK glass with ITO filter, and the microlens fill factor and alignment with the inductors. A model of this entire system, based on a simple atmospheric model, clean Al reflectivity (normalized to 86% at 670 nm), manufacturer's specs for glasses, AR coatings, the SuperCold filter, and measurements of the absorptivity of raw TiN films, is shown as a red solid line in Figure 13. The efficiency of just ARCONS, neglecting the atmosphere and telescope, is shown as the red dashed line.

The theoretical model of ARCONS can be directly compared with laboratory measurements of ARCONS's QE. These tests were done with a QE test bed consisting of a monochromator and light source, integrating sphere, and a rotatable fold mirror to direct the beam to a calibrated photodiode or to the ARCONS optical input. The results of these lab tests are shown as a black dashed line in Figure 13. The roughly factor of 2 discrepancy between the theoretical model and the laboratory measurements is primarily due to a known misalignment of the microlens height above the focal plane, increasing the spot size on the focal plane so it is larger than the inductor, especially at red wavelengths.

The actual system throughput is calculated by doing photometry on a standard star and comparing our results to the object's known spectrum, shown as the solid black line in Figure 13. This curve is about a factor of 2 below the laboratory QE measurements modified by the expected atmospheric transmission and telescope throughput. The cause of this discrepancy is unknown but could include the unknown reflectivity of the Coude M3 flat or angular misalignment of ARCONS with respect to the telescope input.

The MKIDs used in ARCONS have inductors with a preferred orientation of the wires, which could potentially lead to a quantum efficiency that depends on polarization angle. Lab measurements were performed on ARCONS to check whether the MKID arrays exhibited any response to linearly polarized light at different angles. No polarization response was seen to the 5% level.

By going to a Cassegrain focus, improving our antireflection coatings, aligning the microlens well, and simplifying our optics, we should be able to boost our total system throughput in future MKID cameras by a factor of 10. Using MKIDs with black absorbers could eventually yield instruments with total system efficiencies greater than 60%.

7.4. Linearity

The 100 μ s dead time associated with triggering (§ 5) causes some pulses to be missed at high count rates, as shown in Figure 12. This effect can be corrected for analytically by calculating the number of photons that are expected to arrive within a dead-time interval and hence not be counted, $R_c = R_m / (1 - R_m d)$, where R_c is the corrected count rate, R_m is the measured count rate, and d is the dead time, as shown as a green line in Figure 12. The slightly longer effective dead time of 120 μ s used in Figure 12 results from the interaction of

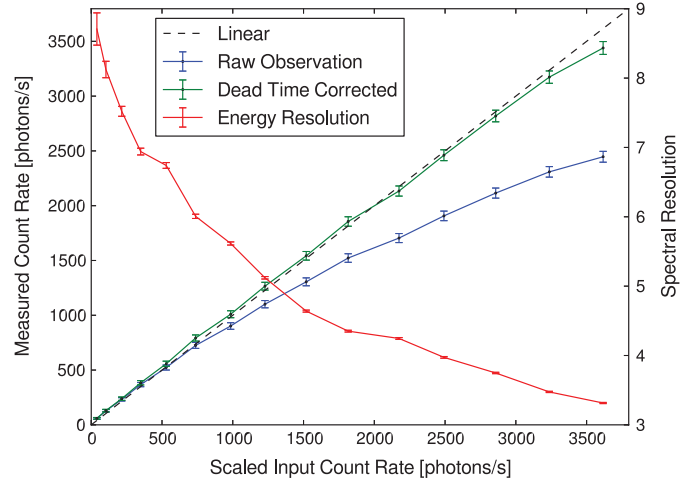


FIG. 13.—System throughput of ARCONS. See the electronic edition of the *PASP* for a color version of this figure.

the firmware dead time and the pulse optimal filters. The firmware of ARCONS is currently being rewritten to remove the dead time and better handle closely spaced photons, which should significantly improve the linearity and spectral resolution at high count rates.




FIG. 14.—Mosaic of the interacting galaxies Arp 147, made with ARCONS on the Palomar 200 inch. The data consist of 36 pointings of 1 minute each. To make this image, the data were cleaned of cosmic rays and hot pixels, wavelength-calibrated, and flat-fielded. These processed data for each of the 36 pointings were offset in R.A. and decl. by the amount requested by our mosaicing control software and combined. No field derotation was performed. The false colors were made by breaking the ARCONS data into three wavelength bands and setting each band to correspond to appropriate red-green-blue (RGB) value. The image was convolved with a Gaussian function to reduce noise below the spatial scale of the seeing, and the image was adjusted in Adobe Photoshop to give an attractive color palette and remove some minor sky noise artifacts. The inset shows a processed Hubble Space Telescope (HST) image of Arp 147. See the electronic edition of the *PASP* for a color version of this figure.




FIG. 15.—Spectra of Landolt 94-42 (Landolt 1992) on top, and the G9V star *COROT*-18 on the bottom, taken with ARCONS on the Palomar 200 inch. These spectra are flux calibrated. Previous *BVRI* photometry is shown as blue points. See the electronic edition of the *PASP* for a color version of this figure.

7.5. Imaging and Spectrophotometry

Imaging and spectroscopic performance of ARCONS is consistent with expectations from the beam map and spectral resolution data. The plate scale was measured to be $0.435''$ per pixel, matching well with our Zemax design. During the run, we experienced one night with seeing measured at $\sim 1.0''$ by the Palomar seeing monitor, and fitting a PSF from this night shows that we recover this seeing value, indicating that our optics are not limiting our imaging performance.

Extremely high quality imaging and spectra will require more of our software pipeline to be in place, but preliminary analyses are presented in Figures 14 and 15. These images and spectra begin to show the capabilities of this unique instrument.

7.6. Timing

ARCONS timing performance was measured by simultaneously observing the Crab pulsar with ARCONS and the Robert C. Byrd Green Bank Telescope (GBT). The data shown in Figure 16 shows that ARCONS sees an optical pulse leading the radio pulse by about $150 \mu\text{s}$, which is consistent with previous observations (Shearer et al. 2003).

Measurement of the arrival time of a single photon should be accurate to $\sim 2 \mu\text{s}$. Currently, there is a delay of $43 \mu\text{s}$ from the time a photon is absorbed by the MKID until the photon packet is ready to be sent over the network, opening up the possibility of extremely fast effective frame rates for certain applications, like wavefront sensing and speckle control in a coronagraph.




FIG. 16.—Alignment of the radio (red line) and optical (blue shading) pulses from the Crab pulsar verify the absolute timing accuracy of ARCONS. See the electronic edition of the *PASP* for a color version of this figure.

8. CONCLUSIONS

ARCONS is a unique, read-noise free photon counting IFS. With a 2024 MKID pixel array covering a $20'' \times 20''$ field of view, ARCONS is the world's largest (and only active) LTD-based instrument in the optical through near-IR. It is uniquely powerful for observations of rapidly variable sources.

The ARCONS instrument has performed astronomical observations, and the first science papers using ARCONS data are currently being written. ARCONS and its successors will continue to improve in pixel count and yield, spectral resolution, system throughput, and detector quantum efficiency. The MKID technology it uses is extremely scalable, allowing arrays approaching megapixels within a decade. MKID-based instruments like ARCONS will start to become a standard part of the UVOIR observer's toolkit in the years to come.

The MKID detectors used in this work were developed under NASA grant NNX11AD55G. S.R. Meeker was supported by a NASA Office of the Chief Technologists Space Technology Research Fellowship, NASA grant NNX11AN29H. This work was partially supported by the Keck Institute for Space Studies. Fermilab is operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy. The authors would also like to thank Shri Kulkarni, Director of the Caltech Optical Observatories, and Jason Prochaska, Associate Director of Lick Observatory, for facilitating this project, as well as the excellent staffs of the Palomar and Lick Observatories for their assistance in getting ARCONS working. Jennifer Milburn's help with the guide camera software was invaluable. This project also greatly benefitted from the support of Mike Werner, Paul Goldsmith, and Jonas Zmuidzinas at JPL. Special thanks to Refree Stephen Shedman for valuable comments

REFERENCES

- Abell, P. A., et al. 2009, preprint (arXiv:0912.0201)
- Auer, L. H., & Standish, E. M. 2000, *AJ*, 119, 2472
- Bintley, D., et al. 2010, *Proc. SPIE*, 7741, 774106
- Burney, J., et al. 2006, *Nucl. Instrum. Methods Phys. Res. A*, 559, 525
- Carlstrom, J. E., et al. 2011, *PASP*, 123, 568
- Day, P. K., Leduc, H. G., Mazin, B. A., Vayonakis, A., & Zmuidzinas, J. 2003, *Nature*, 425, 817
- Doriese, W. B., et al. 2007, *Appl. Phys. Lett.*, 90, 193508
- Doyle, S., Mauskopf, P., Naylor, J., Porch, A., & Duncombe, C. 2008, *J. Low Temp. Phys.*, 151, 530
- Fano, U. 1947, *Phys. Rev.*, 72, 26
- Flaugher, B. L., et al. 2012, *Proc. SPIE*, 8446, 844611
- Hall, D. N. B., et al. 2012, *Proc. SPIE*, 8453, 84530W
- Hijmering, R. A., Verhoeve, P., Martin, D. D. E., Jerjen, I., Kozorezov, A. G., & Venn, R. 2008, *J. Low Temp. Phys.*, 151, 298
- Ho Eom, B., Day, P. K., Leduc, H. G., & Zmuidzinas, J. 2012, *Nat. Phys.*, 8, 623
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, 369, 655
- Huber, M., Angloher, G., Hollerith, C., Ruedig, A., Jochum, J., Potzel, W., & von Feilitzsch, F. 2004, *X-Ray Spectrometry*, 33, 253
- Irwin, K., Hilton, G., Wollman, D., & Martinis, J. 1996, *Appl. Phys. Lett.*, 69, 1945
- Kelley, R. L., et al. 2009, in *AIP Conf. Proc.* 1185, *Thirteenth International Workshop on Low Temperature Detectors*, ed. B. Young, B. Cabrera, & A. Miller (New York: AIP), 757
- Kozorezov, A. G., Wigmore, J. K., Martin, D., Verhoeve, P., & Peacock, A. 2007, *Phys. Rev. B*, 75, 094513
- Landolt, A. U. 1992, *AJ*, 104, 340
- le Grand, J., et al. 1998, *Appl. Phys. Lett.*, 73, 1295
- Leduc, H. G., et al. 2010, *Appl. Phys. Lett.*, 97, 102509
- Li, L., Frunzio, L., Wilson, C., Prober, D. E., Szymkowiak, A. E., & Moseley, S. H. 2001, *J. Appl. Phys.*, 90, 3645
- Marsden, D., Mazin, B. A., Bumble, B., Meeker, S., O'Brien, K., McHugh, S., Strader, M., & Langman, E. 2012, *Proc. SPIE*, 8453, 84530B
- Martin, D., Verhoeve, P., Peacock, A., Kozorezov, A., Wigmore, J., Rogalla, H., & Venn, R. 2006, *Appl. Phys. Lett.*, 88, 123510
- Mattis, D. C., & Bardeen, J. 1958, *Phys. Rev.*, 111, 412
- Mazin, B., Bumble, B., Meeker, S. R., O'Brien, K., McHugh, S., & Langman, E. 2012, *Optics Express*, 20, 1503
- Mazin, B., Day, P., Irwin, K., Reintsema, C., & Zmuidzinas, J. 2006, *Nucl. Instrum. Methods Phys. Res. A*, 559, 799
- McHugh, S., Mazin, B. A., Serfass, B., Meeker, S., O'Brien, K., Duan, R., Raffanti, R., & Werthimer, D. 2012, *Rev. Sci. Instrum.*, 83, 4702
- Monfardini, A., et al. 2010, *A&A*, 521, 29
- Moseley, S. H., Kelley, R. L., Schoelkopf, R. J., Szymkowiak, A. E., McCammon, D., & Zhang, J. 1988, *IEEE Trans. Nucl. Sci.*, NS-35, 59
- Niemack, M. D., et al. 2008, *J. Low Temp. Phys.*, 151, 690
- . 2010, *Appl. Phys. Lett.*, 96, 163509
- Parsons, A., et al. 2009, *Astro2010: The Astronomy and Astrophysics Decadal Survey*, 2010, 21
- Romani, R. W., Miller, A. J., Cabrera, B., Nam, S. W., & Martinis, J. M. 2001, *ApJ*, 563, 221
- Schlaerth, J. A., et al. 2010, *Proc. SPIE*, 7741, 774109
- Shearer, A., Stappers, B., O'Connor, P., Golden, A., Strom, R., Redfern, M., & Ryan, O. 2003, *Science*, 301, 493
- Smith, G. E. 2011, *J. Appl. Phys.*, 109, 102421
- Tecza, M., Thatte, N., Clarke, F., Freeman, D., & Kosmowski, J. 2012, *Proc. SPIE*, 8450, 84501N
- Vissers, M. R., Gao, J., Sandberg, M., Duff, S. M., Wisbey, D. S., Irwin, K. D., & Pappas, D. P. 2012, preprint (arXiv:1209.4853)
- Weinreb, S., Bardin, J., Mani, H., & Jones, G. 2009, *Rev. Sci. Instrum.*, 80, 4702