

Appendix D

QUIET

This appendix contains documents describing some of my contributions to the QUIET project. Most of this work is presented as reproductions of internal specification documents and memoranda. In section D.6, we present a preprint manuscript of the paper describing the first results from the QUIET program.

D.1 Bias Electronics

Here, we present interface specifications for the bias electronics boards to provide bias to the polarimeter modules. Two boards were designed. The *MMIC bias* electronics were used to generate bias signals for the monolithic microwave integrated circuit (MMIC) amplifiers in the modules. The *phase switch bias* electronics were used to provide bias for the phase switches in the modules. The electronics design for the bias boards was developed chiefly by Dr. Michael D. Seiffert and Steve Smith, and the hardware was built and basic functions were tested by Steve Smith. In addition to developing the interface specification, I was responsible for detailed testing, characterization, and debugging of the electronics. The interface specifications are presented here with their original formatting and pagination. In each, a lengthy appendix containing pin assignment tables is not shown here.

QUIET MMIC Bias Interface Specification

Joey Richards*, Mike Seiffert, Steve Smith
Revision 1.5

January 11, 2006

1 Introduction

This document specifies the various interfaces between the MMIC bias cards and the rest of the QUIET system. The MMIC bias card design described in this document is intended to support both 40 GHz (Q-band) and 90 GHz (W-band) receiver modules with at most minor component value changes between versions.

1.1 Schedule

Key dates in the design and production schedule for the MMIC bias cards are listed in Table 1.

Item	Start	Finish
Interface definition	—	6/17/2005
Prototype run (qty 2)	In progress	July
Prototype test/debug	July	7/31/2005
Deliver board to Chicago (qty 1)	—	7/31/2005
Production run (for 19 Q-band elements, 91 W-band elements, plus spares)	TBD	TBD

Table 1: MMIC bias card design/construction schedule.

2 Mode of Operation

In this section, we describe generally how to interact with the MMIC bias card. Details of the operations required are in later sections of this document. Please take note of the warnings regarding circuit operation in this section.

* joey@caltech.edu

2.1 Power-Up Sequence

The MMIC bias card will power-on into a safe “off” state with the MMIC bias is in a low-current mode. The bias can be left in this state indefinitely without risk of module damage.

Incomplete list of Concerns:

- Current balance between MMICs in pre-amp (common drain)
 - Turn both gates on first, then apply drain current

2.1.1 Reset Delay

At power-up or after the DAC /CLR is asserted and de-asserted, an RC delay circuit will hold the /CLR input to the DACs low. A 1 ms delay should be observed after power-up or de-asserting /CLR before attempting to program the DACs.

2.2 Warnings

Please be aware of the following important warnings about MMIC bias card operation.

2.2.1 Limitations of Hardware Failsafes

As much as possible, hardware failsafes (current limits, voltage clamps, etc) have been designed into the MMIC bias circuits. However, it appears that it may not be possible to design the circuits to be completely incapable of damaging the receiver modules. In these cases, software protection will be required to prevent damage to the modules due to user error or software defects. It is strongly advised that this protection be implemented at the lowest level possible in the software drivers to provide maximum protection against software defects.

A detailed list of limitations to hardware failsafes will be included here. These will not be known until the bias cards are designed.

2.2.2 Bias Programming Timing

To prevent contamination of science data, bias values should not be changed while observations are being made. The bias DACs should be updated only during receiver dead time.

2.2.3 Multiplexer Switching Timing

To prevent contamination of science data, the monitor multiplexer address should be changed only during receiver dead time.

2.2.4 Multiplexer Switching Restrictions

In addition to the timing restrictions above, during science data collection, the multiplexer address should only be changed after an even number of phase switch toggles (i.e., synchronously with the 4 kHz phase switch clock). This is to minimize the effect of any bias level offset caused by an interaction with the multiplexer by ensuring such an offset is common mode between the two phase switch states.

Alternatively, sampling may be restricted to the receiver dead time. In this mode, during the dead-time, the multiplexer would switch to the desired signal, the housekeeping board would sample, and the multiplexer would switch back to an isolated state before the end of the dead time.

Interface	Signal Type	Signal Count	Total
Bias Output	Analog	11 per module	154
Bias Control Input	LVDS	5 x 2	10
MAB Ground Sense	Analog	2	2
Monitor/Mux Control Input	LVDS	8 x 2	16
Monitor/Mux Output	Analog	1	1
Power	Power	7	7
Ground	Ground	6	6
Totals:		42 + 11 per mod	196

Table 2: Summary of interface signal counts for the MMIC bias card.

2.2.5 Interactions Between Module Bias Settings

The various bias settings within a MAB will not be entirely independent due to a shared ground return. A ground voltage sense signal is provided to allow software to monitor the voltage of the MAB ground (which will differ from the voltage at the MMIC bias card due to IR drop across the ground wire). It is not expected that this could cause hardware damage, but the bias control software must account for such interactions.

2.2.6 Assumptions About MAB

The gate bias circuit assumes a 1:7.19 ratio voltage divider is installed on the MAB. Use without this divider will damage the module.

3 Electrical Interface

The electrical interface consists of the digital inputs to control the MMIC bias circuit and multiplexers, analog outputs to the MABs, and analog monitor outputs to the housekeeping board. Additionally, the grounding strategy and transient protection requirements are described in this section.

Each MMIC bias card will provide bias for the receiver modules on either two MABs (14 modules). The interface signal requirements are summarized in Table 2.

3.1 Digital Signal Characteristics

All digital inputs to the MMIC bias card are LVDS differential signals. These signals are translated to single-ended digital signals and optoisolated on the MMIC bias card. A power supply and ground must be provided for the LVDS level converter and the input side of the optoisolators (see Section 3.5.2). This power supply and ground must be DC-isolated from the power supplies for the analog sections of the MMIC bias cards (see Section 3.5.3), though it may be common between cards.

3.2 Bias Output

For each 90 GHz receiver module, 10 bias signals must be provided. These are listed in Table 3. The 40 GHz modules will use the same outputs, although they may be connected differently inside the modules.

The descriptions in Table 3 refer to two amps (1 and 2). These correspond to the LNA section (amp 1) and the second gain stage (amp 2) in the module. As seen in the module block diagram in the QUIET

Name	DAC	A3	A2	A1	A0	Description
GX11	0	0	0	0	1	Gate voltage, side X, amp 1, gate 1
GX12	0	0	0	1	1	Gate voltage, side X, amp 1, gate 2
GX2	1	0	0	0	1	Gate voltage, side X, amp 2
DX1	0	0	0	1	0	Drain current, side X, amp 1
DX2	1	0	0	1	0	Drain current, side X, amp 2
GY11	1	0	0	1	1	Gate voltage, side Y, amp 1, gate 1
GY12	0	0	1	0	0	Gate voltage, side Y, amp 1, gate 2
GY2	1	0	1	0	1	Gate voltage, side Y, amp 2
DY1	1	0	1	0	0	Drain current, side Y, amp 1
DY2	0	0	1	0	1	Drain current, side Y, amp 2
RET	—	—	—	—	—	Ground return for module

Table 3: 90 GHz module bias control signals and DAC channel assignments.

Signal	Absolute Maximum			Operating Range			
	V _{MIN}	V _{MAX}	I _{MAX}	V _{MIN}	V _{MAX}	I _{MAX}	ΔV
Gx11	-0.2 V	0.4 V	0.2 mA	-0.1 V	0.3 V	0.1 mA	1 mV
Gx2	-0.2 V	0.4 V	0.2 mA	-0.1 V	0.3 V	0.4 mA	1 mV

Table 4: 40 GHz and 90 GHz module gate bias requirements. ΔV is the gate voltage resolution required to ensure a suitable bias point can be found. These voltages are measured *after* the 1:7.19 voltage divider on the MAB.

90 GHz Module Development memo, the LNA actually consists of two MMIC amplifiers.¹ These each have a separately controlled gate, but their drains are connected to a single pin as indicated.

3.2.1 Module Bias Requirements

The 40 GHz and 90 GHz module bias requirements consist of absolute maximum ratings and operating range ratings. The absolute maximum ratings reflect damage thresholds that must not be exceeded. The operating range ratings reflect the minimum bias range needed to allow proper operation of the module. Ideally, the MMIC bias card hardware will provide bias up to a limit that exceeds the operating range rating but is no greater than the absolute maximum rating.

The gate bias requirements are described in Table 4 and the drain bias requirements are described in Table 5. These tables are based on the QUIET 90 GHz Module Development memo² and personal communications with Todd Gaier. Note that the drain current resolution of 0.1 mA corresponds to a drain voltage resolution of approximately 3 mV or finer.

The MAB circuit design will implement a voltage divider on the gate input with a ratio of 1/7.19. The above specifications refer to the voltages *after* the voltage divider.

¹T. Gaier, “Quiet 90 GHz Module Development,” March 16, 2005, internal memo, 1 (Figure 1).

²*Ibid.*

Signal	Absolute Maximum	Operating Range			
	I_{MAX}	I_{MIN}	I_{MAX}	V_{MAX}	ΔI
Dx1	30 mA	0	25 mA	1.2 V	0.1 mA
Dx2	60 mA	0	50 mA	1.2 V	0.1 mA

Table 5: 40 GHz and 90 GHz module drain bias requirements. ΔI is the drain current resolution required to ensure a suitable bias point can be found.

Signal	Type	V_{MIN}	V_{MAX}	I_{MIN}	I_{MAX}	Output Impedance
Gx11	voltage source	-0.20 V	0.35 V	—	—	TBD
Gx12	voltage source	-0.20 V	0.35 V	—	—	TBD
Gx2	voltage source	-0.20 V	0.35 V	—	—	TBD
Dx1	current source	—	—	1.5 mA	30 mA	—
Dx2	current source	—	—	1.5 mA	30 mA	—

Table 6: MMIC bias card output capabilities. This table reflects the actual capabilities of the MMIC bias card. Circuits should not be designed to these specs; rather, design requirements are listed in Section 3.2.1. V_{MIN} and V_{MAX} capabilities listed are as measured at the module gate input, after a 1:7.19 voltage divider on the MAB.

3.2.2 Bias Output Capabilities

The MMIC bias card’s output capabilities are described in Table 6. Transient suppression / over-voltage protection is discussed in Section 3.7.

The MMIC bias card gate output is designed for use with an MAB gate circuit including a 1:7.19 voltage divider (constructed from a 1.00 k Ω 1% shunt resistor and a 6.19 k Ω 1% series resistor) on the gate input. Without this divider, the MMIC bias card will damage the MMICs.

3.2.3 Bias Output Noise Requirements

Noise on the MMIC bias outputs will contribute to the overall noise of the receiver system. This noise consists of two components. The first is common mode on all the gates and drains, such as would result from variations in the module ground potential. This will be referred to as *ripple*. The second component is uncorrelated between control signals and will be referred to simply as *noise* in the following.

For a complete analysis, interactions between bias noise and the scan strategy should be considered. A ripple component synchronous to the scan could be particularly harmful. We therefore distinguish limits for asynchronous ripple and synchronous ripple.

Bias output noise requirements are listed in Table 7. These are rough estimates extrapolated from the Planck-LFI requirements.³

3.3 Bias Control Interface

The bias control interface consists of the digital inputs used to program the DACs that set the bias levels to the modules. Here we describe the protocol for setting the DACs.

³T. Gaier, “Planck-LFI Bias Electronics Requirements– Technical Note,” December, 1999.

Signal	Noise	Async. Ripple	Sync. Ripple
Gx11	0.58 uV/ $\sqrt{\text{Hz}}$	0.58 uV	5.8 nV
Gx12	0.58 uV/ $\sqrt{\text{Hz}}$	0.58 uV	5.8 nV
Gx2	0.58 uV/ $\sqrt{\text{Hz}}$	0.58 uV	5.8 nV
Dxx	0.58 uV/ $\sqrt{\text{Hz}}$	0.58 uV	5.8 nV

Table 7: Bias output noise requirements. *Noise* refers to uncorrelated noise on each line, assumed Gaussian. *Ripple* specs refer to noise signals correlated between output lines, either synchronous or asynchronous to the scan period.

Name	Description
SCK	Serial clock input. Rising-edge active.
DIN	Serial data input.
/CS	Active-low chip select.
/CLR	Active-low asynchronous clear (resets internal DAC registers and output voltages to 0). See note in text regarding /CLR timing.

Table 8: DAC interface signals.

3.3.1 DAC Control Logic

The bias outputs of the MMIC bias circuits are set by an array of Linear Technologies LTC1660⁴ 10-bit 8-channel DACs on the card. The digital inputs of the DACs are daisy-chained so all DACs on each card are controlled via a single serial interface with chip select plus an asynchronous clear input. The interface signals are summarized in Table 8. These inputs are optoisolated on the MMIC bias card as necessary. The electrical properties of these signals are described in Section 3.1.

Two LTC1660 DACs are used for each channel for a total of 28 DACs. The two DACs for each module are numbered 0 and 1. The bias settings assigned to each DAC are indicated in Table 3. The bias DACs are divided into two sets of 7 (A and B) corresponding to the two MABs biased by the card. A single chip select controls the entire MMIC bias card.

The DACs are programmed by writing 28 16-bit DAC control words to the serial interface. The ordering of the words is shown in Table 9. To load the DACs, begin by raising /CS, then lowering SCK. Next, lower /CS to enable data input to the DACs. Place the first data bit on DIN and raise SCK to clock in the first bit. Lower SCK and continue clocking in the rest of the bits on rising edges of SCK. After all 448 bits, lower SCK and raise /CS to complete the write operation. For a detailed timing diagram, see page 8 of the LTC1660 data sheet.

⁴Data sheet available at <http://www.linear-tech.com>.

MAB B					MAB A				
Mod 6 DAC 1	Mod 6 DAC 0	...	Mod 0 DAC 1	Mod 0 DAC 0	Mod 6 DAC 1	Mod 6 DAC 0	...	Mod 0 DAC 1	Mod 0 DAC 0

Table 9: DAC control word ordering. Words transmitted in left-to-right order.

Address/Control				Input Code										Don't Care	
A3	A2	A1	A0	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0	X1	X0

Table 10: DAC control word format. Bits transmitted in left-to-right order (A3 first, X0 last).

Clock Rate	t_{SINGLE} (usec)	t_{TOTAL} (usec)
10 kHz ($t_{\text{CLK}}=100$ usec)	22,500	112,500
100 kHz ($t_{\text{CLK}}=10$ usec)	2,250	11,250
1 MHz ($t_{\text{CLK}}=1$ usec)	225	1,125
8 MHz ($t_{\text{CLK}}=125$ nsec)	28	141

Table 11: Time required to set DAC control words at various serial clock rates.

The format of each 16-bit data word is described in Table 10. The Address/Control input (A3-A0) specifies which of the 8 DAC outputs should be updated with the data in the Input Code field. The address corresponding to each bias control line is indicated in Table 3. Additionally, the Address/Control value “0000” indicates that the Input Code field should be ignored. By setting Address/Control = “0000” in all but one of the 28 words, it is possible to update only a single DAC. For more information about the Address/Control field, see page 10 of the LTC1660 data sheet.

The Input Code (D9-D0) field is the 10-bit value to write to the DAC output specified by the Address/Control field. D9 is the MSB of the value, D0 the LSB. The last two bits of the DAC control word, X1 and X0, are don't-care bits. Their value is ignored.

In addition to the synchronous serial control interface, there is an asynchronous clear input, /CLR. As described in the LTC1660 data sheet, asserting this signal resets the DAC registers to 0 and immediately sets all DAC outputs to 0 Volts. To prevent DAC glitches at power up, the /CLR input is connected to an RC circuit on the MMIC bias card. To accommodate this, a delay of at least 1 ms following power up or assertion of the /CLR must be observed before beginning to load the DAC.

3.3.2 DAC Control Timing

For detailed timing parameters of the DAC interface, please consult the LTC1660 data sheet. Timing values used in this section are from the table on page 3 of the data sheet.

The maximum serial clock rate supported by the LTC1660 is 16.7 MHz ($t_{\text{CLK}} = t_3 + t_4 = 60$ ns). However, because we are daisy-chaining their outputs, we are limited to a maximum serial clock rate of 8.3 MHz ($t_{\text{CLK}} = t_1 + t_8 = 120$ ns) to guarantee valid data to all DACs. In our calculations below, we will also consider slower clock rates.

To write one set of DAC control words to one MAB on the MMIC bias card requires time $t_{\text{SINGLE}} = t_5 + t_6 + t_7 + 14 \times 16 \times t_{\text{CLK}}$. To update every DAC on the bias card requires that this be repeated 5 times, so $t_{\text{TOTAL}} = 5 \times t_{\text{SINGLE}}$. The results for various clock rates are shown in Table 11. Note that in these calculations, t_6 and t_7 were assumed to be 0, and t_5 was set to one clock period.

Note that the above calculations represent only the time to program the digital interfaces of the MMIC bias DACs. Once these are set, the analog outputs will slew to their new voltage settings at the rates indicated on page 3 of the LTC1660 data sheet. A full-scale slew followed by typical settling time adds less than 50 usec to the above times, however, so it is clear that even a 10 kHz serial clock will allow the MMIC biases for all 91 modules to be updated at about 1 Hz, even if the MMIC bias cards must be programmed in series.

Module	Multiplexer Sub-address	Module	Multiplexer Sub-address
A0	1111 xxxx	B0	0111 xxxx
A1	1110 xxxx	B1	0110 xxxx
A2	1101 xxxx	B2	0101 xxxx
A3	1100 xxxx	B3	0100 xxxx
A4	1011 xxxx	B4	0011 xxxx
A5	1010 xxxx	B5	0010 xxxx
A6	1001 xxxx	B6	0001 xxxx
GND SENSE A	1000 xxxx	GND SENSE B	0000 xxxx

Table 12: Receiver module multiplexer sub-addresses. Note that the ground sense signals for the MABs are assigned here.

3.3.3 DAC Transfer Functions

This section will describe the functions that determine the analog output level corresponding to a particular DAC section. This information will not be available until the bias card design is complete.

3.4 Monitor / Multiplexer Interface

For each receiver module, the MMIC bias card provides several signals to be monitored as part of house-keeping. Multiplexers on each MMIC bias card allow selection of one of the monitor signals on the card to be connected to a single analog output to the backplane. This output should be routed to one of the housekeeping boards for digitization.

3.4.1 Monitor Signal Addressing

The MMIC bias card multiplexer is controlled by 8 digital address lines. The 4 MSBs of the address specify from which receiver module to select a monitor signal and the 4 LSBs specify the sub-address of a particular signal. Table 12 and Table 13 list the sub-addresses of the modules and monitor signals.

For example, suppose IDY2MON from receiver module A5 is desired. From Table 12, for module A5 the MSBs of the address should be binary 1010. From Table 13, the sub-address for IDY2MON is 0010, which specifies the LSBs of the address. Thus, the complete multiplexer address for the desired signal is 10100010.

An unused module address (probably 1111xxxx) may be reserved to monitor MMIC bias signals not related to modules. For example, this may be used to allow MMIC bias card power and ground voltages to be monitored.

3.4.2 Multiplexer Output

The output of the multiplexer is passed through an analog optoisolator before being routed to the connector pin. Power and ground for the output side of the optoisolator must be provided separately from the MMIC bias card main power supplies. Power supply requirements are listed in Section 3.5.2.

The housekeeping output signal range is -2 V to $+2$ V.

3.4.3 Interpreting Monitor Outputs

This section will describe the functions needed to convert the raw monitor signal voltages into operating point measurements. This information will not be available until the bias card design is complete.

Signal Name	Multiplexer Sub-address	Description
IDX1MON	xxxx 1000	Drain current, side X, amp 1
IDX2MON	xxxx 1110	Drain current, side X, amp 2
IDY1MON	xxxx 1011	Drain current, side Y, amp 1
IDY2MON	xxxx 0010	Drain current, side Y, amp 2
VDX1MON	xxxx 1001	Drain voltage, side X, amp 1
VDX2MON	xxxx 1111	Drain voltage, side X, amp 2
VDY1MON	xxxx 1100	Drain voltage, side Y, amp 1
VDY2MON	xxxx 0011	Drain voltage, side Y, amp 2
VGX11MON	xxxx 1010	Gate voltage, side X, amp 1, gate 1
VGX12MON	xxxx 0110	Gate voltage, side X, amp 1, gate 2
VGX2MON	xxxx 0001	Gate voltage, side X, amp 2
VGX11MON	xxxx 1101	Gate voltage, side Y, amp 1, gate 1
VGX12MON	xxxx 0101	Gate voltage, side Y, amp 1, gate 2
VGX2MON	xxxx 0100	Gate voltage, side Y, amp 2
+2.5V	xxxx ????	+2.5 V rail
None	xxxx ????	Unallocated

Table 13: MMIC bias monitor signals for each receiver module.

Supply	Current	Max Ripple
+3.3 V	500 mA	5 mV
-5 V	25 mA	5 mV

Table 14: Main power supplies for each MAB. Each supply must be DC-isolated from all other supplies in the system. Two sets are required per MMIC bias card.

3.5 Power and Ground

3.5.1 Main Power Supplies

The MMIC bias card design will completely isolate the power and ground for the two MABs it biases. Separate power supply and ground return pins for each MAB will be provided. These should be powered by a separate, DC-isolated set of power supplies for each MAB (i.e., two sets of power supplies per MMIC bias card).

Each MAB requires two power supplies, one positive and one negative. These are summarized in Table 14.

3.5.2 Optoisolator Power Supplies

In addition to the main power supplies, a separate power supply is required for the input level converters and optoisolators and for the output analog optoisolator. These power supplies may be shared between different MMIC and phase switch bias cards. The specifications are listed in Table 15.

Supply	Voltage	Current	Max Ripple.
Input Opto	+5 V	50 mA	5 mV
Output Opto (+)	+5 V	50 mA	5 mV
Output Opto (-)	-5 V	50 mA	5 mV

Table 15: Optoisolator power supplies for each MMIC bias card. Power supplies may be shared between cards.

3.5.3 Grounding Strategy

Because they provide power and ground connections to the receiver modules, the MMIC bias cards are a critical piece of the QUIET instrument grounding scheme. The grounding scheme must address two issues. First, radiometer signal contamination due to ground loops or other effects must be minimized. Second, interactions between various bias settings should be minimized.

Each MMIC bias card has two separate ground sections, corresponding to the two MABs it biases. Each ground will be connected to all the modules in the corresponding MAB. Because the receiver modules are case-grounded to the feed array, all 91 module grounds will be connected at the top of the dewar. Therefore, unless an isolated supply is used for each module, some ground loops will be introduced. As a compromise, one power supply per MAB bias card allows the area of the ground-loop to be minimized without requiring the expense (power, pin count, financial, etc.) of 91 isolated power supplies.

Because the ground loops are contained to a single MAB, effect of a shared ground at the bias card can be minimized. This should be done by keeping the signal and ground connections between the bias card and its MAB or MABs as close together as possible as they are routed through the backplane and the FPC connector. All connections between a MAB and its bias card should be routed on a single FPC if possible. If multiple FPCs are required per MAB, care should be taken to route the FPCs destined for a particular MAB in close proximity along their entire lengths.

Another concern is the interaction between bias settings on different modules. Because ground traces and wires have non-zero resistance, return current in the ground path between the module and the bias card causes the ground voltage at the module to differ from that at the bias card. If two modules share a ground return, changes in the bias current of one module will cause changes in the bias card's ground voltage relative to the shared ground at the feed array. This will change the bias settings of the other module.

This interaction will be minimized in two ways. First, simply by limiting the number of modules with a shared ground return, the number of possible interactions between modules is limited. Second, a ground voltage sense wire will be connected to the common ground at each MAB. This sense wire is available to the housekeeping system as a monitor output signal at the address listed in Table 12.

3.6 FPC Connection Requirements

The FPC connections to the MABs should be designed to accommodate the requirements in Table 16.

3.7 Transient Protection

Because the transient protection diodes on the MABs will be not provide protection when cooled in the cryostat, transient protection during cooled operation will be provided on the MMIC bias card as follows.

- Gates
 - Zener diode protection on outputs
- Drains
 - Shunt capacitors at outputs to roll off transients

Signal	Max Current	Resistance
Gates	0.1 mA	20 Ω
Dx1	25 mA	2 Ω
Dx2	50 mA	2 Ω
GND	150 mA	0.5 Ω
GND Sense	< 0.1 mA	20 Ω

Table 16: FPC connection requirements.

- No protection diodes — would interfere with bias circuit
- Overall
 - Zener diode protection on regulator output in case of regulator failure

4 Mechanical Interface

The mechanical interface includes the physical size of the cards, the connector type and placement, and the location of any mounting holes or brackets on the board.

4.1 Card Dimensions

The MMIC bias card is 207.95 mm high by 160 mm deep (8.187 inches by 6.299 inches). The board dimensions, connector locations, etc., are based on the IEEE 1101.2-1992 conduction-cooled Eurocard specification, modified to reduce the board height. A mechanical drawing is shown in Figure 1.

4.2 Connectors

There are two 160 pin VME connectors (P1 and P2) along one long edge of the PCB. All connections to the MABs are routed to P2. All other connections (digital controls, power inputs, housekeeping outputs) are routed to P1. Connector pin outs are listed in Appendix A.1

5 Thermal Specifications

5.1 Power Budget

The Complete QUIET Electronics memo estimates 24 mW of “overhead” power will be dissipated by the MMIC bias cards for each module.⁵ This did not account for linear power regulation on the MMIC bias cards, however. Measurements of early prototype MMIC bias circuits showed an actual power overhead of approximately 350 mW per module. It is believed this can be reduced to approximately 220 mW in the final design. This is in addition to power delivered to (and dissipated in) the receiver module itself.


5.2 Heat Sinking

Heat will be conducted out of the MMIC bias card through wedgelock connectors attached to the bare copper regions along the top and bottom (160 mm edges).

⁵C. Bischoff, et al, “Complete QUIET Electronics, 91-Element Prototype Edition, V.2,” March 11, 2005, internal memo, 15-16.

FOR CONDUCTION-COOLED EUROCARDS

IEEE
Std 1101.2-1992



NOTE: P1 and P2 connectors are M55302/131-01 (as specified in MIL-C-55302/131 (4)). These mate with connectors specified in MIL-C-55302/132 (5).

Fig 1
Double-Height Board, Component Side View

Figure 1: Mechanical drawing of board outline and keep-outs.

Modules/Card	Est. Power Dissipation (mW)
7	1,540
14	3,080

Table 17: Per-card overhead power dissipation estimates.

5.3 Thermal Stability

To maintain bias stability, the temperature of the MMIC bias card must be maintained. The target gate output voltage stability is a variation of less than 10 uV. To maintain this output stability, the card temperature must be maintained within 1 degree C.

A Appendices

A.1 Connector Pin Outs

A.1.1 P1 Pin Out

The P1 connector (see Figure 1) carries all signals not connected to the MABs. The pin out is in Tables 18 – 22.

A.1.2 P2 Pin Out

The P2 connector (see Figure 1) carries all signals destined for the MABs. The pin out is in Tables 23 – 27.

A.2 Revision History

- Revision 1.5 (December 14, 2005)
 - Fixed DAC and MUX address tables to match actual hardware.
 - Changed MMIC leg designation from A/B to X/Y.
 - Fixed typos in P2 pinout (AORET→A0RET, second pin labelled A5DA1 corrected to A5DA2).
 - Began this revision history section.

QUIET Phase Switch Bias Interface Specification

Joey Richards*, Mike Seiffert, Steve Smith
Revision 1.4

January 11, 2006

1 Introduction

This document specifies the various interfaces between the phase switch bias cards and the rest of the QUIET system. This phase switch bias design is intended for use in both the 40 GHz (Q-band) and 90 GHz (W-band) systems.

1.1 Schedule

Key dates in the design and production schedule for the MMIC bias card are listed in Table 1.

Item	Start	Finish
Interface definition	—	6/17/2005
Prototype run (qty 2)	In Progress	July
Prototype test/debug	July	7/31/2005
Deliver board to Chicago (qty 1)	—	7/31/2005
Production run (for 19 Q-band elements, 91 W-band elements, plus spares)	TBD	TBD

Table 1: Phase switch bias card design/construction schedule.

2 Mode of Operation

In this section, we describe generally how to interact with the phase switch bias card. Details of the operations required are in later sections of this document. Please take note of the warnings regarding circuit operation in this section.

2.1 Power-Up Sequence

This section will describe any special procedures or precautions necessary to safely power up the phase switch bias card. This is not expected to be particularly sensitive.

*joey@caltech.edu

Interface	Signal Type	Signal Count	Total
Bias Output	Analog	5 per module	105
Bias Control Input	LVDS	7 x 2	14
Monitor/Mux Control Input	LVDS	6 x 2	12
Phase Toggle	LVDS	4 x 2	8
Housekeeping	Analog	1	1
Power	Power	9	9
Ground	Ground	8	8
Total:			157

Table 2: Summary of interface signals for the phase switch bias card.

2.1.1 Reset Delay

At power-up or after the DAC /CLR is asserted and de-asserted, an RC delay circuit will hold the /CLR input to the DACs low. A 1 ms delay should be observed after power-up or de-asserting /CLR before attempting to program the DACs.

2.2 Warnings

2.2.1 Receiver Dead Time

Because the phase switches will produce spurious output during switching events, data collected during switching events should be discarded.

2.2.2 Multiplexer Switching Timing

To prevent contamination of science data, the housekeeping monitor multiplexer address should be changed only during receiver dead time.

2.2.3 Multiplexer Switching Restrictions

In addition to the timing restrictions above, during science data collection the multiplexer address should only be changed after an even number of phase switch toggles (i.e., synchronously with the 4 kHz phase switch clock). This is to minimize the effect of any bias level offset caused by an interaction with the multiplexer by ensuring such an offset is common mode between the two phase switch states.

3 Electrical Interface

The electrical interface consists of digital inputs to control the phase switch bias levels and polarity, a set of analog bias outputs destined for the MABs, and digital inputs and an analog output to select and monitor housekeeping signals. Additionally, we describe the grounding strategy and transient protection in this section.

Each phase switch card will provide bias for the phase switches in the modules on three MABs (21 modules per phase switch card). The interface signal requirements are summarized in Table 2.

Name	Description
PA1	Side A phase switch bias 1
PA2	Side A phase switch bias 2
PB1	Side B phase switch bias 1
PB2	Side B phase switch bias 2
GND	Ground return for module

Table 3: Phase switch bias control signals.

Signal	Absolute Maximum		Operating Range		
	V_{MIN}	I_{MAX}	V_{MIN}	V_{MAX}	I_{MAX}
Pxx	-2.5 V	2 mA	-1.8 V	0.85 V	820 μ A

Table 4: Phase switch bias requirements, from personal communication with T. Gaier.

3.1 Digital Signal Characteristics

All digital inputs to the phase switch bias card are LVDS differential signals. These signals are translated to single-ended digital signals and optoisolated on the phase switch bias card. A power supply and ground must be provided for the LVDS level converter and the input side of the optoisolators (see Section 3.6.2). This power supply and ground must be DC-isolated from the power supplies for the analog sections of the phase switch bias cards (see Section 3.6.3), though it may be common between phase switch cards.

3.2 Bias Output

Each of the two phase switches on each receiver module requires two bias connections. The phase switch bias card outputs for each module are listed in Table 3.

3.2.1 Phase Switch Bias Requirements

The phase switch bias requirements consist of absolute maximum ratings and operating range ratings. The absolute maximum ratings reflect damage thresholds that must not be exceeded. The operating range ratings reflect the minimum bias range needed to allow proper operation of the module. Ideally, the phase switch bias card hardware will provide bias up to a limit that exceeds the operating range rating but is no greater than the absolute maximum rating. The ratings are summarized in Table 4.

3.2.2 Bias Output Capabilities

The phase switch bias card’s output capabilities are described in Table 5. These will be specified to meet the phase switch bias requirements when connected to an MAB through a flexible printed circuit (FPC) that meets the specifications in Section 3.7.

Signal	V_{MIN}	V_{MAX}	I_{MAX}
Pxx	?	?	1.6 mA

Table 5: Phase switch bias card output capabilities.

Name	Description
SCK	Serial clock input. Rising-edge active.
DIN	Serial data input.
/CS	Active-low chip select.
/CLR	Active-low asynchronous clear (resets internal DAC registers and output voltages to 0). See note in text regarding /CLR timing.

Table 6: DAC interface signals.

3.2.3 Bias Output Noise Estimates

TBD

3.3 Bias Control Interface

The bias control interface consists of the digital inputs used to program the DACs that set the bias levels to the phase switches. Here we describe the protocol for setting the DACs.

3.3.1 DAC Control Logic

The bias outputs of the phase switch bias circuits are set by an array of Linear Technologies LTC1660 10-bit 8-channel DACs on the card. The digital inputs of the DACs are daisy-chained so all DACs on each card are controlled via a single serial interface with chip select plus an asynchronous clear input. The interface signals are summarized in Table 6. These inputs are optoisolated on the phase switch bias card as necessary. The electrical properties of these signals are described in Section 3.1

Each phase switch requires only one programmable bias setting so only one DAC channel is required for each module. Thus, there are three eight-channel DACs to support 21 modules per card. The DACs are numbered 0–2. DAC channel assignments are listed in Table 7. The DACs are programmed by transmitting three 16-bit DAC control words to the serial interface. The DAC serial interfaces are daisy-chained so that the first word transmitted sets DAC 2, the second DAC 1, and the third DAC 0. Any time the DACs are updated, three words must be transmitted, even if only one DAC is to be updated.

To load the DACs, begin by raising /CS, then lowering SCK. Next, lower /CS to enable data input to the DACs. Place the first data bit on DIN and raise SCK to clock in the first bit. Lower SCK and continue clocking in the rest of the bits on rising edges of SCK. After all bits have been transmitted, lower SCK and raise /CS to complete the write operation. For a detailed timing diagram, see page 8 of the LTC1660 data sheet.

The format of each 16-bit control word is described in Table 8. The Address/Control input (A3-A0) specifies which of the 8 DAC outputs should be updated with the data in the Input Code field. The address corresponding to each bias control line is indicated in Table 7. Additionally, the Address/Control value “0000” indicates that the Input Code field should be ignored. By setting Address/Control = “0000” in all but one word, it is possible to update only a single DAC. For more information about the Address/Control field, see page 10 of the LTC1660 data sheet.

The Input Code (D9-D0) field is the 10-bit value to write to the DAC output specified by the Address/Control field. D9 is the MSB of the value, D0 the LSB. The last two bits of the DAC control word, X1 and X0, are don’t-care bits. Their value is ignored.

In addition to the synchronous serial control interface, there is an asynchronous clear input, /CLR. As described in the LTC1660 data sheet, asserting this signal resets the DAC registers to 0 and immediately sets all DAC outputs to 0 Volts. To prevent DAC glitches at power up, the /CLR input is connected to an

Module	DAC	Addr	Module	DAC	Addr	Module	DAC	Addr
A0	0	0001	B0	1	0001	C0	2	0001
A1	0	0010	B1	1	0010	C1	2	0010
A2	0	0011	B2	1	0011	C2	2	0011
A3	0	0100	B3	1	0100	C3	2	0100
A4	0	0101	B4	1	0101	C4	2	0101
A5	0	0110	B5	1	0110	C5	2	0110
A6	0	0111	B6	1	0111	C6	2	0111

Table 7: DAC interface signals. Each DAC controls the modules on one MAB (labeled A–C).

Address/Control				Input Code										Don't Care	
A3	A2	A1	A0	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0	X1	X0

Table 8: DAC control word format. Bits transmitted in left-to-right order (A3 first, X0 last).

RC circuit on the MMIC bias card. To accommodate this, a delay of at least 1 ms following power up or assertion of the $\overline{\text{CLR}}$ must be observed before beginning to load the DAC.

3.3.2 DAC Control Timing

For detailed timing parameters of the DAC interface, please consult the LTC1660 data sheet. Timing values used in this section are from the table on page 3 of the data sheet.

The maximum serial clock rate supported by the LTC1660 is 16.7 MHz ($t_{CLK} = t_3 + t_4 = 60$ ns). However, because we are daisy-chaining their outputs, we are limited to a maximum serial clock rate of 8.3 MHz ($t_{CLK} = t_1 + t_8 = 120$ ns) to guarantee valid data to all DACs. Writing one set of DAC control words to the phase shift cards requires time $t_{SINGLE} = t_5 + t_6 + t_7 + 3 \times 16 \times t_{CLK}$ (assuming 21 modules per card). Updating all DAC channels requires time $t_{TOTAL} = 7 \times t_{SINGLE}$.

For example, even with a very slow 10 kHz clock, $t_{CLK} = 100$ usec. Neglecting t_6 and t_7 and assuming $t_5 = t_{CLK}$, we find $t_{SINGLE} = 4.9$ msec and $t_{TOTAL} = 34.3$ msec. In the worst case, even if the phase shifter cards must be programmed in series, programming all five phase shift bias cards requires less than 180 msec.

3.3.3 DAC Transfer Functions

This section will describe a function to predict the bias current that will result from a particular DAC setting.

3.4 Phase Switch Toggle Input

Each phase switch has two inputs that are used to bias diodes inside the switch. At any time, one bias is forward-biased and the other reverse-biased. By switching which diode is forward-biased, the phase switch is flipped by 180 degrees.

This is accomplished in the phase switch bias circuit by switching the two bias outputs between a fixed negative reference voltage (-1.9 V) and a forward bias current set by the bias DAC. Which of the two inputs of the phase switch is forward-biased is controlled using the PCLK toggle input. Table 9 describes the effect of the PCLK input on the phase switch bias.

Each receiver module has two independently controlled phase switches. During ordinary QUIET operation, one phase switch in each module will be kept in a constant state (i.e., its PCLK line will not toggle)

PCLKA/PCLKB	PA1/PB1	PA2/PB2
HI	Fwd biased	Rev biased
LO	Rev biased	Fwd biased

Table 9: PCLK truth table. PCLKA controls PA1 and PA2 on the module, PCLKB controls PB1 and PB2.

Module	Multiplexer Sub-address	Module	Multiplexer Sub-address	Module	Multiplexer Sub-address
A0	1111x	B0	1011x	C0	0111x
A1	1110x	B1	10110x	C1	01110x
A2	11101x	B2	10101x	C2	10010x
A3	11100x	B3	10100x	C3	01100x
A4	11011x	B4	10011x	C4	01011x
A5	11010x	B5	10010x	C5	01010x
A6	11001x	B6	10001x	C6	01001x

Table 10: Receiver module multiplexer sub-addresses.

while the other is switched at 4 kHz. During testing and calibration, it may be necessary to toggle the state of both switches on each module, so this capability is provided.

Each phase switch bias board provides 2 pairs of PCLK lines (4 total), labeled PCLK0a/b and PCLK1a/b. PCLK0 controls even-numbered modules on all three MABs (0, 2, 4, 6) and PCLK1 controls odd-numbered modules on all three MABs (1, 3, 5). The -a and -b designation indicates which phase switch (side a or side b) in each module is controlled.

3.4.1 Phase Switch Toggle Timing

During ordinary QUIET operation, one phase switch in each receiver module will be toggled at 4 kHz. Because the phase switch will generate spurious output during each transition, no data can be collected during the toggle. For reasonable operation, the phase switch bias circuit switching time, t_{SWITCH} , will not exceed 10% of the switching period. There are two transitions in each period of the 4 kHz clock, so $t_{SWITCH} \leq 12.5$ usec.

3.5 Monitor / Multiplexer Interface

The bias current to the forward-biased diode in each phase switch is available as an analog output. Multiplexers on each phase switch bias card allow one phase switch bias current monitor signal on the card to be connected to a single analog output to the backplane. This output should be routed to one of the housekeeping boards for digitization.

3.5.1 Monitor Signal Addressing

The phase switch bias card multiplexer is controlled by 6 digital address lines. The 5 MSBs of the address specify from which receiver module to select a monitor signal and the LSB specifies the sub-address of the particular signal. Table 10 and Table 11 list the sub-addresses of the monitor signals and the modules.

For example, suppose Ib from receiver module B5 is desired. From Table 10, the MSBs of the address for module B5 are binary 01101. From Table 11, the sub-address for Ib is 1, which specifies the LSB of the address. Thus, the complete multiplexer address for the desired signal is 011011.

Signal Name	Multiplexer Sub-address	Description
Ia	xxxxx1	Forward bias current, phase switch A
Ib	xxxxx0	Forward bias current, phase switch B

Table 11: Phase switch bias monitor signals for each receiver module.

Supply	Current	Max Ripple
+5 V	50 mA	5 mV
-5 V	25 mA	5 mV

Table 12: Main power supplies for each phase switch bias card. Use of a single DC-isolated supply for each phase switch bias card is allowed.

3.5.2 Multiplexer Output

The output of the multiplexer is passed through an analog optoisolator before being routed to the connector pin. Power and ground for the output side of the optoisolator must be provided separately from the phase switch bias card main power supplies. Power supply requirements are listed in Section 3.6.2.

The housekeeping output signal range is -2 V to $+2$ V.

3.5.3 Interpreting Monitor Outputs

This section will describe the function to calculate bias current given the measured housekeeping voltage.

3.6 Power and Ground

3.6.1 Main Power Supplies

The phase switch bias card design will completely isolate the power and ground for the three MABs it biases. Separate power supply and ground return pins for each MAB will be provided.

It is allowable to power all three MABs on each phase switch bias card with a single set of power supplies. Each phase switch bias card must have its own set of DC-isolated power supplies, however, and all phase switch bias card power supplies must be DC-isolated from all other power supplies in the system. If a single set of supplies is used, the backplane must route these supplies to the input pins for all three MABs.

Each MAB requires two power supplies, one positive and one negative. These are summarized in Table 12.

3.6.2 Optoisolator Power Supplies

In addition to the main power supplies, a separate power supply is required for the input level converters and optoisolators and for the output analog optoisolator. These power supplies may be shared between different phase switch and MMIC bias cards. The specifications are listed in Table 13.

3.6.3 Grounding Strategy

Because they provide power and ground connections to the receiver modules, the phase switch bias cards are a critical piece of the QUIET instrument grounding scheme. Because there are fewer settings for each module and the phase switch bias current to each module will be very nearly constant, interactions between modules are less of a concern than for the MMIC bias cards. However, preventing noise pick-up or coupling to other circuit elements is still important.

Supply	Voltage	Current	Max Ripple
Input Opto	+5 V	50 mA	5 mV
Output Opto (+)	+5 V	75 mA	5 mV
Output Opto (-)	-5 V	75 mA	5 mV

Table 13: Optoisolator power supplies for each phase switch bias card. Power supplies may be shared between cards.

Signal	Max Current	Resistance
GND	4 mA	2 Ω
Pxx	1 mA	2 Ω

Table 14: FPC connection requirements.

Each phase switch bias card should be powered from a dedicated, DC-isolated power supply with an isolated ground. This ground will be connected to all the modules biased by the card. Because the receiver modules are case-grounded to the feed array, all 91 module grounds will be connected at the top of the dewar. Therefore, unless an isolated supply is used for each module, some ground loops will be introduced. As a compromise, one power supply per phase switch bias card allows the area of the ground-loop to be minimized without requiring the expense (power, pin count, financial, etc.) of 91 isolated power supplies.

Because each phase switch bias card biases the modules on only three MABs, the effect of a shared ground at the bias card can be minimized. This should be done by keeping the signal and ground connections between the bias card and its MABs as close together as possible as they are routed through the backplane and the FPC connector. The three MABs biased by each phase switch bias card should be neighbors in the array to minimize loop area.

3.7 FPC Connection Requirements

The FPC connections to the MABs should be designed to accommodate the requirements in Table 14.

3.8 Transient Protection

The following transient protection will be included on the phase switch bias card. This is intended to protect the phase switches from damage due to transients.

- Forward bias circuit
 - Current-limiting resistor
- Reverse bias circuit
 - Zener diode voltage protection
- Overall
 - Zener diode protection on regulator output in case of regulator failure

4 Mechanical Interface

The mechanical interface includes the physical size of the cards, the connector type and placement, and the location of any mounting holes or brackets on the board.

4.1 Card Dimensions

The MMIC bias card is 207.95 mm high by 160 mm deep (8.187 inches by 6.299 inches). The board dimensions, connector locations, etc., are based on the IEEE 1101.2-1992 conduction-cooled Eurocard specification, modified to reduce the board height. A mechanical drawing is shown in Figure 1.

4.2 Connectors

There are two 160 pin VME connectors (P1 and P2) along one long edge of the PCB. All connections to the MABs are routed to P2. All other connections (digital controls, power inputs, housekeeping outputs) are routed to P1. Pin outs for these connectors are listed in Appendix A.1.

5 Thermal Specifications

5.1 Power Budget

TBD, but small.

5.2 Heat Sinking

Heat will be conducted out of the MMIC bias card through wedgelock connectors attached to the bare copper regions along the top and bottom (160 mm edges).

5.3 Thermal Stability

To maintain bias stability, the temperature of the phase switch bias card must be maintained. The allowable temperature swing is to be determined.

A Appendices

A.1 Connector Pin Outs

A.1.1 P1 Pin Out

The P1 connector (see Figure 1) carries all signals not connected to the MABs. The pin out is in Table 15 – Table 19.

A.1.2 P2 Pin Out

The P2 connector (see Figure 1) carries all signals destined for the MABs. The pin out is in Tables 20 – 24.

D.2 Module Protection Circuitry

In this section we present a memorandum describing a recommended design for protection circuitry to prevent damage to the QUIET polarimeter modules when the cryostat was cold. The polarimeter modules are installed in the cold section of the cryostat and installed on a printed circuit board called a module attachment board (MAB) with connectors for the flexible printed circuit (FPC) cabling that extended out to the cryostat interface plate. Circuitry to protect the polarimeter module from static discharge was installed on the MAB, but much of the protection was provided by diodes that were not functional at cryogenic temperatures. Thus, if the connectors on the warm side of the interface plate were exposed to a static discharge during maintenance, polarimeter module damage could result. In this memorandum, Robert Dumoulin and I propose an additional set of “warm protection” circuitry to be installed on the FPC cabling to address this risk. It is presented with its original formatting and pagination.

Static Protection for Cold QUIET Modules

Robert Dumoulin* and Joey Richards†

July 11, 2006

Abstract

We propose that static protection circuits be integrated into a modified FPC design to protect the QUIET modules from damage due to static discharges when the MAB protection circuits are cooled.

1 Introduction

The QUIET modules contain several circuit elements that are susceptible to damage due to static shocks. The module attachment boards (MABs) and bias cards each contain protection circuitry to mitigate the risk of damage. However, the protection in these areas is insufficient to fully protect the modules during operation. The protection diodes on the MABs are not rated for cryogenic operation and are expected to fail when the modules are cooled. The redundant protection on the bias cards will continue to function, but will not provide any protection if the flexible printed circuit (FPC) connections are unplugged for maintenance.

We propose that the FPCs be modified to include protection circuitry that will be kept in the warm region of the cryostat. This will guarantee protection of the modules at all times, even when cooled for operation.

2 Existing Protection Circuits

Currently, protection is installed on both the MAB and on the bias cards. We propose that this protection be kept in place. For clarity, we summarize here the existing protection as we understand it.

2.1 MAB Protection

Figures 1–3 show the protection circuits to be included on the MAB for the gate, drain, and phase switch pins. The diode types are not specified here, we assume they will be selected appropriately to provide protection without impeding the operating range of the bias signals. This information is based on the schematics in the July 7 QUIET MAB Interface Specification¹ and recent communication with Dan Kapner about the voltage divider ratio.

It should be noted that the protection diodes on the gate circuits are on the input side of the 5.1:1 voltage dividers on those lines. This is a change from the initial MAB implementation which placed the diodes directly on the module gate. This change was previously suggested in order to reduce the size of voltage spikes that could appear on the module gate.

* *robert@phys.columbia.edu*

† *joey@caltech.edu*

¹The QUIET MAB Interface Specification indicates that BAV99 diodes will be used. However, it's not clear that these will provide sufficient gate voltage range in the circuit as drawn.




Figure 1: Schematic of gate protection to be installed on MAB. Note that the diode protection appears on the input side of the 5.1:1 voltage divider.




Figure 2: Schematic of drain protection to be installed on MAB.




Figure 3: Schematic of phase switch protection to be installed on MAB.




Figure 4: Schematic of gate protection on the MMIC bias card. The LED is used to provide a sharper turn-on knee when the circuit begins limiting the voltage.




Figure 5: Schematic of drain protection circuit on the MMIC bias card.

In our proposal, we add an additional reason to place these diodes in the position shown in Figure 1. Diode protection placed upstream of the MAB will necessarily occur before the voltage divider². Thus, diodes compatible with this position must be specified for the added protection and it is sensible to use the same components on the MAB itself.

2.2 Bias Card Protection

The protection circuits on the MMIC bias card are shown in Figures 4 and 5. These circuits come from the June 26 schematics used in the recent design review.

Figure 6 shows the protection circuit on the phase switch bias card. This circuit is from the schematics used for the 2005 phase switch bias card prototype build.

²It is undesirable to move the voltage divider on to the FPC because this will increase the impact of noise picked up by the portion of the FPC following the divider. The divider should be placed as near to the module gate pin as possible to minimize noise susceptibility.




Figure 6: Schematic of protection circuit on the phase switch bias card. The resistors are both 1 k Ω .

3 Proposal for Additional Protection

We propose incorporating an extra set of “warm-protection” circuits for the modules integrated into the FPCs themselves. These circuits should be located on the portion of the FPC that is inside the cryostat to minimize flexure of the component-bearing portion of the FPC. Because we do not have very many options for small protection diodes, we will most likely have to widen the cable to accommodate all of the components.

3.1 Widening the Cable

Each trace for the FPC is currently about 0.5 mm. In order to accommodate a diode in a SOT-363 package, we would have to widen each trace (at a local region/point) to about 2.5 mm to accommodate them. We would also have to put resistors/capacitors and ground links for the cable to function as protection circuitry. According to engineers at Cirexx, our FPC manufacturer, it is not a problem to widen the cable for any region/length (leaving the end connectors the same size). This can probably be done neatly, leaving most of the cable unchanged with a small thick region where the components are.

3.2 Rigidity of the Cable

Since the cable will have quite a few components in a particular region, the cable must to be fairly rigid for this small component bearing of the cable (perhaps several inches). This should not present too many problems, as we should be able to find a small warm section where the cable does not need very much flexibility.

3.3 Additional Cost of Protection on FPC

Until we submit designs to Cirexx, we cannot get a specific quote for the cost of this additional work. According to two engineers at Cirexx, widening the cable and adding components (we would have to stuff the FPCs ourselves) will probably cost an additional 15 – 20%, so long as we do not change the number of layers on the FPC. The maximum additional cost of this would be about an additional 25% over the present cost. According to the November 15, 2005, FPC Memo, the cost of 5 sets of 5 FPCs is about \$3400, so the additional cost for a component-bearing FPC will not likely exceed \$850 for this quantity.

If additional FPC layers must be added, the cost increase will be much greater, probably resulting in a 100% increase in price. It is our understanding that the existing FPC uses two routing layers with signals on one layer and ground connections on the other. This is convenient for routing because all protection circuits will require a connection to ground.

3.4 FPC Protection Circuits

We suggest that the same protection circuits (minus resistors) used on the MAB be replicated on the FPC. During cold operation, this should pose no problems. During warm operation, the parallel protection diodes on the FPC and on the MAB will draw extra current (because the current through the “off” diodes is not exactly zero). This may be inconvenient for the drains in particular, but should only affect warm operation.

4 Rejected Option: Additional PCB

Initially, we considered the option of introducing a new PCB that would contain the new warm protection circuitry. This PCB would sit inside the cryostat and connect to a shortened FPC from the MAB on one side and connect either to another FPC or to another type of high-density cable on the other. This introduces either two extra sets of FPC connections or the challenge of identifying a more convenient, mechanically robust high-density cable. Because of this and because we believe the cost of integrating the protection circuits directly on the FPC to be similar to (or less than) that of an additional PCB with connectors, we

have rejected this option. However, in the unlikely event that routing the protection circuitry will require additional FPC layers, this option may be worth reconsidering.

5 Conclusions

We believe that additional protection circuitry is necessary to guard against static damage to the QUIET modules when the modules are cooled in the cryostat. It appears that a duplicate set of protection circuits can be integrated into the FPC “cables” at a reasonable cost. We therefore propose that the FPC design be modified to include these protection circuits.

D.3 Housekeeping Measurement Procedures

This section contains a memorandum describing a recommended set of procedures for measuring and calibrating the housekeeping data outputs of the bias electronics. This work was based on experience from testing prototype bias electronics cards in the lab. This memorandum is presented with its original formatting and pagination.

QUIET MMIC Bias Housekeeping Calibration Procedure

Joey Richards*
Revision 1.1

February 1, 2008

1 Introduction

This note describes a basic procedure for calibrating the QUIET MMIC bias card housekeeping output. At this time, careful analysis of unit-to-unit variation, temperature sensitivity, and other potentially important effects has not been done. It is not entirely clear what degree of absolute accuracy is required for these outputs. Module performance is quite likely sensitive to variation within the errors that will result from the calibration procedure described here. Based on testing with one module site, it appears that this method will result in measurement accuracy of approximately 2-3 mV for gate voltages, 5 mV for drain voltages, and 1 mA for drain currents (up to 1 V drain voltage).

2 Calibration Procedure

For quick reference, the following calibration procedure is suggested.

1. Measure analog optoisolator voltage transfer function, f_{opto} by injecting a signal into the GNDSENSE input.
2. Measure impedance of each MAB's ground connection and each drain bias line. For the accuracy described here, 1 Ω precision is sufficient—it may be possible to achieve this without measurement. If it's convenient to measure gate bias lines as well, these may be useful for improving precision in the future.
3. Use f_{opto}^{-1} and equations (1), (8), and (9) to convert housekeeping measurements into actual bias parameters.

3 Circuit Analysis

The key component to be calibrated is the analog optoisolator. This is common to all modules on a MAB, so only two sets of measurements per MMIC bias card are required. Once this curve is known, it is straightforward to compute estimates of the actual gate and drain bias parameters from the housekeeping measurements. For the level of precision described here, this measurement need only be made once per MAB-worth of circuits; no measurements of individual bias circuits are required.

*joey@caltech.edu




Figure 1: Housekeeping analog optoisolator voltage transfer curve, f_{opto} measured by injecting a signal into the GNDSENSE input. The Y axis indicates the voltage on the housekeeping output line.

3.1 Optoisolator

As of revision 2 of the MMIC bias card design, the optoisolator circuit’s voltage response is nonlinear. The circuit consists of a separate LED/photo-diode circuits for positive and negative voltage swings. These circuits are not perfectly matched, so the positive and negative regions have somewhat different slopes and offsets, with a non-linear transition region slightly below zero output. Figure 1 shows an example of the curve.

This curve should be measured twice for each MMIC bias board: once for each MAB. This can most easily be done by injecting signal on the GNDSENSE input to the bias board. This signal should be referenced to the bias board ground, not the MAB ground. This ground should be one of the bias circuit grounds, labeled as A0–A9 on the MMIC bias board schematic. A connection can be obtained by soldering a lead to the bias board or by connecting to, e.g., pin Z15 of P2 (A0GND) for MAB A or pin Z18 of P2 (B0GND) for MAB B. To cover the entire range of the optoisolator, the input should be swept from -3 V to +3 V . Measuring with small step size is suggested, particularly in the transition region near zero output.

At JPL, we swept the input with a slow triangle wave and recorded the housekeeping output over several periods while recording the actual input voltage on another ADC channel. A table was constructed from these data. Samples are linearly interpolated between nearest neighbors to fill in missing points. This appears to be sufficient.

3.2 Gate Calibration

Figure 2 shows a simplified schematic of a single gate output. R1 and R2 are on the MMIC bias board, R_{lead} and R_{gnd} represent the FPC resistances, and R3 and R4 make up the voltage divider on the MAB. Table 1 lists the circuit component values. R_{lead} and R_{gnd} will depend on the test setup and should ideally be measured for each channel. R_{gnd} is common to an entire MAB.

Circuit analysis yields the following relationship between the housekeeping output and the actual gate

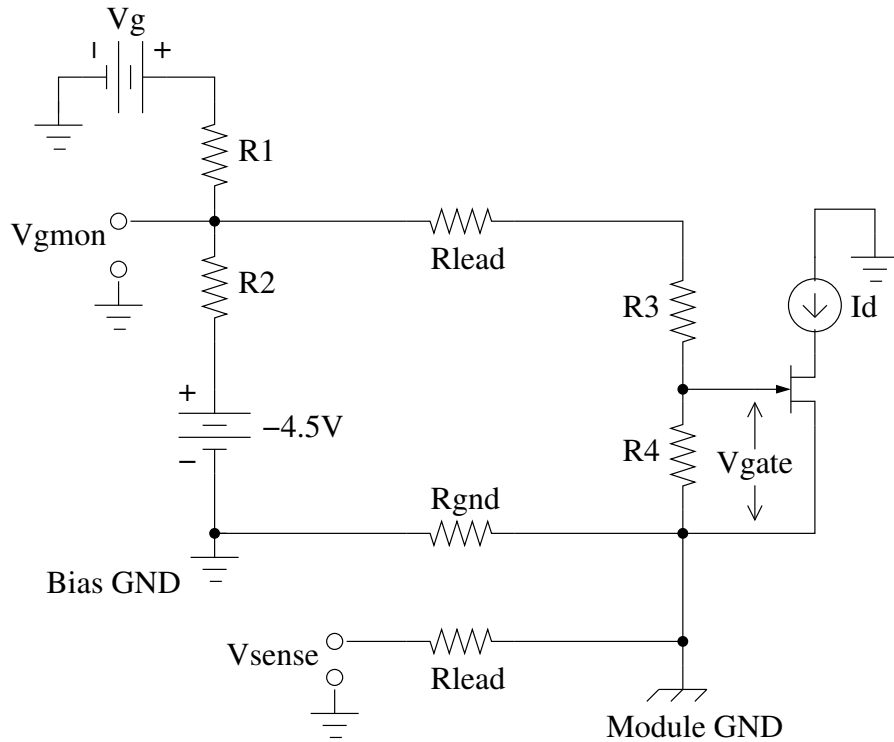


Figure 2: Schematic of a single gate control and housekeeping circuit, including lead resistances. V_g represents the DAC output that sets the gate voltage. The labeled -4.5 V power supply value is actually -4.56 V for W-band and -2.81 V for Q-band.

Resistance	W Band	Q Band
R1	499 Ω	1.00 k Ω
R2	3.74 k Ω	2.00 k Ω
R3	4.12 k Ω	499 Ω
R4	1.00 k Ω	1.00 k Ω

Table 1: Gate circuit resistor values for W- and Q-band. R_1 and R_2 values are quoted from JPL schematics. R_3 and R_4 values for W-band are from the March 1, 2007, Version 2 W-Band MAB schematic posted on the QUIET Wiki. For Q-band, these are from *Revised Q-band MAB Design*, R. Dumoulin, January 31, 2008.

Quantity	Value
\mathbf{R}_{ee} (Drain 1)	24.9 Ω
\mathbf{R}_{ee} (Drain 2)	11.3 Ω
\mathbf{V}_{supply}	2.5 V

Table 2: Drain circuit values for both W- and Q-bands. V_{supply} can be monitored on a housekeeping channel if desired.

voltage.

$$V_{gate} = G \times (V_{gmon} - V_{sense}) - \epsilon \times V_{sense} , \quad (1)$$

where

$$G = \frac{R_4}{R_{lead} + R_3 + R_4 + R_{gnd}} , \quad (2)$$

and

$$\epsilon = G \times \frac{R_{gnd}}{R_{lead} + R_3 + R_4 + R_1 \parallel R_2} . \quad (3)$$

Remember that both V_{gmon} and V_{sense} were measured through the optoisolator, so must be corrected using the inverse of the optoisolator transfer function described in Section 3.1. That is,

$$V_{gmon} = f_{opto}^{-1}(V_{gmon,measured}) \quad (4)$$

$$V_{sense} = f_{opto}^{-1}(V_{sense,measured}) . \quad (5)$$

Figure 3 shows the error in gate voltage measured using this approach. Each gate on one output was swept several times with varying MAB ground offsets injected with a floating lab power supply. The effect of ground offset has clearly been removed quite effectively.

3.3 Drain Calibration

Figure 4 shows a simplified drain circuit. Again, the optoisolator transfer function should be used to correct the measured housekeeping outputs.

$$V_{Idmon} = f_{opto}^{-1}(V_{Idmon,measured}) \quad (6)$$

$$V_{Vdmon} = f_{opto}^{-1}(V_{Vdmon,measured}) . \quad (7)$$

The drain current should be estimated first because it is needed to estimate the drain voltage. It can be found from the following equation.

$$I_{drain} = \frac{V_{supply} - V_{Idmon}}{(1 + h_{FE}^{-1}) \times R_{ee}} , \quad (8)$$

where h_{FE} is the current source transistor's forward current transfer ratio. For the 2N2907A resistor used in the revision 2 bias circuits, h_{FE} is at least 100, so the h_{FE}^{-1} term can be safely ignored.

The drain voltage is simply V_{Vdmon} corrected for the voltage drops across R_{lead} and R_{gnd} . V_{sense} allows us to directly measure the R_{gnd} drop, and we know the current through R_{lead} is the same as the drain

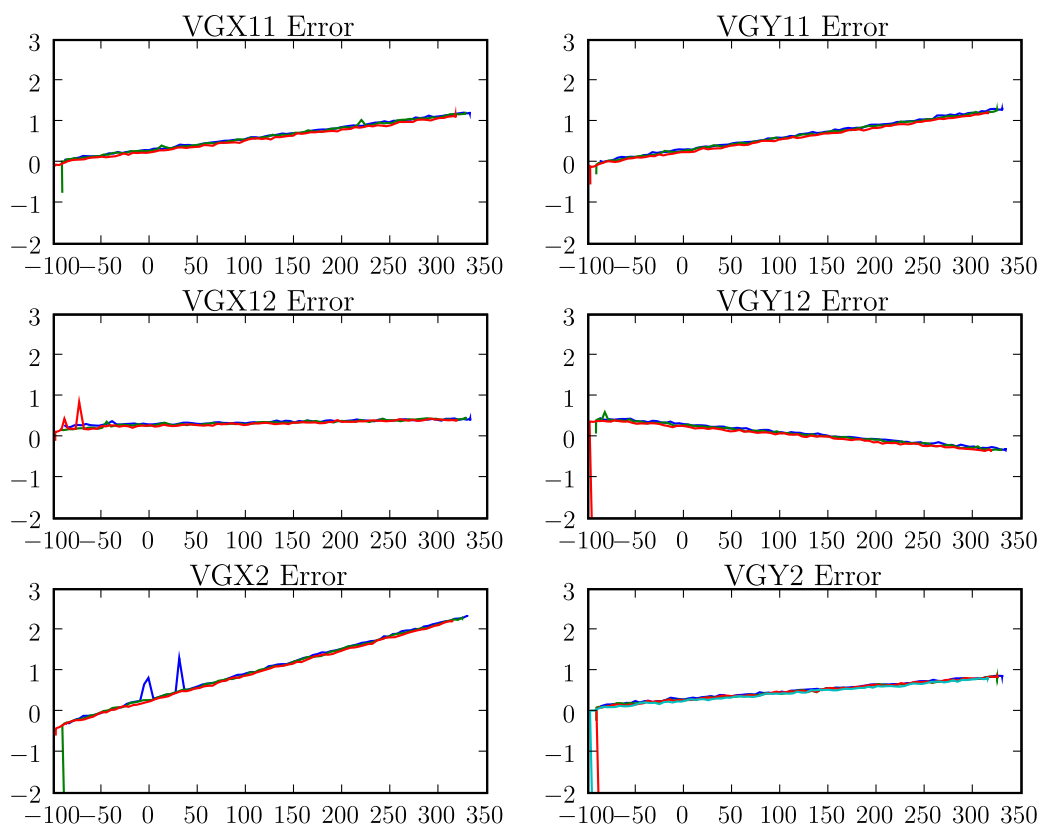


Figure 3: Gate measurement errors for a W-band single-board MAB with no module installed. The X axes indicate the actual gate voltage in mV. The Y axes indicate the error, also in mV. Each plot contains three sweeps, corresponding to different MAB ground voltage offsets. These were, approximately, 17 mV for blue, 50 mV for green, and 120 mV for red.

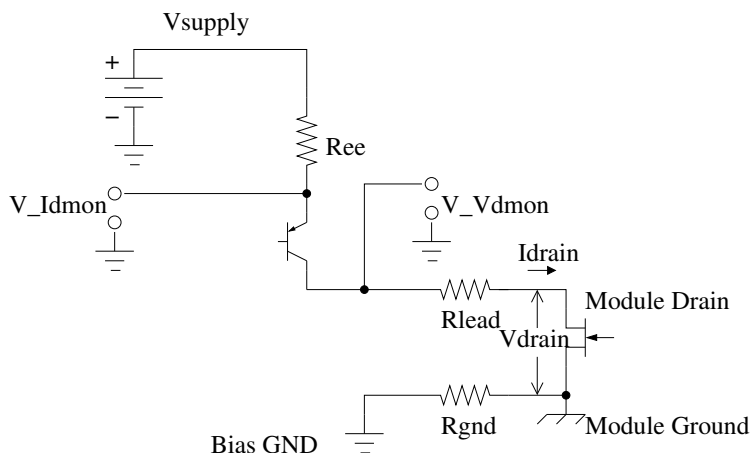


Figure 4: Simplified drain schematic. V_{Idmon} and V_{Vdmon} are the housekeeping outputs. The DAC output that controls the current is not shown; it sets the voltage at the base of the transistor shown here.

current, ignoring current through the protection diodes (not shown). For drain voltages below about 1 V, this is a reasonable assumption. The drain voltage estimate is then

$$V_{\text{drain}} = V_{Vdmon} - V_{\text{sense}} - I_{\text{drain}} \times R_{\text{lead}} . \quad (9)$$

A reasonable estimate of R_{lead} is necessary since significant current flows through the lead wire. However, since R_{lead} should only be 2 Ω according to the MMIC bias card specification, it should not be difficult to know this to within, say, 1 Ω . That would result in an error of 1 mV per mA of drain current, which should be acceptable.

Figure 5 shows measured errors between the actual drain current and voltages and the estimates from the housekeeping output. The large deviations at high drain voltage result from the protection diodes turning on. These cause overestimates of the drain current, which result in errors in the drain voltage as well. Several ground offset voltages were used in the test, which is why there are several different protection turn-on curves in each plot.

This test was run with a MMIC bias card connected to a warm MAB with protection via 2 Ω cryowires. Because of the drops in the cryowires, the voltage seen by the protection on the MMIC bias card is higher than the drain voltage. As a result, current through that protection will be much greater than that through the on-MAB protection. This will be even more true when the MAB's protection is cold.

Figure 6 plots the drain errors as a function of V_{Vdmon} . A single function of V_{Vdmon} can describe these error curves quite well, which confirms the expectation that the MMIC board protection is the dominant current “thief.”

If operation above a drain voltage of 1 V is necessary, it may be necessary to correct for the current lost to the protection circuits. Fortunately, this should be straightforward as long as all warm protection circuits are connected to the MMIC bias card via low impedance wires. It should apparently be possible to reach a few mA accuracy on the drain current using a single correction curve with V_{Vdmon} as its input.

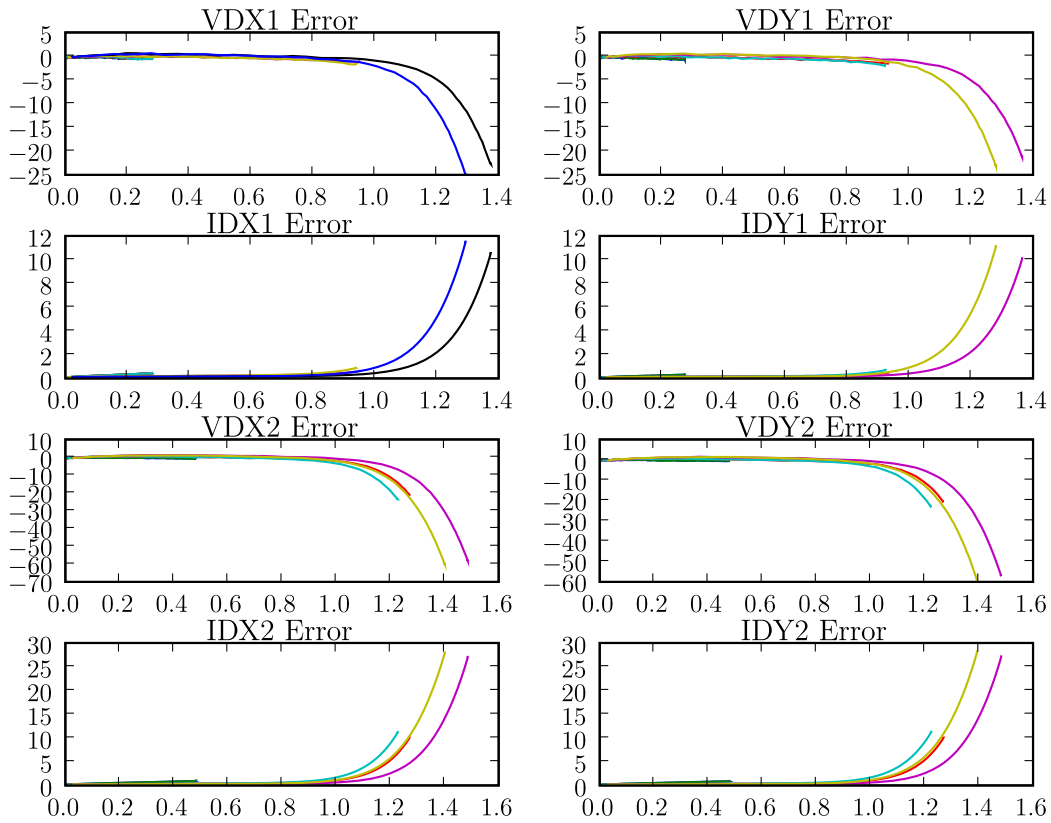


Figure 5: Drain voltage and current estimate errors as a function of actual drain voltage. Voltages are plotted in mV, currents in mA, against drain voltage in Volts. Drains were simulated with a resistor to ground. Resistors of 10.0Ω , 34.8Ω , and 81.2Ω were used. Additionally, the MAB ground was offset by several values between 3 and 124 mV for each resistor value.




Figure 6: Same data from Figure 5, now plotted versus the housekeeping drain voltage monitor, $V_{V_{dmon}}$. It is seen that the deviation, caused by diode protection stealing current from the drains, is well-described as a single function of $V_{V_{dmon}}$. This suggests that the protection on the MMIC bias board is primarily responsible for the “missing” drain current, as is expected.

D.4 Receiver Characterization

In this section we present four memoranda describing detailed study of QUIET polarimeter module performance, with emphasis on study of the $1/f$ noise produced by the QUIET polarimeter modules. These memoranda discuss measurement procedures that were carried out using a test apparatus including a cryostat, bias electronics, and digital readout electronics that I assembled with assistance from Dr. Michael D. Seiffert. Components for this test apparatus included a prototype QUIET bias electronics board set with a heavily modified interface to connect to a single computer workstation running a custom bias board control interface program I developed. A preamplifier I constructed, based on the design described in section D.5, was used for some tests. Existing preamplifiers and a manually controlled bias electronics box were used for some tests. In some cases an existing data acquisition and bias control system were used rather than one of my design and construction. The memoranda describe work that was primarily carried out by me, with guidance and consultation with Dr. Seiffert, particularly with regard to the measurement design and method. These memoranda are presented with their original formatting and pagination.

Module Noise Performance Draft Report

W-Band Module 9

Revision 1.0

Joey Richards* and Mike Seiffert[†]

February 23, 2006

1 Introduction

This draft report describes the results of noise performance testing on QUIET W-band receiver module 9. Demodulated time series data were collected and analyzed to characterize the performance of the module. We have measured the noise-effective bandwidth, the 1/f noise knee frequency, and we have begun to investigate processing techniques for correcting imbalances in the RF system. In these tests, the module was at room temperature with its RF inputs unconnected, looking into the room.

From our noise measurements, we find the module's noise-effective bandwidth to be 12–13 GHz, compared to an estimated 16–17 GHz calculated from its gain curve.

We estimated the 1/f noise knee frequency for several configurations. For a single diode with no differencing, the knee is at about 30 Hz. For a single diode differenced between phase switch states, the knee drops to about 300 mHz. For a weighted double-difference between two diodes, we find a knee frequency of about 100 mHz.

Weighting the two phase states for each diode substantially lowers the 1/f noise knee frequency. With a weighted difference, single-diode knee frequencies are reduced to about 100 mHz. The doubly-weighted double-difference 1/f knee is at about 50 mHz.

2 Data Acquisition System

Time series data were collected from two of the four detector diodes on the module. These were the two diodes corresponding to Q polarization measurements, designated diode 1 (D1) and diode 4 (D4). These data were collected using a data acquisition system originally designed for the Planck project.

2.1 Data Acquisition System

The data acquisition system consists of a low-noise preamp for each diode, followed by an analog integration circuit connected to an ADC in a PC. The PC stores digitized data from the integrator and controls the phase switch clock.

Each preamp consists of an Analog Devices OP37 low-noise op amp gain stage (voltage gain of 100), followed by an OP27 unity gain buffer. The preamp input is DC-coupled directly to the detector diode. No DC bias is provided to the diode for room-temperature operation.

The integrator circuit is based on a Burr-Brown (now TI) IVC102 integrator chip. The circuit integrates the two phase switch states separately. The phase switch clock is driven at 4096 kHz with 50% duty cycle. A blanking time of approximately 10 microseconds is inserted at each phase switch clock transition. Following

**joey@caltech.edu*

[†]*Michael.D.Seiffert@jpl.nasa.gov*




Figure 1: Rising edge of phase switch clock (top trace) followed by falling voltage on phase switch diode (bottom trace). Delay from rising edge to stable bias is approximately 18 microseconds.

this blanking interval, the active integration channel is reset and the preamp output is then integrated, held, and digitized by the data acquisition PC.

The integrator provides additional voltage gain of about 10. Additionally, each integration channel has a DC offset trim pot. Before each data run, this is adjusted to zero the DC level of each channel.

2.2 Module Bias

Module bias is set using a pair of QUIET bias cards, one for the MMICs and one for the phase switches. These are set using a Visual Basic application running on the data acquisition PC. For these tests, the bias levels were set at the start of data collection and were not monitored or adjusted during operation.

The MMIC bias values were set to approximate the recommended room-temperature bias settings provided by Todd Gaier.¹ The phase switch bias was adjusted to balance the DC detector output between the two phase states as much as possible.

2.3 Known Issues

There are several known problems with the data acquisition process used in tests so far. The effect of these problems is to reduce (worsen) the figures of merit we present in this report. The results presented here are thus lower bounds on the actual performance of the module.

2.3.1 Integrator Blanking Incorrectly Positioned

The 10 microsecond blanking interval of the integrator circuit is positioned to coincide with the phase switch clock transition. However, due to slow optoisolators on the phase switch bias card, there is a long (approximately 20 microsecond) delay between this transition and the actual phase switch event. As a result, the blanking interval probably misses the phase switch interval entirely. Figures 1 and 2 show this delay.

2.3.2 DC Level Underestimated

In some of the calculations that follow, the ratio $\frac{\Delta V}{V}$ is used to estimate the noise-effective bandwidth. Because of the incorrect position of the blanking interval, the DC voltage is underestimated by our acquisition system. This results in an underestimate of the noise-effective bandwidth.

¹Data from *module9_02-06-2006_300K-2.xls*.




Figure 2: Falling edge of phase switch clock (top trace) followed by rising voltage on phase switch diode (bottom trace). Delay from falling edge to stable bias is approximately 20 microseconds.

2.3.3 Open-Loop Biasing

We have not connected to the housekeeping outputs of the bias cards or otherwise instrumented the bias circuits to monitor the bias settings during operation. Although care was taken to achieve the recommended bias settings for the module, it is possible that these settings have drifted with temperature.

Also, the recommended drain bias was specified in terms of both current and voltage settings. We biased the drains by first setting the gate voltage, then adjusting the bias current until the drain voltage reached the recommended level. We did not directly measure the drain current, however.

3 Noise-Effective Bandwidth

The sensitivity of the QUIET radiometer is given by

$$\frac{\Delta T}{T} = \frac{\Delta V}{V} = \frac{2}{\sqrt{\beta\tau}} \quad (1)$$

where τ is the integration time. β is the noise-effective bandwidth (also known as the reception bandwidth) of the receiver. The noise-effective bandwidth can also be estimated from the receiver's gain curve,

$$\beta = \frac{[\int G(f) df]^2}{\int [G(f)]^2 df} \quad (2)$$

where $G(f)$ is the gain of the receiver at frequency f .

We calculated the noise-effective bandwidth both from noise measurements and from a measured gain curve of the module. The results are in Table 1 and notes on the calculation methods follow.

3.1 Bandwidth from Noise Measurements

To calculate the noise-effective bandwidth from noise measurements, Equation 1 was solved for β , giving

$$\beta = \frac{4}{\tau \left(\frac{\Delta V}{V}\right)^2}. \quad (3)$$

Diode	β_{noise}	β_{gain}
D1	13.2 GHz	17.4 GHz
D4	12.1 GHz	15.9 GHz

Table 1: Noise-effective bandwidths. β_{noise} was calculated from noise data, β_{gain} was calculated from the gain curve.

Our data acquisition system reports the DC voltage level measured on the detector diode.² This was used for V in the equation. For ΔV , we computed the noise spectrum and used the white noise amplitude in $\text{nV}/\sqrt{\text{Hz}}$. Because of this choice of units, our effective total integration period is 1 second. This is divided into two phase states, so this is divided by two. Finally, approximately 10% of the integration time is lost to blanking during the phase state transition, so we use $\tau = 0.45$ seconds in our calculations.³ The data used for these calculations is given in Table 2.

Diode	V	ΔV	τ
D1	0.709 mV	18.4 nV	0.45 sec
D4	0.947 mV	25.7 nV	0.45 sec

Table 2: Data for noise-effective bandwidth calculation.

3.2 Bandwidth from Gain Curves

The gain curves for the diodes are shown in Figures 3 and 4. These were used to approximate the integrals in Equation 2. Only the region between 83 GHz and 103 GHz was included in the integral. For calculation, the gain curves were approximated by piecewise linear functions connecting local extrema.

4 Noise Data

The following analysis used two noise data sets. Each data set represents approximately 15 minutes of data.

For the first data set (*6_2_17_11_38.dat*), the data acquisition system was connected to D1 and D4 of module 9. The module was biased to the recommended room-temperature settings. The module’s RF inputs were unconnected, looking into the room. This data set was used for most of the measurements described below.

For the second data set (*6_2_17_14_47.dat*), the data acquisition system was again connected to D1 and D4 of module 9. In this case, the MMIC and phase switch biases were set to 0. This data set was used for the baseline spectrum computations only.

4.1 Data Acquisition Baseline

Figures 5 through 7 were computed from input-free data and show the baseline noise characteristics of the data acquisition system. The single-difference noise floor of each input channel is at approximately $5 \text{ nV}/\sqrt{\text{Hz}}$.

Figure 7 illustrates that the noise on these channels is at least reasonably uncorrelated. Differencing the two input streams results in a signal with almost exactly $\sqrt{2}$ times the individual channel white noise levels.

²This is probably an underestimate, see Section 2.3.2.

³The actual integration time is actually probably a bit lower than this, see Section 2.3.1.




Figure 3: Gain curve for D1. Relevant curve is the dark blue plot, “A1 & B1” in the legend.

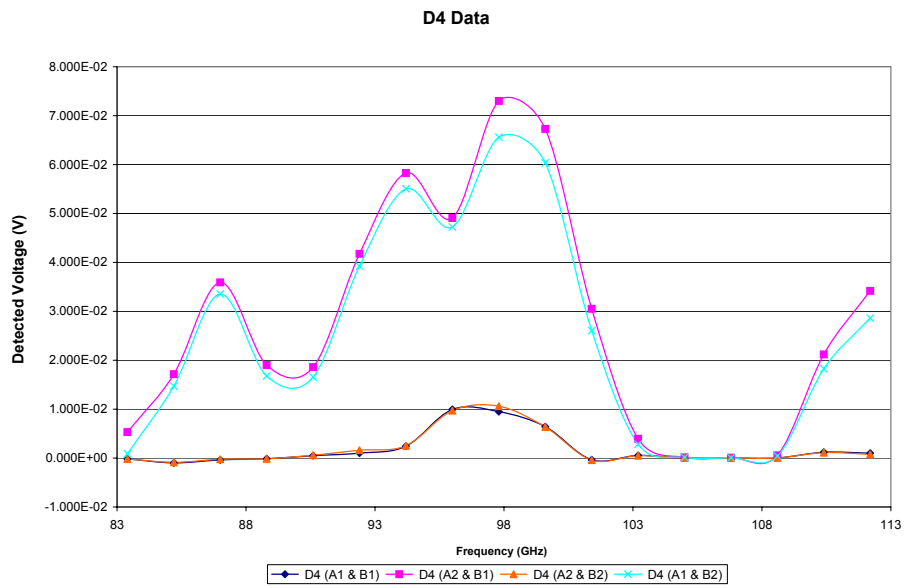


Figure 4: Gain curve for D4. Relevant curve is the light blue plot, “A1 & B2” in the legend.




Figure 5: Channel A data acquisition system noise baseline. Input connected to D1 of unpowered module. Computed from difference of “plus” and “minus” phase states. White noise at 5.2 nV/sqrtHz.




Figure 6: Channel B data acquisition system noise baseline. Input connected to D4 of unpowered module. Computed from difference of “plus” and “minus” phase states. White noise at 5.1 nV/sqrtHz.




Figure 7: A-B difference noise spectrum. White noise level increased by approximately $\sqrt{2}$ from individual channels, indicating that the channel noise is uncorrelated. White noise at 7.2 nV/sqrtHz.

4.2 Time Series Comparison

Figure 8 shows the effect of single-differencing on the data streams. The D1 data shown are not significantly different from the D4 data, so only one plot is included.

4.3 Single Detector Diode Noise Spectra

Figures 9 and 10 show the noise spectra collected in the “plus” phase state on each diode. These did not differ significantly from the “minus” phase state spectra.

4.4 Single-Difference Noise Spectra

The spectra in Figures 11 and 12 were computed from singly-differenced data. No weighting was done to correct imbalances in the phase switch states.

Figures 13 to 16 show single differences between the two diodes. Weighting did not substantially improve the noise level or knee frequencies, so was not used in these plots.

4.5 Double-Difference Noise Spectra

Figure 17 was computed from a double-difference between diodes D1 and D4. A single data set for each diode was computed by differencing the two phase states. These data were then differenced between D1 and D4 to produce the spectrum.

Because the signal levels from the two diodes were not identical, the signal from each diode was weighted prior to combining into the difference signal. It seems that inversely weighting each diode by its mean white noise level yields the best combination, both in terms of double difference white noise level and $1/f$ knee frequency. Inversely weighting by the DC diode level gives an almost identical result.

No weighting was applied to the individual phase state signals on each diode. Only the second difference in the double-difference was weighted.




Figure 8: Comparison of time series data from D1 in the two phase states. Also shown is the single-difference data, demonstrating the reduction in $1/f$ noise. Red (top) is the “plus” phase state, blue (middle) is the “minus” phase state, and black (bottom) is the difference. The data are plotted with arbitrary DC offset and have been downsampled to 10 Hz.




Figure 9: Noise spectrum of D1 “plus” state with no differencing. “Minus” state spectrum is essentially identical. Spectrum computed with resolution of 100 mHz.




Figure 10: Noise spectrum of D4 “plus” state with no differencing. “Minus” state spectrum is essentially identical. Spectrum computed with resolution of 100 mHz.




Figure 11: Single-difference noise spectrum from D1, “plus” state minus “minus” state. White noise level is $18.4 \text{ nV}/\sqrt{\text{Hz}}$, indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.




Figure 12: Single-difference noise spectrum from D4, “plus” state minus “minus” state. White noise level is $25.7 \text{ nV}/\sqrt{\text{Hz}}$, indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.




Figure 13: Single-difference noise spectrum of D1 “plus” state minus D4 “plus” state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.




Figure 14: Single-difference noise spectrum of D1 “plus” state minus D4 “minus” state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.




Figure 15: Single-difference noise spectrum of D1 “minus” state minus D4 “plus” state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.




Figure 16: Single-difference noise spectrum of D1 “minus” state minus D4 “minus” state. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 10 mHz.




Figure 17: Double-difference noise spectrum between D1 and D4. Diode signals inversely weighted by single-difference white noise level. Y-axis scale in arbitrary units (proportional to $V/\sqrt{\text{Hz}}$). White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 1 mHz.




Figure 18: Weighted D1 single-difference noise spectrum. Weights adjusted to minimize white noise level and knee frequency. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 1 mHz.

4.6 Weighted Phase State Differencing

By properly weighting the two phase states in the single difference, a substantial reduction in the $1/f$ noise knee frequency can be achieved. Figure 18 shows the noise spectrum for a weighted single-difference using the same data used in Figure 11. The weighting was adjusted to minimize the resulting white noise level — this coincided with minimizing the knee frequency. In this case, the D1 “plus” phase state was weighted by 0.9894 and the “minus” phase state was weighted by 1.0106. The $1/f$ knee frequency has been reduced from about 300 mHz to about 100 mHz compared to the unweighted spectrum. A similar effect was seen on the D4 diode.

When combined into a double-difference, the final knee frequency is also reduced. Figure 19 shows that the knee frequency has dropped from about 100 mHz in Figure 17 to about 50 mHz.




Figure 19: Double-difference noise spectrum between D1 and D4 using weighted single-differences. White noise level indicated by red (lower) horizontal line. Blue (upper) horizontal line indicates $\sqrt{2}$ times the white noise level. Spectrum computed with resolution of 1 mHz.

Module Noise Performance Draft Report

W-Band Module 9 Part 2

Revision 0.1

Joey Richards* and Mike Seiffert†

March 15, 2006

1 Introduction

This is an addendum to the previous W-Band Module 9 draft report. These documents remain a work in progress and will ultimately be reduced to a single report. For now, however, the new data are presented separately here.

2 Data Acquisition System

Since the previous report, we have made two changes to the hardware configuration. First, we have bypassed the optoisolators on the phase switch bias card. Figure 1 shows the alignment of the blanking period with the phase switch transition. Second, we have connected the 3.3 V power supply to the MMIC bias card to the same supply as the digital sections of the bias cards and increased that voltage to 3.6 V. This is a temporary measure that was necessary because of a failing power supply. This deviation from the final power supply configuration should not affect the tests described here.

We are in the process of characterizing the data acquisition system. At present, we are working to confirm the calibration used to generate absolute amplitude measurements. Measurements appear reasonable in order of magnitude, but there are some outstanding questions of $\sqrt{2}$ factors that must be resolved. This does not affect knee frequency measurements, but does affect the reception bandwidth calculation.

3 Bare Module

3.1 Corrected Blanking Period

Figure 2 through Figure 4 show the effect of correcting the blanking period. It appears that there is very little change in the undifferenced data, but both white noise and $1/f$ noise are reduced slightly in the single difference data.

3.2 Spectra and Knee Frequencies

Figure 5 through Figure 11 show spectra computed from a 19-hour run with a nominally-biased room temperature module with nothing connected to its RF input ports. The data acquisition system was used in “34x Compression” mode to reduce the size of the data set. In this mode, 34 consecutive samples are averaged and represented by a single point in the data set (each phase state is averaged separately). The resulting Nyquist frequency is just over 60 Hz.

**joey@caltech.edu*

†*Michael.D.Seiffert@jpl.nasa.gov*




Figure 1: This scope trace shows the alignment of the blanking period with the actual phase switch transition. The magenta (top) trace is the phase switch clock, the cyan (middle) trace is the blanking pulse, and the blue (bottom) trace is the voltage on the phase switch diode. The integrator is held in reset when the blanking pulse is high. The phase switch transition is clearly within the blanking period.




Figure 2: This plot compares the D1 “plus” phase state spectrum before and after the blanking fix. The black (noisier) plot is the “before” and the red (smoother) plot is the “after.” The reduced variance on the spectrum is because the “after” data set is substantially longer.




Figure 3: This plot shows the same “before” spectrum as in Figure 2 (black plot) along with three spectra computed using subsets of the same “after” data set (red, green, and blue plots). All data sets were the same length and were smoothed by a 34-point boxcar integration and downsampled by a factor of 34. This plot illustrates the similarity of the “before” and “after” undifferenced data.




Figure 4: This plot shows the unweighted single-difference data from D1 for the same “before” and “after” data. The white noise level is about 5% lower for the “after” (red, smooth) spectrum and it appears that the $1/f$ noise level is somewhat lower as well.




Figure 5: Undifferenced D1 "plus" state spectrum.

To reduce the size of this document, the spectra have generally been plotted over a limited frequency range. No features have been hidden by this. In most spectra, a horizontal red bar indicates the white noise level and a horizontal blue bar indicates $\sqrt{2}$ times the white noise level, the amplitude at the $1/f$ knee frequency.




Figure 6: D1 unweighted single difference spectrum. $1/f$ knee at about 180 mHz.




Figure 7: D1 weighted single difference spectrum. $1/f$ knee at about 150 mHz. Weighted by inverse standard deviation.




Figure 8: D4 unweighted single difference spectrum. $1/f$ knee at about 220 mHz.




Figure 9: D4 weighted single difference spectrum. $1/f$ knee at about 150 mHz. Weighted by inverse standard deviation.




Figure 10: Unweighted double difference spectrum. $1/f$ knee at about 140 mHz.




Figure 11: Doubly-weighted double difference spectrum. $1/f$ knee at about 80 mHz. This is a weighted difference of weighted single differences. All weights were inverse standard deviation.




Figure 12: These are the spectra of undifferenced data plotted separately for each diode and phase switch state. Black and red plots (virtually identical, on bottom) are D1 “plus” and “minus,” respectively. Blue and green plots (also virtually identical, on top) are D4, “plus” and “minus.” The module was connected to an OMT and feedhorn.

4 Module with OMT

We connected an OMT and a feedhorn to the module and collected data sets. For these tests, the horn was aimed at a room-temperature (290K) Eccosorb target.

4.1 The “Bulge”

Many of the spectra, both with and without the OMT, show an unexpected “bulge” in the $1/f$ -dominated region. This is generally suppressed in unweighted differences and often *not* suppressed in weighted differences. Furthermore, there is no bulge apparent in the doubly-weighted double difference spectrum Figure 18, which is the weighted difference between Figure 14 and Figure 16. Each of these spectra clearly shows the bulge. Thus, it appears that this effect is common to both phase states and both diodes. More investigation is required.




Figure 13: D1 unweighted single difference spectrum. $1/f$ knee at about 180 mHz. The module was connected to an OMT and feedhorn.




Figure 14: D1 weighted single difference spectrum. $1/f$ knee at about 260 mHz, which is higher than the unweighted knee. Weight used here was the inverse standard deviation of the spectrum data; other weights were tried, but all increased the knee frequency relative to the unweighted single difference knee. The module was connected to an OMT and feedhorn.




Figure 15: D4 unweighted single difference spectrum. $1/f$ knee at about 290 mHz. The module was connected to an OMT and feedhorn.




Figure 16: D4 weighted single difference spectrum. $1/f$ knee at about 240 mHz. Weighted by inverse standard deviation. The module was connected to an OMT and feedhorn.




Figure 17: Unweighted double-difference spectrum. $1/f$ knee at about 230 mHz. The module was connected to an OMT and feedhorn.




Figure 18: Doubly-weighted double-difference spectrum. $1/f$ knee at about 100 mHz. This is a weighted difference of weighted single differences. All weights were inverse standard deviation. The module was connected to an OMT and feedhorn. Also, it is not plotted, but a singly-weighted double difference was constructed from unweighted single difference data. It appears that this can reduce the $1/f$ knee frequency to between 100 and 150 mHz.

Effect of Bias on $1/f$ Gain Fluctuations in a QUIET MMIC

Joey Richards* and Mike Seiffert

April 8, 2008

Abstract

This memo describes preliminary results from measurements of the $1/f$ gain fluctuations in a single QUIET MMIC. By using a method similar to that used by Jarosik [1], we minimize the $1/f$ and white noise contributions of our test set. Our preliminary results show a correlation between the level of $1/f$ fluctuations at 1 Hz and the MMIC's drain bias voltage. These measurements were performed using MMIC W-26.

1 Test Apparatus

Our test apparatus consists of two nominally identical back-end MMIC amplifier chains, each connected to a wide-band detector diode. The configuration is shown in Figure 1. The device under test (DUT) is connected to one input a magic-tee. The other magic-tee input is terminated. Its outputs each pass through an adjustable attenuator and into one of the back-end amplifier chains.

The back-end amplifier chains, labeled A and B, each consist of two MMIC amplifier modules. MMIC A1 (B1) is the input to the chain and is followed by a band-pass filter, an isolator, a high-pass filter, then MMIC A2 (B2). The second MMIC's output passes through an isolator and into the wide-band detector diode.

The outputs of the detector diodes are amplified by a low-noise preamp with a voltage gain of 100 and an input impedance of approximately $1\text{ k}\Omega$.¹ Each preamp output connects to an integrator circuit and is further amplified, then integrated and sampled by a National Instruments ADC card in the test PC. The total low-frequency voltage gain from the preamp output to the ADC input is approximately 110. Unless otherwise specified, in our tests reported here, back-end chain A connects to ADC channel 1 and chain B to channel 0.

Before connecting a DUT to the system, we measured the background noise floor with the inputs terminated. Figure 2 shows a typical spectrum. The white noise floor is due to the test set—thus far, we are unable to measure the noise from a warm load above the noise floor of our data set. There are significant $1/f$ fluctuations apparent in the correlated (red) spectrum. We believe these correlations are due to noise from the integrator circuits, but more investigation is necessary. For tests with a MMIC, these fluctuations are low enough to make acceptable measurements, however.

Our DUT is a singly-packaged QUIET MMIC, serial number W-26. It is biased using a separate, DC-isolated bias channel in the same bias box providing bias to the back-end MMICs. This box provides both fixed gate and drain voltages, unlike the QUIET MMIC bias board which provides fixed gate voltage and fixed drain current. The attenuators on both chains were adjusted to produce 1 mV on each detector diode with a nominal bias of the DUT. Biases for the DUT were chosen to vary the drain voltage while keeping the DC level on each detector diode close to 1 mV, without adjusting the attenuators.

*joey@caltech.edu

¹The relatively low input impedance significantly loads the detector diode output, but this does not appear to be a limiting factor for our measurements so far.

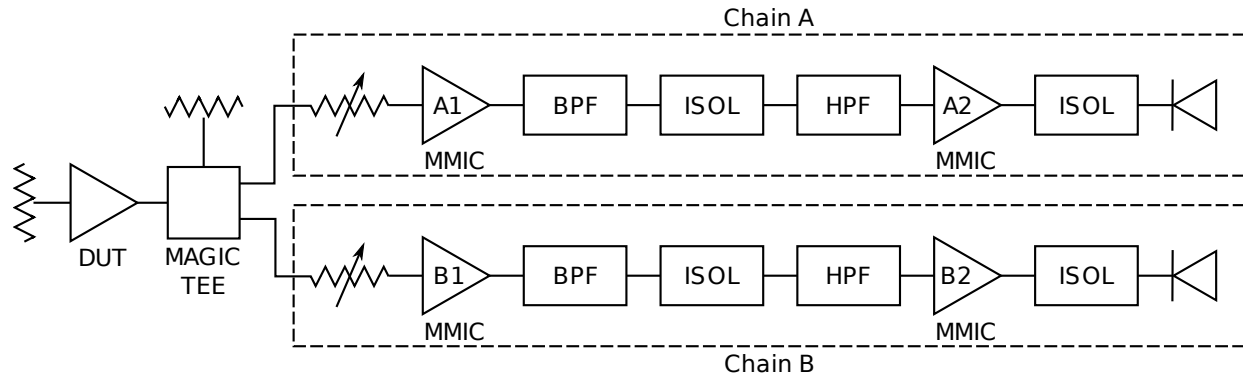


Figure 1: Schematic of the RF portion of the test apparatus. All components are at room temperature. Resistor symbols indicate terminators and potentiometers represent variable attenuators. “BPF” and “HPF” indicate band-pass and high-pass filter, respectively. “ISOL” indicates an isolator. The output of each detector diode is amplified and integrated/sampled by a base-band test set that is not shown.




Figure 2: Typical spectrum of test set noise. The DUT input to the magic-tee is terminated. The attenuator on each chain input is set to slight attenuation, but the noise spectra do not depend on this attenuation. The blue and green data are ADC channels 0 and 1, respectively. The red data are the absolute value of the real cross-spectrum. In this test, chain A is connected to ADC channel 0 and chain B is connected to channel 1.

2 Data Analysis

2.1 Cross-Spectra

To reduce the noise from the back-end amplifiers and electronics, we compute the cross-spectrum of the two channels. An ordinary power spectrum is the Fourier transform of the auto-correlation series of a single input channel. The cross-spectrum is the generalization to two different input series: the Fourier transform of the cross-correlation series.

Unlike a power spectrum, the cross-spectrum is not a real-valued spectrum. The power spectrum is real because the auto-correlation series of a real series is a symmetric, real series, so its Fourier transform is also symmetric and real. The cross-correlation series, however, is merely real, not symmetric. As a result, the cross-spectrum has a symmetric real part and an anti-symmetric imaginary part.

We eliminate the imaginary part of the cross-spectrum by folding the two-sided cross-spectrum into a one-sided spectrum. When we add each negative-frequency component to its positive-frequency counterpart, the imaginary parts cancel (due to anti-symmetry) and we are left with a purely real one-sided spectrum.

This one-sided spectrum is not necessarily positive, since the mean value of the product of two series can be positive or negative. If there is a large correlated component between the two input series, it will dominate and the result will be positive. However, if uncorrelated noise is at a similar level to the correlated signal, the result will occasionally dip below zero. For our purposes, we simply plot the absolute value of the one-sided cross-spectrum—this will not result in significant distortions as long as we work with a signal that dominates the uncorrelated noise.

2.2 Estimating Gain Fluctuations

The cross-spectrum computed as discussed above is an unbiased estimate of the power spectrum of the correlated noise common to the two input chains. Because the correlated signal due to the DUT dominates the background correlated signal, the cross-spectrum of the two chains as an estimate of the power spectrum of the DUT.

We use the spectral density at 1 Hz as a fiducial measure of the gain fluctuations in the DUT. When divided by the DC detector value, this yields the gain fluctuations in Hz^{-1} . The background noise in our test set with the DUT powered off produces a DC level approximately 15% the DC level when the DUT is powered on. This background DC level is subtracted from the value used to divide the spectral density.

The variable attenuators were adjusted to produce equal 1.025 mV responses at both detector diodes at a nominal bias in the middle of the range tested. The attenuators were not adjusted, but as we varied the DUT bias, the two detector diodes did not always agree. At the maximum, they differed by about 0.13 mV. We used the geometric mean of the two detector diodes to normalize the gain fluctuations. Other methods of averaging differed only slightly and did not affect the qualitative features of the plots.

For each bias setting, a run of at least 45 minutes was collected. The longest run was a 20-hour overnight run. Comparisons of long runs with shorter runs at the same DUT bias did not show a significant effect on the estimated noise level. Each data run was divided into 10 second sections and the cross-spectra from these separate sections were averaged to produce a low-noise spectrum as shown in Figure 3.

From each averaged spectrum, the spectral density at 1 Hz was estimated both by eye and by fitting a $1/f^\alpha$ to the region between 0.5 Hz and 5 Hz. These methods produced consistent results. Error bars were crudely estimated by examining the “fur” on the spectrum. The error estimates were examined both on log-log and semi-log scales to avoid distortion by the logarithmic Y-axis scale. A more careful error analysis is intended in the future.

3 Results

See figures that follow. Only plots as a function of drain voltage and drain current bias values are included. Plots against other bias values (gate voltage and gate current) are qualitatively similar to the drain current plots here.




Figure 3: Typical spectrum from a 5000-second (1.4-hour) data set. The blue and green data are detector channels 0 and 1, respectively. The red data are the real cross-spectrum of the two channels. The large peak is at 60 Hz.




Figure 4: Gain fluctuations as a function of MMIC drain voltage. Note that the y-axis label units are incorrect. Y-axis units are arbitrary, but proportional to Hz^{-1} .




Figure 5: Gain fluctuations as a function of MMIC drain current. Note that the y-axis label units are incorrect. Y-axis units are arbitrary, but proportional to Hz^{-1} .




Figure 6: Gain fluctuations as a function of DC detected voltage. Here, the y-axis units are correct. The DC detector voltage is the geometric mean of the DC detected voltage at the two detector diodes. The difference between the two was less than about 0.1 mV, so the averaging method has little effect.




Figure 7: Spectral density at 1 Hz as a function of MMIC drain voltage. These are the same data as plotted in Figure 4 before dividing by the detected DC voltage.




Figure 8: Spectral density at 1 Hz as a function of MMIC drain current. These are the same data as plotted in Figure 5 before dividing by the detected DC voltage.

4 Discussion / Future Plans

These results indicate a trend of lower $1/f$ gain fluctuations as the drain voltage of the DUT is decreased. It does not appear that there are systematic correlations with other bias variables. Several avenues of future measurements are immediately clear. These results come from tests with a single sample of a MMIC. A second MMIC should be tested similarly for comparison. Tests of these MMICs at cryogenic temperatures should be performed—these results are only interesting if they hold true at cryogenic temperatures as well as at room temperature.

Assuming this effect is real, it also remains to be verified that the reduction in $1/f$ gain fluctuations on a single MMIC tested in this manner translates into an effective $1/f$ noise reduction in a QUIET module in actual operation. This depends on the cause of residual $1/f$ noise after demodulation and will shed light on the origins of this residual $1/f$ whether or not the reduction is realized.

References

- [1] N. C. Jarosik, “Measurements of the low-frequency-gain fluctuations of a 30-GHz high-electron-mobility-transistor cryogenic amplifier,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 44, pp. 193–197, Feb. 1996.

1/f Fluctuation Studies: Uncorrelated MMIC Tests

Joey Richards* and Mike Seiffert

September 26, 2011

1 Apparatus

For these measurements we completely isolated the two legs of our test set, as shown in Figure 1. The back-end biases were set to approximately the same values as in our previously reported measurements. The MMIC devices under test (DUTs) were biased as follows.

	DUT A		DUT B	
	V	I	V	I
Gate A	0.106 V	30 μ A	0.044 V	-8 μ A
Gate B	0.102 V	30 μ A	0.052 V	-5 μ A
Drain	0.970 V	10.1 mA	0.739 V	9.9 mA

2 Spectra

In this document, we plot spectra on semilog or linear plots rather than log-log plots as we have done in the past. This is more appropriate for cross-spectra because it does allow both positive and negative values to be plotted. In particular, when a cross-spectrum is near zero-mean, plotting the absolute value (as is necessary for a log-log plot) tends to make its mean value appear significantly larger.

*joey@caltech.edu

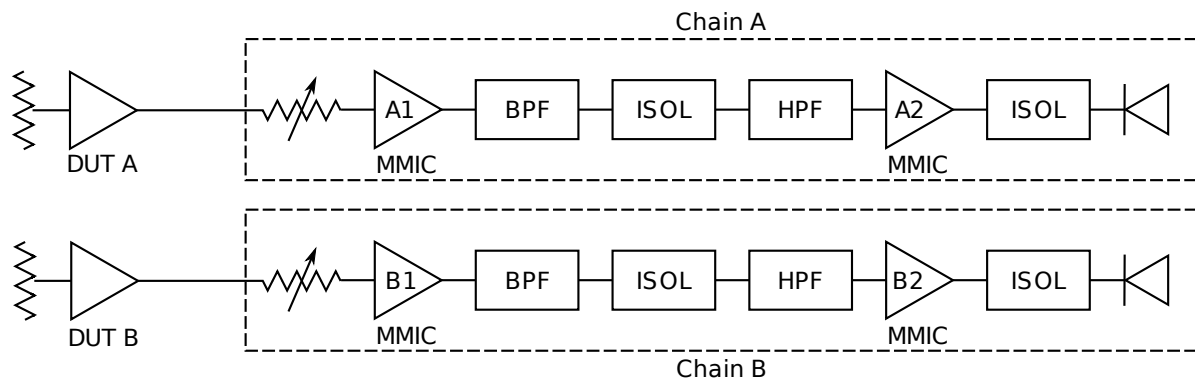


Figure 1: RF test set schematic for the uncorrelated MMIC tests. MMIC A is serial number L401. MMIC B is serial number W-26 and was used in the previous single-DUT tests. All components are room temperature.




Figure 2: Simulated cross-spectrum of two independent $1/f$ -noise series.




Figure 3: Power spectra of A and B signals (blue and green) and the cross-spectrum (red) using the fully independent two-DUT test set shown in Figure 1. Due to a systematic processing error, data above 60 Hz is invalid.




Figure 4: High-frequency spectra of A and B signals (blue and green) and the cross-spectrum (red). The RMS spectral densities over these data are $24.1 \text{ nV}/\sqrt{\text{Hz}}$, $22.5 \text{ nV}/\sqrt{\text{Hz}}$, and $2.3 \text{ nV}/\sqrt{\text{Hz}}$ for blue, green, and red, respectively.

3 Two Notes

3.1 Statistical Note

The cross-spectrum in Figure 4 appears to have a near-zero mean, but it appears to the eye that its values are clustered around non-zero positive and negative values with relatively few points very close to zero. A histogram of these data is shown in Figure 5, clearly demonstrating the absence of points near zero.

The voltage spectral density plot is calculated as the point-wise square root of the power spectral density. Because the cross-spectrum can be positive or negative, we take the absolute value of the power spectral density before the square root, then multiply by its sign to produce a real-valued voltage spectral density. Figure 6 shows the power spectral density before this processing. There is clearly no absence of points near the mean of the distribution.

The absence of points near zero in the voltage spectral density plot is an artifact of the square-rooting process we use. Let x be a Gaussian random variable with zero mean and unit standard deviation. The probability density for x is

$$p_x(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right). \quad (1)$$

Now consider $y = \text{sign}(x)\sqrt{|x|}$, the same transformation we use to produce our voltage spectral density plot. To find its probability density, we just need the derivative with respect to x . We'll ignore the discontinuity in the derivative at $x = y = 0$ that comes from the sign function, as it is inconsequential for our




Figure 5: Histogram of the cross-spectrum data from Figure 4. Clearly, there are two separate clusters of values, one positive and one negative.




Figure 6:




Figure 7: Probability density function of $y = \text{sign}(x)\sqrt{|x|}$ for a zero-mean, unity-variance Gaussian random variable x . This qualitatively matches the features of the histogram of the voltage spectral density plot in Figure 5.

purposes.

$$p_y(y)dy = p_x(x)dx , \quad (2)$$

$$p_y(y) = p_x(x) \left(\frac{dx}{dy} \right) , \quad (3)$$

$$p_y(y) = p_x(x) (2|y|) , \quad (4)$$

$$p_y(y) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) (2|y|) , \quad (5)$$

$$p_y(y) = \sqrt{\frac{2}{\pi}} |y| \exp\left(-\frac{y^4}{2}\right) . \quad (6)$$

This probability distribution is plotted in Figure 7. It features a null at zero and is qualitatively quite similar to our observed histogram. The imbalance between the positive and negative lobes is not understood, but probably reflects correlated noise or spurious interference of some sort.

3.2 Knee Frequency Note

We use the 1 Hz gain fluctuation level rather than the knee frequency as a parameter because it is independent of the radiometer bandwidth. The conversion between these parameters is straightforward, but depends on the type of radiometer in use. Following Wollack and Pospieszalski (1998), for our test set, the total fluctuations at frequency f are given by

$$\frac{\delta V}{V_{dc}} = \sqrt{\frac{2}{\beta} + \delta g^2(f)} . \quad (7)$$

Here, $\delta g^2(f) = \delta g_0^2 \times (1 \text{ Hz}/f)$ is the $1/f$ gain fluctuation in 1/Hz and β is the RF bandwidth. The first term gives the white noise floor. Solving this to find the knee frequency yields

$$f_{knee} = \frac{\beta \cdot \delta g_0^2}{2} . \quad (8)$$

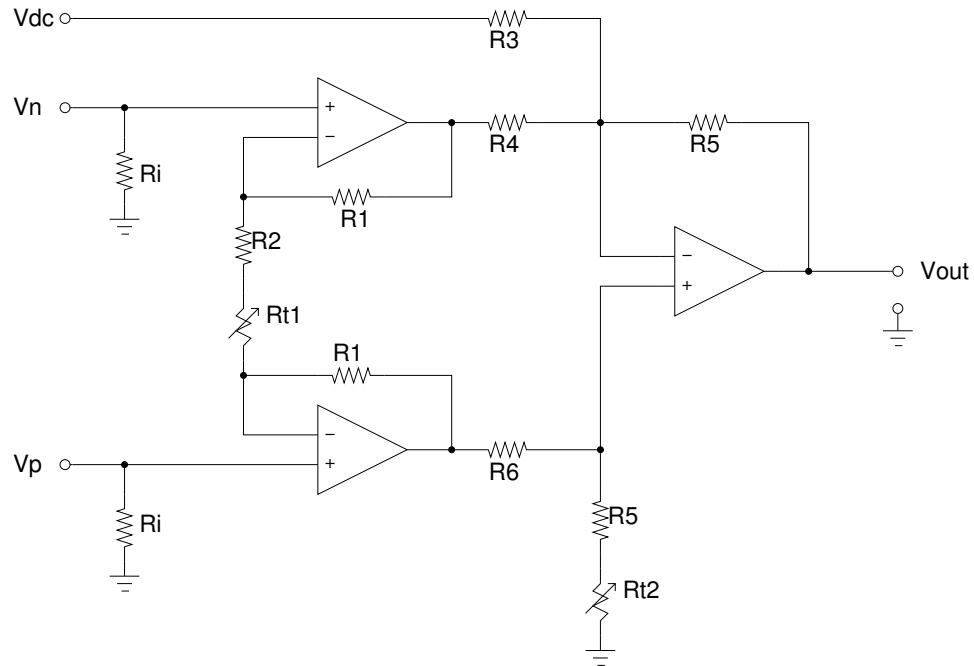


Figure D.66. Full preamp schematic. Note that resistor value R_6 should be $R_4 + \frac{R_4 R_5}{R_3}$. When this is the case, the output of the amplifier will be $V_{out} = \frac{R_5}{R_4} \left(1 + 2 \frac{R_1}{R_2}\right) (V_p - V_n) - \frac{R_5}{R_3} V_{DC}$. The trim pots R_{t1} and R_{t2} tune the gain and common-mode rejection ratio (CMRR), respectively.

D.5 Low-Noise Preamplifier Design

In this section, my design for a low-noise, high-input-impedance differential preamplifier with an independent subtraction input is presented. Results from testing of a prototype of the circuit are briefly described as well. A four-channel version was constructed on a custom printed circuit board and used extensively for radiometer characterization. This circuit was designed as an option for use in the QUIET back-end, and a scheme for connecting this amplifier to a biased RF detector diode is presented and analyzed. Ultimately a different design was selected for deployment, although this work influenced the final design.

The circuit shown in figure D.66 is a standard three-op amp instrumentation amplifier with the addition of a subtraction input in the second stage. It can be configured to have a wide bandwidth (> 1 MHz), low input-referenced (RTI) noise (~ 6 nV/ $\sqrt{\text{Hz}}$), and an input impedance limited only by the input impedance of the op amps (> 1 M Ω).

D.5.1 Analysis

The standard instrumentation amplifier is built from two pieces: an input stage consisting of a pair of non-inverting amplifiers connected to a differential amplifier output stage. The design described here includes an additional subtraction node on the differential amplifier to allow a DC offset voltage to be removed from the input in order to avoid dynamic range problems due to detector diode biasing.




Figure D.67. Schematic of input amplifier stage. Note that the R_{t1} trim pot has been included in R_2 for analysis.

D.5.1.1 Input Stage

Figure D.67 shows the input stage of the preamp. The input impedance of the amplifier is $2R_i$. These shunt resistors are required to provide DC bias paths for the op amp inputs. Because these resistors are in parallel with the source, they will not contribute to the noise floor of the amplifier as long as the source impedance $R_s < R_i$. Furthermore, as long as $R_s \ll R_i$, the impedance of the source will not affect the effective gain of the amplifier.

The outputs of this stage, V_{op} and V_{on} , are related to the inputs as follows.

$$V_{on} = \left(1 + \frac{R_1}{R_2}\right) V_n - \frac{R_1}{R_2} V_p, \quad (\text{D.1})$$

$$V_{op} = \left(1 + \frac{R_1}{R_2}\right) V_p - \frac{R_1}{R_2} V_n. \quad (\text{D.2})$$

This is equivalent to a differential gain of

$$V_{op} - V_{on} \equiv G_{d1} (V_p - V_n) = \left(1 + 2\frac{R_1}{R_2}\right) (V_p - V_n). \quad (\text{D.3})$$

D.5.1.2 Differential Stage

Figure D.68 shows the differential / DC offset subtraction stage of the preamp. First, we will assume V_{on} and V_{DC} are grounded, and calculate the output response due to V_{op} alone. Since the current into the op amp




Figure D.68. Schematic of differential amplifier stage. Note that the Rt2 trim pot has been included in R7 for analysis. R7 was labeled R5 in figure D.66. In our analysis, we will show that R7 should have the same value as R5 to achieve balanced gain.

inputs is negligible, the voltage at the noninverting op amp input is

$$V_+ = V_{op} \times \frac{R_7}{R_6 + R_7}, \quad (\text{D.4})$$

where R_7 in this equation equals the trimmed value $R_7 + R_{t7}$ on the schematic. Assuming that the feedback is working, this will be equal to the voltage at the inverting input, which is

$$V_- = V_{out} \times \frac{R_3 \parallel R_4}{(R_3 \parallel R_4) + R_5}. \quad (\text{D.5})$$

Solving for the output voltage, this yields

$$V_{out} = \left(\frac{(R_3 \parallel R_4) + R_5}{R_6 + R_7} \right) \left(\frac{R_7}{R_3 \parallel R_4} \right) V_{op}. \quad (\text{D.6})$$

Next, we will calculate the response due to V_{on} alone. The noninverting input is now at ground, so the inverting input will be a virtual ground. Using Kirchhoff's current law at the inverting input node, we have

$$\frac{V_{on}}{R_4} + \frac{V_{out}}{R_5} = 0, \quad (\text{D.7})$$

so

$$V_{out} = -\frac{R_5}{R_4} V_{on}. \quad (\text{D.8})$$

Likewise, the response due to V_{DC} is

$$V_{out} = -\frac{R_5}{R_3}V_{DC}. \quad (D.9)$$

The combined output is then

$$V_{out} = \left(\frac{(R_3 \parallel R_4) + R_5}{R_6 + R_7} \right) \left(\frac{R_7}{R_3 \parallel R_4} \right) V_{op} - \frac{R_5}{R_4}V_{on} - \frac{R_5}{R_3}V_{DC}. \quad (D.10)$$

If we set $R_6 = R_4 + \frac{R_4 R_5}{R_3}$ and $R_7 = R_5$, then this becomes

$$V_{out} = \frac{R_5}{R_4} (V_{op} - V_{on}) - \frac{R_5}{R_3}V_{DC} \equiv G_{d2} (V_{op} - V_{on}) - G_{DC}V_{DC}. \quad (D.11)$$

D.5.1.3 Full Amplifier

Combining the results of the preceding sections, we find that the total response of the preamplifier is as follows.

$$V_{out} = G_{d2}G_{d1} (V_p - V_n) - G_{DC}V_{DC}, \quad (D.12)$$

$$V_{out} = \frac{R_5}{R_4} \left(1 + 2\frac{R_1}{R_2} \right) (V_p - V_n) - \frac{R_5}{R_3}V_{DC}. \quad (D.13)$$

D.5.2 Prototype Circuit

A prototype circuit has been assembled using Analog Devices OP37 op amps and the following component values. All resistors have 1% tolerance.

Component	Value
Ri	100 k Ω
R1	4.87 k Ω
R2	1.00 k Ω
R3	1.00 k Ω
R4	1.00 k Ω
R5	10.0 k Ω
R6	11.0 k Ω
Rt1	200 Ω single-turn potentiometer
Rt2	200 Ω single-turn potentiometer

This prototype was tuned using Rt1 to have a total DC gain of 100, divided equally between the two stages. The DC offset gain, G_{DC} is 10 in this prototype. Rt2 can be used to tune the common mode rejection ratio of the preamp, although this has not been tested.

The bandwidth of the differential inputs exceeds 1 MHz. The gain appears to be flat from DC to about 400 kHz. Starting at 400 kHz, there is a rise in the gain, peaking at a gain of about 150 at 1 MHz, then

dropping back to a gain of 100 at about 1.5 MHz and decreasing steadily above 1.5 MHz. It should be possible to eliminate this gain peak by adding some capacitors to the circuit.

A version of the prototype without the DC offset input was measured to have a noise floor of about $6 \text{ nV}/\sqrt{\text{Hz}}$. The inclusion of the DC offset input should not affect the noise floor significantly. The $1/f$ knee frequency appears to occur at a few hertz on the prototype circuit.

D.5.3 Detector Diode Biasing

A scheme using a single digital-to-analog converter (DAC) to provide both the diode bias current and a DC level for removal is shown in figure D.69. In this circuit, a DAC is operated with an output range $10\times$ the desired bias voltage range for V_{bias} , then divided down and low-pass filtered, and buffered to minimize the DAC noise contribution. Although a simple RC low-pass filter is shown in the schematic, a more aggressive filter could be integrated in the buffer amp if necessary.




Figure D.69. One possible scheme for biasing the detector diode and generating a DC offset signal using a single DAC channel. The amplifier block in the lower-right represents the preamp described in this note.

This buffered voltage, V_{bias} , biases the detector diode D1 through the R_b resistors. These resistors must be chosen to allow sufficient bias current with the available V_{bias} range. Additionally, these resistors will determine the effective input impedance of the preamp circuit, so making these as large as possible will maximize transfer of the diode signal into the preamp.

The V_{bias} signal is also used to generate the DC offset subtraction signal for the preamp. In the circuit shown, it is amplified by a noninverting amplifier with gain $G = 1 + \frac{R_1}{R_2}$. These resistors should be set such

that

$$G = \frac{G_{d1}G_{d2}}{G_{DC}} \times \frac{R_{D1}}{R_{D1} + 2R_b}, \quad (\text{D.14})$$

where G_{DC} is the DC offset gain of the preamp, G_{d1} and G_{d2} are the differential input gains in the first and second stages of the preamp, and R_{D1} is the DC resistance of the detector diode at its bias point. Error in the gain matching, caused by, e.g., an incorrect estimate of R_{D1} , will leave a residual DC signal contribution in V_{out} .

Because the DC offset gain of the preamp can be independently configured, it is possible to eliminate the noninverting amplifier and inject V_{bias} directly into the V_{DC} input of the preamp. In this case, G_{DC} can be set to compensate for the voltage drops in the bias resistors. This would allow the entire bias circuit and preamp to be constructed from a single quad op amp package.

D.5.4 Conclusions

This circuit is a promising candidate for a preamplifier. It exhibits a large input impedance, eliminating concerns that variations in detector diode impedance will affect the gain of the preamp. The noise floor, while slightly higher than the previous circuit ($6 \text{ nV}/\sqrt{\text{Hz}}$ versus about $5 \text{ nV}/\sqrt{\text{Hz}}$), should be acceptably low. Finally, the inclusion of a DC offset input can eliminate the need for a separate unity-gain DC path while still allowing measurement of the DC level of the diode. Although only an OP37-based circuit has been tested so far, it should be possible to build an equivalent circuit using Linear Technologies LT1125 or LT1127 quad op amps. This would allow the entire preamp to be constructed using a single integrated circuit and a few resistors.

D.6 QUIET 43 GHz Paper

This section contains the preprint manuscript of the QUIET first-results paper, which describes the operation of and results from operation of the Q-band instrument on the Chajnantor plateau in the Atacama Desert in Chile (QUIET Collaboration et al. 2011). The present author's contributions to this work consisted of development, testing, and characterization of the polarimeter modules and support electronics, including the work described in the previous sections of this chapter.

FIRST SEASON QUIET OBSERVATIONS: MEASUREMENTS OF CMB POLARIZATION POWER SPECTRA
 AT 43 GHz IN THE MULTIPOLE RANGE $25 \leq \ell \leq 475$

QUIET COLLABORATION—C. BISCHOFF^{1,22}, A. BRIZIUS^{1,2}, I. BUDER¹, Y. CHINONE^{3,4}, K. CLEARY⁵, R. N. DUMOULIN⁶,
 A. KUSAKA¹, R. MONSALVE⁷, S. K. NÆSS⁸, L. B. NEWBURGH^{6,23}, R. REEVES⁵, K. M. SMITH^{1,23}, I. K. WEHUS⁹,
 J. A. ZUNTZ^{10,11,12}, J. T. L. ZWART⁶, L. BRONFMAN¹³, R. BUSTOS^{7,13,14}, S. E. CHURCH¹⁵, C. DICKINSON¹⁶,
 H. K. ERIKSEN^{8,17}, P. G. FERREIRA¹⁰, T. GAIER¹⁸, J. O. GUNDERSEN⁷, M. HASEGAWA³, M. HAZUMI³,
 K. M. HUFFENBERGER⁷, M. E. JONES¹⁰, P. KANGASLAHTI¹⁸, D. J. KAPNER^{1,24}, C. R. LAWRENCE¹⁸, M. LIMON⁶, J. MAY¹³,
 J. J. MCMAHON¹⁹, A. D. MILLER⁶, H. NGUYEN²⁰, G. W. NIXON²¹, T. J. PEARSON⁵, L. PICCIRILLO¹⁶, S. J. E. RADFORD⁵,
 A. C. S. READHEAD⁵, J. L. RICHARDS⁵, D. SAMTLEBEN^{2,25}, M. SEIFFERT¹⁸, M. C. SHEPHERD⁵, S. T. STAGGS²¹,
 O. TAJIMA^{1,3}, K. L. THOMPSON¹⁵, K. VANDERLINDE^{1,26}, R. WILLIAMSON^{6,27}, B. WINSTEIN¹

Submitted to ApJ—This paper should be cited as “QUIET (2010)”

ABSTRACT

The Q/U Imaging Experiment (QUIET) employs coherent receivers at 43 GHz and 95 GHz, operating on the Chajnantor plateau in the Atacama Desert in Chile, to measure the anisotropy in the polarization of the CMB. QUIET primarily targets the B modes from primordial gravitational waves. The combination of these frequencies gives sensitivity to foreground contributions from diffuse Galactic synchrotron radiation. Between 2008 October and 2010 December, over 10,000 hours of data were collected, first with the 19-element 43-GHz array (3458 hours) and then with the 90-element 95-GHz array. Each array observes the same four fields, selected for low foregrounds, together covering ≈ 1000 square degrees. This paper reports initial results from the 43-GHz receiver which has an array sensitivity to CMB fluctuations of $69 \mu\text{K}/\sqrt{\text{s}}$. The data were extensively studied with a large suite of null tests before the power spectra, determined with two independent pipelines, were examined. Analysis choices, including data selection, were modified until the null tests passed. Cross correlating maps with different telescope pointings is used to eliminate a bias. This paper reports the EE, BB, and EB power spectra in the multipole range $\ell = 25\text{--}475$. With the exception of the lowest multipole bin for one of the fields, where a polarized foreground, consistent with Galactic synchrotron radiation, is detected with $3\text{-}\sigma$ significance, the E-mode spectrum is consistent with the Λ CDM model, confirming the only previous detection of the first acoustic peak. The B-mode spectrum is consistent with zero, leading to a measurement of the tensor-to-scalar ratio of $r = 0.35^{+1.06}_{-0.87}$. The combination of a new time-stream “double-demodulation” technique, Mizuguchi–Dragone optics, natural sky rotation, and frequent boresight rotation leads to the lowest level of systematic contamination in the B-mode power so far reported, below the level of $r = 0.1$.

Subject headings: cosmic background radiation—Cosmology: observations—Gravitational waves—Inflation—Polarization

¹ Kavli Institute for Cosmological Physics, Department of Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA; send correspondence to A. Kusaka, akito@kicp.uchicago.edu

² Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

³ High Energy Accelerator Research Organization (KEK), 1-1 Oho, Tsukuba, Ibaraki 305-0801, Japan

⁴ Astronomical Institute, Graduate School of Science, Tohoku University, Aramaki, Aoba, Sendai 980-8578, Japan

⁵ Cahill Center for Astronomy and Astrophysics, California Institute of Technology, 1200 E. California Blvd M/C 249-17, Pasadena, CA 91125, USA

⁶ Department of Physics and Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

⁷ Department of Physics, University of Miami, 1320 Campo Sano Drive, Coral Gables, FL 33146, USA

⁸ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

⁹ Department of Physics, University of Oslo, P.O. Box 1048 Blindern, N-0316 Oslo, Norway

¹⁰ Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

¹¹ Oxford Martin School, 34 Broad Street, Oxford OX1 3BD, UK

¹² Department of Physics and Astronomy, University College London, Gower Street, London WC1E, UK

¹³ Departamento de Astronomía, Universidad de Chile, Casilla 36-D, Santiago, Chile

¹⁴ Departamento de Astronomía, Universidad de Concepción, Casilla 160-C, Concepción, Chile

¹⁵ Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Varian Physics Building, 382 Via Pueblo Mall, Stanford, CA 94305, USA

¹⁶ Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester M13 9PL, UK

¹⁷ Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway

¹⁸ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA 91109

¹⁹ Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA

²⁰ Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

²¹ Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

²² Current address: Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 43, Cambridge, MA 02138, USA

²³ Current address: Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

²⁴ Current address: Micro Encoder Inc., Kirkland, WA 98034, USA

1. INTRODUCTION

The inflationary paradigm resolves several outstanding issues in cosmology, including the flatness, horizon, and monopole problems, and it provides a compelling explanation for the origin of structure in the Universe (e.g. Liddle & Lyth 2000, and references therein). So far all cosmological data, including measurements of Cosmic Microwave Background (CMB) anisotropies, support this paradigm; still the underlying fundamental physics responsible for inflation is unknown. Inflation produces a stochastic background of gravity waves that induce odd-parity tensor “B modes” at large angular scales in the CMB polarization. If these primordial B modes, parametrized by the tensor-to-scalar ratio r , are detected, one can learn about the energy scale of inflation. In many attractive slow-roll models, this scale is given approximately by $r^{1/4} \times 10^{16}$ GeV. For large-field models, the energy scale is near the Grand Unification Scale in particle physics, so that $r \gtrsim 0.01$. A new generation of experiments aims for good sensitivity in this range of r . Establishing the existence of primordial B modes would both verify an important prediction of inflation and provide access to physics at an incredibly high energy scale.

The most stringent limit to date is $r < 0.20$ at the 95% confidence level (Komatsu et al. 2010) set by a combination of CMB–temperature-anisotropy measurements, baryon acoustic oscillations, and supernova observations, but cosmic variance prohibits improvements using only these measurements.

E-mode polarization has now been detected by many experiments (e.g., Kovac et al. 2002; Leitch et al. 2005; Montroy et al. 2006; Sievers et al. 2007; Wu et al. 2007; Bischoff et al. 2008; Larson et al. 2010). These measurements are consistent with predictions from CMB–temperature-anisotropy measurements, and they provide new information on the epoch of reionization. Only BICEP has accurately measured E-mode polarization in the region of the first acoustic peak (Chiang et al. 2010); that paper also reports the best limit on r coming from cosmological B modes: $r < 0.72$ at the 95% confidence level.

Experiments measuring B-mode polarization in the CMB should yield the best information on r , but this technique is still in its infancy. B modes are expected to be at least an order of magnitude smaller than the E modes so control of systematic errors and foregrounds will be particularly critical. Below ≈ 90 GHz, the dominant foreground comes from Galactic synchrotron emission, while at higher frequencies, emission from thermal dust dominates. Most planned or operating CMB polarization experiments employ bolometric detectors observing most comfortably at frequencies $\gtrsim 90$ GHz, so they cannot estimate synchrotron contamination from their own data.

²⁵ Current address: Nikhef, Science Park, Amsterdam, The Netherlands

²⁶ Current address: Department of Physics, McGill University, 3600 Rue University, Montreal, Quebec H3A 2T8, Canada

²⁷ Current address: Kavli Institute for Cosmological Physics, Enrico Fermi Institute, The University of Chicago, Chicago, IL 60637, USA

The Q/U Imaging Experiment (QUIET) is one of two CMB polarization experiments to observe at frequencies suitable for addressing synchrotron contamination, making observations at 43 GHz (Q band) and 95 GHz (W band) and with sufficient sensitivity to begin to probe primordial B modes. The other is *Planck* (Tauber et al. 2010).

QUIET uses compact polarization-sensitive modules based upon High–Electron–Mobility Transistor (HEMT) amplifiers, combined with a new time-stream “double-demodulation” technique, Mizuguchi–Dragone (MD) optics (for the first time in a CMB polarization experiment), natural sky rotation, and frequent rotation about the optical axis to achieve a very low level of contamination in the multipole range where a primordial–B-mode signal is expected.

Between 2008 October and 2010 December, QUIET collected over 10,000 hours of data, split between the Q-band and W-band receivers. Here we report first results from the first season of 3458 hours of Q-band observation. After describing the instrument, observations, and detector calibrations (Sections 2, 3, and 4), we discuss our analysis techniques and consistency checks (5 and 6). CMB power spectra are then presented together with a foreground detection (7). We evaluate our systematic errors (8) and then conclude (9).

2. THE INSTRUMENT

The QUIET instrument comprises an array of correlation polarimeters cooled to 20 K and coupled to a dual-reflector telescope, installed on a three-axis mount inside a comoving ground screen. The instrument is illustrated in Figure 1. Further details are given below and in Newburgh et al. (2010), Kusaka et al. (2010), and Buder (2010).




FIG. 1.— Overview of the QUIET instrument. The cryostat and 1.4-m telescope mirrors are enclosed in a rectangular comoving absorbing ground screen; in this figure its walls are transparent. The telescope, cryostat and electronics are mounted on a single platform attached to the deck bearing, which allows rotations around the instrument’s optical axis.

The Q-band QUIET receiver is a 19-element array containing 17 low-noise correlation polarimeters, each simultaneously measuring the Stokes Q, U, and I parameters, and two CMB differential-temperature monitors.

The first element in the QUIET optical chain is a 1.4-m crossed Mizuguchi–Dragone dual-reflective telescope

(Mizugutch et al. 1976; Dragone 1978). The crossed MD configuration is very compact, with low cross polarization and a large diffraction-limited field of view. The telescope is described in detail in Imbriale et al. (2010). Light incident on the mirrors is focused into an array of corrugated circular feed horns (Gundersen & Wollock 2009), yielding a full-width half-maximum (FWHM) beam size of $27'3$ and a roughly circular field of view of 7° diameter. Radiation from each feed horn enters a septum polarizer (Bornemann & Labay 1995) which separates left and right circularly-polarized components (L and R) into two waveguide ports which mate to a QUIET correlation module, detailed below.

The module array and feed horns are cooled to 20 K in a cryostat to reduce instrumental noise. An electronics enclosure mounted next to the cryostat houses the electronics necessary for biasing the modules and recording their data. The cryostat, electronics, and telescope are installed on the former CBI mount (Padin et al. 2002). This mount provides three-axis motion: azimuth, elevation, and rotation about the optical axis. This last is called “deck” rotation.

The cryostat and telescope are enclosed by an absorbing comoving ground screen. The ground screen was designed to have two parts, but the upper section (not shown in Fig. 1) was not installed until after the Q-band instrument was removed. Its absence was correctly anticipated to result in two far sidelobes, which were mapped with a high-power source by the QUIET W-band instrument in the field and measured to be $\lesssim -60$ dB with the QUIET Q-band instrument when the Sun passed through them. The effects of these sidelobes are mitigated through filtering and data selection (Sections 5.1.3 and 5.2). Section 8.4 shows that any residual contamination is small.

Each QUIET Q-band correlation module, in a footprint of only $5.1 \times 5.1 \text{ cm}^2$, receives the circular polarization modes of the celestial radiation and outputs Stokes Q , U and I as follows. Each input is independently amplified and passed through a phase switch. One phase switch alternates the sign of the signal voltage at 4 kHz, while the other switches at 50 Hz. The two signals are combined in a 180° hybrid coupler, with outputs proportional to the sum and difference of the inputs. Since the module inputs are proportional to $(L, R) = (E_x \pm iE_y)/\sqrt{2}$, where E_x and E_y are orthogonal components of the incident electric field, the coupler outputs are amplified versions of E_x and iE_y , with the phase switch reversing their roles. Half of each output is bandpass filtered and rectified by a pair of detector diodes, while the other half passes into a 90° hybrid coupler. A second pair of bandpass filters and detector diodes measures the power from this coupler’s outputs (Kangaslahti et al. 2006).

Synchronous demodulation of the 4-kHz phase switching yields measurements of Stokes $+Q$ and $-Q$ on the first two diodes and Stokes $+U$ and $-U$ on the remaining two. This high-frequency differencing suppresses low-frequency atmospheric fluctuations as well as $1/f$ noise from the amplifiers, detector diodes, bias electronics, and data-acquisition electronics. Subsequent demodulation of the 50-Hz phase switching removes spurious instrumental polarization generated by unequal transmission

coefficients in the phase-switch circuits. The resulting four “double-demodulated” time streams are the polarization channels.

Averaging the output of each diode rather than demodulating it results in a measurement of Stokes I , hereafter called total power, denoted “TP.” The TP time streams are useful for monitoring the weather and the stability of the detector responsivities, but suffer too much contamination from $1/f$ noise to constrain the CMB temperature anisotropy. Therefore, the Q-band instrument includes two correlation modules that are coupled to a pair of neighboring feed horns to measure the temperature difference between them, in a scheme similar to the *WMAP* differencing assemblies (Jarosik et al. 2003). These differential-temperature modules provide calibration data for the telescope pointing, beams, and sidelobes, as well as CMB data. Their feed horns are in the outer ring of the close-packed hexagonal array, $\approx 3^\circ$ from the center.

Here we summarize several array-wide characteristics of the polarimeters. Bandpass measurements in the lab and at the start of the observing season find that the average center frequency is 43.1 ± 0.4 GHz, and the average bandwidth is 7.6 ± 0.5 GHz. We calculate the noise power spectra of the double-demodulated polarimeter time streams from each 40–90-minute observation to assess their $1/f$ knee frequencies and white-noise levels (see Section 5.1). The median $1/f$ knee frequency is 5.5 mHz, well below the telescope scan frequencies of 45–100 mHz.

From the white-noise levels and responsivities (Section 4.1) we find an array sensitivity²⁸ to CMB fluctuations of $69 \mu\text{K}\sqrt{\text{s}}$, such that the mean polarized sensitivity per module is $280 \mu\text{K}\sqrt{\text{s}}$.




FIG. 2.— The CMB and Galactic patches, in equatorial coordinates, superimposed on a Q-band all-sky *WMAP* 7-year temperature map (Jarosik et al. 2010). Note that the Galactic-plane temperature signal saturates the color scale. Patch G-2 is the Galactic center.

3. OBSERVATIONS

QUIET is located on the Chajnantor plateau in the Atacama Desert of northern Chile ($67^\circ 45' 42'' \text{ W}$, $23^\circ 01' 42'' \text{ S}$). A combination of high altitude (5080 m) and extreme dryness results in excellent observing conditions for most of the year. During the eight months of QUIET Q-band observations, the median precipitable

²⁸ This is the sensitivity for 62 polarization channels. Six of 68 polarization channels are non-functional—an array yield of 92%.

water vapor (PWV) measured at the nearby APEX site (Güsten et al. 2006) was 1.2 mm.

We began observations with the Q-band receiver on 2008 October 24, and took 3458 hours of data until 2009 June 13 (when the receiver was replaced on the telescope by the 90-element W-band receiver). Of these data, 77% are for CMB, with 12% of the observing time used for Galactic fields, 7% for calibration sources, and 4% cut due to obvious instrumental problems such as lack of telescope motion. We observe 24 hours a day, except when interrupted. Our full-season operating efficiency is 63%; causes of downtime include occasional snow, power outages, and mechanical failures.

TABLE 1
PATCH LOCATIONS AND INTEGRATION TIMES

Patch	RA (J2000)	Dec. (J2000)	Integration Hours
CMB-1	12 ^h 04 ^m	−39°00′	905
CMB-2	05 ^h 12 ^m	−39°00′	703
CMB-3	00 ^h 48 ^m	−48°00′	837
CMB-4	22 ^h 44 ^m	−36°00′	223
G-1	16 ^h 00 ^m	−53°00′	311
G-2	17 ^h 46 ^m	−28°56′	92

NOTE. — The central equatorial coordinates and integration times for each observing patch. G-1 and G-2 are Galactic patches.

3.1. Field Selection

We observe four CMB fields, referred to henceforth as “patches.” Table 1 lists their center positions and total integration times, while Figure 2 indicates their positions on the sky. The number of patches is determined by the requirement to always have one patch above the lower elevation limit of the mount (43°). The specific positions of each patch were chosen to minimize foreground emission using *WMAP* 3-year data. The area of each patch is $\approx 250 \text{ deg}^2$. In addition to the four CMB patches, we observe two Galactic patches. These allow us to constrain the spectral properties of the polarized low-frequency foregrounds with a high signal-to-noise ratio. The results from the Galactic observations will be presented in a future publication.

3.2. Observing Strategy

Scanning the telescope modulates the signal from the sky, converting CMB angular scales into frequencies in the polarimeter time streams. Since QUIET targets large angular scales, fast scanning ($\approx 5^\circ \text{ s}^{-1}$ in azimuth) is critical to ensuring that the polarization modes of interest appear at higher frequencies than the atmospheric and instrumental $1/f$ knee frequencies.

So that each module sees a roughly-constant atmospheric signal, each QUIET scan is a constant-elevation scan (CES): periodic motion solely in azimuth with both the elevation and deck-rotation axes fixed. Each CES has an amplitude of 7.5° on the sky, with period 10–22 s. Typical CESes last 40–90 minutes. We replot the telescope when the patch center has moved by 15° in order to build up data over an area of $\approx 15^\circ \times 15^\circ$ for each patch. Note that a central region $\approx 8^\circ$ across is observed by all polarimeters since the instrument’s field of view has a diameter of $\approx 7^\circ$. Diurnal sky rotation and weekly

deck rotations provide uniform parallactic-angle coverage of the patch, and ensure that its peripheral regions are also observed by multiple polarimeters.

TABLE 2
REGULAR CALIBRATION OBSERVATIONS

Source	Schedule	Duration (min.)
sky dips	every 1.5 hours	3
Tau A	every 1–2 days	20
Moon	weekly	60
Jupiter	weekly	20
Venus	weekly	20
RCW38	weekly	20

4. CALIBRATION

Four quantities are required to convert polarimeter time streams into polarization power spectra: detector responsivities, a pointing model, detector polarization angles, and beam profiles. To this end, a suite of calibration observations is performed throughout the season using astronomical sources (Taurus A—hereafter Tau A, Jupiter, Venus, RCW38, and the Moon); atmospheric measurements (“sky dips,” which typically consist of three elevation nods of $\pm 3^\circ$); and instrumental sources (a rotating sparse wire grid and a polarized broadband noise source). From these we also measure instrumental polarization, as described below. QUIET’s regular calibration observations are summarized in Table 2.

We typically use two or more methods to determine a calibration constant, taking the spread among the methods as an indication of the uncertainty. We show in Section 8 that aside from the case of absolute responsivity, all calibration uncertainties lead to estimates of systematic effects on the power spectra well below statistical errors. This immunity comes from having a large number of detectors and highly-crosslinked polarization maps.




FIG. 3.— Polarimeter responses from the central feed horn to the polarization of Tau A at four parallactic angles. These data were collected with one correlation module in about 20 minutes. The errors are smaller than the points. From top to bottom, responses are shown for the detector diodes sensitive to the Stokes parameters +Q, −Q, +U, and −U, respectively. For each, the fitted model is plotted as a dashed line.

4.1. Responsivity

The polarized flux from Tau A provides a 5 mK signal which we observe at four parallactic angles. The sinusoidal modulation of the signal induced by the changing parallactic angles is fitted to yield responsivity coefficients for each detector. Figure 3 shows the response of the four polarization channels from the central feed horn to Tau A. A typical responsivity is 2.3 mV K^{-1} , with a precision from a single set of observations of 6%. The absolute responsivity from Tau A was measured most frequently for the central feed horn. We choose its +Q diode detector to provide the fiducial absolute responsivity.

The responsivities of other detectors relative to the fiducial detector are determined with the sky dips as described below. We have three independent means of assessing the relative responsivities among polarimeters: from nearly-simultaneous measurements of the Moon, from simultaneous measurements of responses to the rotating sparse wire grid in post-season tests, and from Tau A measurements. The errors from these methods are 4%, 2%, and 6% respectively, while the error from the sky-dip method is 4%. All the methods agree within errors.

Sky dips generate temperature signals of several 100 mK and thus permit measurement of the TP responsivities. The signals vary slightly with PWV. We estimate the slope from the data as $4\% \text{ mm}^{-1}$ and correct for it. This slope is consistent with the atmospheric model of Pardo et al. (2001). Because the ratios of the responsivities for the TP and polarized signals from each detector diode are stable quantities within a few percent of unity, we use sky dips performed at the beginning of each CES to correct short-term variations in the polarimeter responsivities. The responsivities vary by $\lesssim 10\%$ over the course of a day, due to changing thermal conditions for the bias electronics. Further post-season tests provide a physical model: the relevant temperatures are varied intentionally while the responsivities are measured with sky dips. We confirm the results with the polarized broadband source.

We bound the uncertainty in the absolute responsivity of the polarimeter array at 6%. The largest contributions to this estimate are uncertainties in (1) the beam solid angle (4%, see below), (2) the response difference between polarized and TP signals for each diode detector (3%), and (3) the Tau A flux (3%, Weiland et al. 2010). The first enters in converting the flux of Tau A into μK , while the second enters because although one fiducial diode detector is calibrated directly from Tau A, for the rest we find relative responsivities from sky dips and normalize by the fiducial diode’s responsivity.

For the differential-temperature modules, all detectors observe the signal from Jupiter simultaneously, providing the absolute responsivity for all channels upon comparison with the Jupiter flux from Weiland et al. (2010). Observations of Venus (Hafez et al. 2008) and RCW38 agree with the Jupiter measurements within errors, and sky dips track short-term variations. We calibrate the absolute responsivity with 5% accuracy.




FIG. 4.— Map of the polarization of the Moon from one detector diode. The amplitude of the quadrupole polarization visible here is $\approx 400 \text{ mK}$. Similar maps are produced for all 17 polarization modules in the array with a single \approx hour-long observation. The dotted line indicates the polarization orientation of the detector. Contours are spaced at intervals of 100 mK, with negative contours indicated by dashed lines.

4.2. Pointing

The global pointing solution derives from a physical model of the 3-axis mount and telescope tied to observations of the Moon with the central feed horn in the array, as well as Jupiter and Venus with the differential-temperature feed horns. Optical observations are taken regularly with a co-aligned star camera and used to monitor the time evolution of the pointing model.

During the first two months in the season, a mechanical problem with the deck-angle encoder results in pointing shifts. The problem was subsequently repaired. Based on pointing observations of the Moon and other astronomical sources, we verify that these encoder shifts are less than 2° . Systematic uncertainties induced by this problem are discussed in Section 8.1.

After the deck-angle problem is fixed, no significant evolution of the pointing model is found. The difference in the mean pointing solution between the start and the end of the season is smaller than $1'$. Observations of the Moon and Jupiter also provide the relative pointing among the feed horns. The root mean square (RMS) pointing error in the maps is $3'.5$.

4.3. Detector Polarization Angles

Our primary measurement of the polarization angle for each detector comes from observing the radial polarization of the Moon, as illustrated in Figure 4. The polarization angles are stable, changing by $< 0.2^\circ$, except during the period with the deck-angle–encoder problem mentioned above.

Two other less precise methods also give estimates of the detector angles: fits to the Tau A data, and determination of the phases of the sinusoidal responses of all the detectors to rotation of the sparse wire grid. In each case, the differences between the detector angles determined by the secondary method and the Moon are described by a standard deviation of $\approx 3^\circ$. However, we find a mean shift between the Tau A-derived and Moon-derived angles of 1.7° . To estimate the errors in the angles in light of this shift, we use an empirical approach: in Section 8.2 we estimate the impact on the power spectra from using the Tau A results instead of the Moon results, and find it to be small.




FIG. 5.— Top panel: Polarization beam profile from Tau A observations with the central feed horn. The data are overplotted with the expansion in Gauss-Hermite polynomials described in the text. Bottom panel: Beam window function with errors shown by the gray band.

4.4. Beam Profile and Window Function

The polarization and differential-temperature beams are obtained from maps created using the full data sets of Tau A and Jupiter observations respectively, with square pixels of $1/8$ on a side. For polarization, this process produces the main and leakage beam maps simultaneously, with the latter describing the instrumental polarization. The average FWHM for the beams across the array is $27/3$, measured with $0/1$ precision for the central feed horn and for the differential-temperature feed horns at the edge of the focal plane. The non-central-polarization-horn FWHMs are measured less frequently and thus are less precisely known, with an uncertainty of $1/5$. The beam elongation is typically small (1%), and its effect is further reduced by the diurnal sky rotation and weekly deck rotations which result in a symmetrized effective beam in the CMB maps. We compute 1-dimensional symmetrized beam profiles, with a resolution of $0/6$. These profiles are modeled as a sum of six even Gauss-Hermite terms (Monsalve 2010). The main-beam solid angles are computed by integrating these models out to $54'$ (roughly -28 dB), yielding $78.0 \pm 0.4 \mu\text{sr}$ for the differential-temperature horns and $74.3 \pm 0.7 \mu\text{sr}$ for the central horn. An average gives $76 \mu\text{sr}$ for all horns in the array. We also examine alternative estimates such as integrating the raw beam map instead of the analytical fit. We assign a systematic uncertainty of 4% based on the differences among these different estimates. The systematic error includes possible contributions from sidelobes, which we constrain to $0.7 \pm 0.4 \mu\text{sr}$ with antenna range measurements carried out before the observation season.

The window functions, encoding the effect of the finite resolution of the instrument on the power spectra, are computed from the central-horn and the temperature-horn-profile models. The central-horn beam profile and window function are shown in Figure 5. The uncertainty accounts for statistical error and differences between polarization and differential-temperature beams,

as described in Section 8.1.

4.5. Instrumental Polarization

Instrumental imperfections can lead to a spurious polarization signal proportional to the unpolarized CMB temperature anisotropy. We call this the I to Q (or U) leakage term. In our instrument, a fraction of the power input on one port of the correlation module is reflected because of a bandpass mismatch to the septum polarizer, and a fraction of the reflected power re-enters the other port. The dominant monopole term comes from this effect. We measure the monopole term from the polarimeter responses to temperature changes, using sky dipoles; Moon, Tau A, and Galactic signals; as well as variations from the weather. The average magnitude is 1.0% (0.2%) for the Q (U) diodes. Note that the discrepancy in the Q and U averages was predicted from measurements of the properties of the septum polarizers and confirmed in the field.

5. DATA ANALYSIS PROCEDURE

QUIET employs two independent analysis pipelines to derive CMB power spectra. We present the methods used for analysis in each pipeline, including data selection, filtering, map making, and power-spectra estimation.

Pipeline A is based on the pseudo- C_ℓ analysis framework, first described by Hivon et al. (2002), which is used by numerous experiments (Netterfield et al. 2002; Brown et al. 2009; Chiang et al. 2010; Larson et al. 2010; Lueker et al. 2010). This pipeline made all analysis choices in accordance with a strict (blind) analysis validation policy described in Section 6. An advantage of the pseudo- C_ℓ framework is computational efficiency, which is critical for completing the more than 30 iterations of the null-test suite. For the same reason, this pipeline is used for the systematic-error evaluations found in Section 8. Pseudo- C_ℓ analysis also enables us to perform cross correlation, making the resultant power spectra immune to possible misestimation of noise bias.

Pipeline B implements a maximum-likelihood framework (e.g., Tegmark 1997; Bond et al. 1998), which has a long history of use by CMB experiments (e.g., Mouskoff et al. 2000; Page et al. 2007; Wu et al. 2007; Bischoff et al. 2008). This framework yields minimum-variance estimates of the power spectra, naturally accounts for E/B mixing, and directly provides the exact CMB likelihood required for estimation of cosmological parameters, without the use of analytical approximations. In addition to power spectra, it produces unbiased maps with full noise-covariance matrices, useful for comparisons with other experiments. On the other hand, this approach is also computationally more expensive than the pseudo- C_ℓ framework, and a reduced set of null tests is therefore used to evaluate data consistency.

The processing of the time-ordered data (TOD) and the methodology used for data selection are treated in Sections 5.1 and 5.2, respectively. Brief descriptions of the pseudo- C_ℓ and maximum-likelihood techniques are found in Section 5.3. TOD processing, data selection, and analysis for temperature-sensitive modules are discussed in Section 5.4.

5.1. Time-Ordered-Data Processing

To prepare the TOD for map making, we execute three steps: pre-processing, noise modeling, and filtering. Of these steps, only the filtering is significantly different between the two pipelines.

5.1.1. Pre-processing

The first data-processing step is to correct for a small non-linearity that was discovered in the analog-to-digital converter (ADC) system. The non-linearities occur every 1024 bits; roughly 14% of the data are affected. Systematic uncertainty from this effect is estimated in Section 8.5. Next, the receiver data are synchronized with the telescope pointing. The double-demodulation step, described in Section 2, is applied, reducing the sample rate from 100 Hz to 50 Hz. A model of the detectors' polarized responsivities converts the data from ADC counts into thermodynamic temperature. The two pipelines use different responsivity models. Pipeline A applies a constant responsivity throughout each CES, addressing possible variability within a CES as part of the systematic error (Section 8); pipeline B updates responsivities on 2-minute timescales (Dumoulin 2010).

5.1.2. Noise Model

After pre-processing, the time streams for each detector diode in each CES are Fourier-transformed and their noise power spectra are fit to a model²⁹ with three parameters: the amplitude of white noise, the $1/f$ knee frequency, and the power-law slope of the $1/f$ noise. We also compute the white-noise correlations among detector diodes in the same module: the most important are between the two Q or the two U detector diodes (with an average coefficient of 0.22). A small fraction of the noise spectra contain features not accounted for in the noise model: beam sidelobes (see Section 2) scanning across features on the ground create a narrow spike at the scan frequency; slowly-changing weather patterns during a CES create a broader peak also at the scan frequency; and there are some narrow spikes at high ($\gtrsim 6$ Hz) frequencies. To prevent these features from biasing the noise model, the fit excludes a region around the scan frequency as well as frequencies above 4.6 Hz. In addition to the noise-model parameters, several statistics quantifying the agreement between the data and noise model are also used for data selection as described in Section 5.2.

5.1.3. Filtering

In pipeline A, three filters are applied. These were chosen from the results of many runs of the null-test suite (see Section 6). First, to remove the high-frequency narrow spikes, we apply a low-pass filter that cuts signals off sharply above 4.6 Hz³⁰. Second, to suppress contamination from atmospheric fluctuations and detector $1/f$ noise, we subtract a linear function from each telescope half scan (left-going or right-going) removing modes below twice the scan frequency³¹. The third filter, designed

to eliminate signal from ground emission, removes any azimuthal structure that remains after summing over all half scans in the CES.

In pipeline B, an apodized bandpass filter is used that accepts modes from 2.5 times the scan frequency to 4.5 Hz; the highpass component of this filter is designed to suppress scan-synchronous contamination. Further, a time-independent ground-emission model is subtracted. The model of ground emission is generated by building low-resolution and high-signal-to-noise maps in horizon coordinates from the full-season data for each deck angle and module, using large (55') pixels. Only features that are stable in time, azimuth, elevation, and deck angle contribute to this model. The amplitude of the ground correction is $\lesssim 1 \mu\text{K}$.

5.2. Data Selection

The fundamental unit of data used for analysis is the double-demodulated output of one detector diode for a single CES, referred to as a "CES-diode." Selecting only those CES-diodes that correspond to good detector performance and observing conditions is a critical aspect of the data analysis. The data-selection criteria began with a nominal set of cuts and evolved into several distinct configurations, as many as 33 in the case of pipeline A. For each configuration, analysis validation (see Section 6) was performed yielding statistics quantifying the lack of contamination in the data set. The final data set was chosen when these statistics showed negligible contamination and were little affected by changes to the cuts.

Cut efficiencies, defined as the fractions of CES-diodes accepted for the analysis, are given for both pipelines in Table 3. While each pipeline applies its own cuts uniformly to all four patches, the efficiencies among patches are non-uniform because of differences in weather quality. Over the course of the eight month observing season, patch CMB-1 is primarily visible at night, when the atmosphere tends to be more stable; patch CMB-3 is mostly observed during the day.

The first step of the data selection is simply to remove known bad data: data from six non-functional detector diodes, data during periods of mount malfunctions, and CESes lasting less than 1000 s. Further, we cut individual CES-diodes that show deviation from the expected linear relationship between the demodulated and TP signals. This cut removes data with poor thermal regulation of the electronics or cryostat, or residual ADC non-linearity.

The beam sidelobes, described in Section 2, introduce contamination to the data if the telescope scanning motion causes them to pass over the ground or the Sun. Ground pickup is dealt with by filtering as described in Section 5.1.3. The less frequent cases of Sun contamination are handled by cutting those CES-diodes for which the Sun's position overlaps with the measured sidelobe regions for each diode.

Additional cuts are specific to each pipeline. Pipeline A removes data taken during bad weather using a statistic calculated from fluctuations of the TP data during 10-s periods, averaged across the array. This cut removes entire CESes. Several more cuts remove individual CES-diodes. While these additional cuts are derived from the noise modeling statistics, they also target residual bad weather. During such marginal weather conditions only some channels need to be cut, since the sensitivity for

²⁹ At the level of a single CES, the TOD of each detector diode are dominated by noise; the contribution of the CMB is negligible.

³⁰ For QUIET's beam size and scanning speed a low-pass filter of 4.5–4.6 Hz results in a minimal loss of sensitivity to the CMB.

³¹ Typical scan frequencies range from 45 mHz to 100 mHz.

a given detector diode to atmospheric fluctuations depends on its level of instrumental polarization. Next, we reject CES-diodes with poor agreement between the filtered data and the noise model in three frequency ranges: a narrow range (only 40 Fourier modes) about the scan frequency, from twice the scan frequency to 1 Hz, and from 1 Hz to 4.6 Hz. We also cut CES-diodes that have higher than usual $1/f$ knee frequencies, or large variations during the CES in the azimuthal slopes of the double-demodulated time streams; both these cuts help eliminate bad weather periods. Finally, we also remove any CES-diodes with an outlier greater than 6σ in the time domain on three timescales (20 ms, 100 ms, and 1 s).

For pipeline B, the weather cut rejects CESes based on a statistic computed from fluctuations of the double-demodulated signals from the polarization modules on 10-s and 30-s timescales. Three cuts are applied to remove individual CES-diodes. The first is a cut on the $1/f$ knee frequency, similar to that of pipeline A. Second, a cut is made on the noise model χ^2 in the frequency range passed by the filter, and third, we reject CES-diodes having a large χ^2 in the azimuth-binned TOD. This cut rejects data with possible time variation in the ground signal. Finally, an entire CES is removed if more than 40% of its detectors have already been rejected.

5.3. Map Making and Power-Spectra Estimation

After filtering, the TOD for all diodes are combined to produce Q and U maps for each of the QUIET patches. The maps use a HEALPix $N_{\text{side}} = 256$ pixelization (Gorski et al. 2005). This section describes the map making and power-spectra estimation from the maps for each of the pipelines.

5.3.1. Pipeline-A Map Making

Polarization maps (Q and U) are made by summing samples into each pixel weighted by their inverse variance, calculated from the white-noise amplitudes. The full covariance matrix is not calculated. Two polarized sources, Centaurus A and Pictor A, are visible in the maps and are removed using circular top-hat masks with radii of 2° and 1° , respectively.

Separate maps are made for each range of telescope azimuth and deck-angle orientations. The coordinates are binned such that there are 10 divisions in azimuth³² and six distinct ranges of deck-angle orientation. Making separate maps for different telescope pointings enables the cross correlation described in the next section.

TABLE 3
TOTAL HOURS OBSERVED AND DATA-SELECTION EFFICIENCIES

Patch	Total Hours	A %	B %	Common %
CMB-1	905	81.7	84.3	76.7
CMB-2	703	67.3	70.0	61.2
CMB-3	837	56.0	61.4	51.4
CMB-4	223	70.6	74.2	65.9
All Patches	2668	69.4	72.9	64.2

NOTE. — Selection efficiencies for each pipeline. “Common” gives the efficiencies if both sets of cuts were applied.

³² The azimuth divisions are the same for all patches, which means that not all divisions are populated for patches CMB-3 and CMB-4.

5.3.2. Power-Spectra Estimation in Pipeline A

The MASTER (Monte Carlo Apodized Spherical Transform Estimator) method is used in pipeline A (Hivon et al. 2002; Hansen & Gorski 2003); it is based on a pseudo- C_ℓ technique and takes account of effects induced by the data processing using Monte Carlo (MC) simulations. The pseudo- C_ℓ method allows estimation of the underlying C_ℓ using spherical-harmonics transformations when the observations do not cover the full sky uniformly (Wandelt et al. 2001). The pseudo- C_ℓ spectrum, designated by \tilde{C}_ℓ , is related to the true spectrum C_ℓ by:

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle. \quad (1)$$

There is no term corresponding to noise bias, which would arise if we did not employ a cross-correlation technique. Here B_ℓ is the beam window function, described in Section 4.4, and $M_{\ell\ell'}$ is a mode-mode-coupling kernel describing the effect of observing only a small fraction of the sky with non-uniform coverage. It is calculable from the pixel weights, which are chosen to maximize the signal-to-noise ratio (Feldman et al. 1994). We bin in ℓ and recover C_ℓ in nine band powers, C_b , and F_ℓ is the transfer function (displayed in Section 7) due to filtering of the data; its binned estimate, F_b , is found by processing noiseless CMB simulations through pipeline A and used to obtain C_b . For the polarization power spectra, equation (1) is generalized for the case where \tilde{C}_ℓ contains both \tilde{C}_ℓ^{EE} and \tilde{C}_ℓ^{BB} .

In the power-spectra estimates, we include only the cross correlations among pointing-division maps, excluding the auto correlations. Because the noise is uncorrelated for different pointing divisions, the cross-correlation technique allows us to eliminate the noise-bias term and thus the possible residual bias due to its misestimate. Cross correlation between different pointing divisions also suppresses possible effects of ground contamination and/or time-varying effects. Dropping the auto correlations creates only a small increase in the statistical errors ($\approx 3\%$) on the power spectra.

The errors estimated for the pipeline-A power spectra are frequentist two-sided 68% confidence intervals. A likelihood function used to compute the confidence intervals is modeled following Hamimeche & Lewis (2008) and calibrated using the MC simulation ensemble of more than 2000 realizations with and without CMB signal. We also use the likelihood function to put constraints on r and calculate the consistency to Λ CDM.

The partial sky coverage of QUIET generates a small amount of E/B mixing (Challinor & Chon 2005), which contributes an additional variance to the BB power spectrum. We incorporate it as part of the statistical error. This mixing can be corrected (Smith & Zaldarriaga 2007) in future experiments where the effect is not negligible compared to instrumental noise.

5.3.3. Pipeline-B Map Making

In pipeline B, the pixel-space sky map $\hat{\mathbf{m}}$ ($N_{\text{side}} = 256$) is given by

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{d}, \quad (2)$$

where \mathbf{P} is the pointing matrix, \mathbf{N} is the TOD–noise-covariance matrix, \mathbf{F} corresponds to the apodized band-pass filter discussed in Section 5.1.3, and \mathbf{d} denotes the TOD. This map is unbiased, and for the case $\mathbf{F} = \mathbf{1}$ it is additionally the maximum-likelihood map, maximizing

$$\mathcal{L}(\mathbf{m}|\mathbf{d}) = e^{-\frac{1}{2}(\mathbf{d}-\mathbf{Pm})^T \mathbf{N}^{-1}(\mathbf{d}-\mathbf{Pm})}. \quad (3)$$

The corresponding map–noise-covariance matrix (e.g., Tegmark 1997; Keskitalo et al. 2010) is

$$\mathbf{N}_{\hat{\mathbf{m}}} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{P})^{-1} (\mathbf{P}^T \mathbf{F}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{P}) (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{P})^{-1}. \quad (4)$$

Note that one often encounters the simplified expression $\mathbf{N}_{\hat{\mathbf{m}}} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{F} \mathbf{P})^{-1}$ in the literature. This corresponds effectively to assuming that $\mathbf{F} = \mathbf{F}^2$ in the Fourier domain, and is strictly valid for top-hat–filter functions only. For our filters, we find that the simplified expression biases the map-domain χ^2 ($\equiv \hat{\mathbf{n}}^T \mathbf{N}_{\hat{\mathbf{m}}}^{-1} \hat{\mathbf{n}}$, where $\hat{\mathbf{n}}$ is a noise-only map) by $\approx 3\sigma$, and we therefore use the full expression, which does lead to an unbiased χ^2 .

Equations (2–4) apply to both polarization and temperature analysis. The only significant difference lies in the definition of the pointing matrix, \mathbf{P} . For polarization, \mathbf{P} encodes the detector orientation, while for temperature it contains two entries per time sample, +1 and –1, corresponding to the two horns in the differential-temperature assembly.

After map making, the maps are post-processed by removing unwanted pixels (i.e., compact sources and low-signal-to-noise edge pixels). All 54 compact sources in the 7-year *WMAP* point source catalog (Gold et al. 2010) present in our four patches are masked out, for a total of 4% of the observed area. We also marginalize over large-scale and unobserved modes by projecting out all modes with $\ell \leq 5$ ($\ell \leq 25$ for temperature) from the noise-covariance matrix using the Woodbury formula, assigning infinite variance to these modes.

5.3.4. Power-Spectra Estimation in Pipeline B

Given the unbiased map estimate, $\hat{\mathbf{m}}$, and its noise-covariance matrix, $\mathbf{N}_{\hat{\mathbf{m}}}$, we estimate the binned CMB power spectra, C_b , using the Newton–Raphson optimization algorithm described by Bond et al. (1998), generalized to include polarization. In this algorithm one iterates towards the maximum-likelihood spectra by means of a local quadratic approximation to the full likelihood. The iteration scheme in its simplest form is

$$\delta C_b = \frac{1}{2} \sum_{b'} \mathcal{F}_{bb'}^{-1} \text{Tr} [(\hat{\mathbf{m}} \hat{\mathbf{m}}^T - \mathbf{C})(\mathbf{C}^{-1} \mathbf{C}_{,b} \mathbf{C}^{-1})], \quad (5)$$

where b denotes a multipole bin, \mathbf{C} is the signal-plus-noise pixel-space covariance matrix, and $\mathbf{C}_{,b}$ is the derivative of \mathbf{C} with respect to C_b . The signal component of \mathbf{C} is computed from the binned power spectra, C_b , and the noise component is based on the noise model described in Section 5.1.2, including diode-diode correlations. Finally,

$$\mathcal{F}_{bb'} = \frac{1}{2} \text{Tr}(\mathbf{C}^{-1} \mathbf{C}_{,b} \mathbf{C}^{-1} \mathbf{C}_{,b'}) \quad (6)$$

is the Fisher matrix. Additionally, we introduce a step length multiplier, α , such that the actual step taken at

iteration i is $\alpha \delta C_b$, where $0 < \alpha \leq 1$ guarantees that \mathbf{C} is positive definite. We adopt the diagonal elements of the Fisher matrix as the uncertainties on the band powers.

We start the Newton–Raphson search at $C_\ell = 0$, and iterate until the change in the likelihood value is lower than 0.01 times the number of free parameters, corresponding roughly to a 0.01- σ uncertainty in the position of the multivariate peak. Typically we find that 3 to 10 iterations are required for convergence.

Estimation of cosmological parameters, θ , is done by brute-force grid evaluation of the pixel-space likelihood,

$$\mathcal{L}(\theta) \propto \frac{e^{-\frac{1}{2} \mathbf{d}^T \mathbf{C}^{-1}(\theta) \mathbf{d}}}{\sqrt{|\mathbf{C}(\theta)|}}. \quad (7)$$

Here $\mathbf{C}(\theta)$ is the covariance matrix evaluated with a smooth spectrum, C_ℓ , parametrized by θ . In this paper, we only consider 1-dimensional likelihoods with a parametrized spectrum of the form $C_\ell = a C_\ell^{\text{fid}}$, a being a scale factor and C_ℓ^{fid} a reference spectrum; the computational expense is therefore not a limiting factor. Two different cases are considered, with a being either the tensor-to-scalar ratio, r , or the amplitude of the EE spectrum, q , relative to the Λ CDM model.

5.4. Temperature Data Selection and Analysis

As described in Section 2, we dedicate one pair of modules to differential-temperature measurements. While these modules are useful for calibration purposes, when combined with our polarization data they also enable us to make self-contained measurements of the TE and TB power spectra.

For temperature, both pipelines adopt the pipeline-A data-selection criteria used for polarization analysis (see Section 5.2). The temperature-sensitive modules, however, are far more susceptible to atmospheric contamination than the polarization modules. Thus, these cuts result in reduced efficiencies: 12.4%, 6.9%, and 6.8% for patches CMB-1, CMB-2, and CMB-3, respectively³³. More tailoring of the cuts for these modules would improve efficiencies.

In pipeline A, the analysis proceeds as described in Sections 5.1.3, 5.3.1, and 5.3.2 except for two aspects. First, in the TOD processing a second-order polynomial is fit and removed from each telescope half scan instead of a linear function. This suppresses the increased contamination from atmospheric fluctuations in the temperature data. Second, we employ an iterative map maker based on the algorithm described by Wright et al. (1996). Map making for differential receivers requires that each pixel is measured at multiple array pointings or crosslinked. In order to improve crosslinking we divide the temperature data into only four maps by azimuth and deck angle, rather than the 60 divisions used for polarization analysis. To calculate TE and TB power spectra, polarization maps are made for these four divisions, plus one additional map that contains all polarization data with pointings not represented in the temperature data.

For pipeline B the algorithms for making temperature maps and estimating power spectra are identical to the polarization case, as described in Sections 5.3.3 and 5.3.4.

³³ Patch CMB-4 is excluded due to low data-selection efficiency and a lack of sufficient crosslinking.

6. ANALYSIS VALIDATION

The QUIET data analysis follows a policy of not looking at the power spectra until the analysis is validated using a set of predefined tests for possible systematic effects³⁴. The validation tests consist of a suite of null tests, comparisons across multiple analysis configurations, and consistency checks among power spectra from different CMB patches. Data-selection criteria, filtering methods, and the division of data into maps for cross correlation in pipeline A are all evaluated based on the test results. We finalize all aspects of the data analysis including calibration and evaluation of the systematic error before unveiling the power spectra (blind analysis). The risk of experimenter bias is thereby eliminated.

Details of tests found in this section describe pipeline A. While the pipeline B analysis follows a similar program of null tests to verify the result, the increased computational requirements of the maximum-likelihood framework limit the number of tests that could be performed and require those tests to be run using lower-resolution maps than for the non-null analysis. The bulk of this section treats validation of the polarization analysis; at the end, we briefly describe the temperature analysis validation.

In a null test, the data are split into two subsets. Maps, m_1 and m_2 , are made from each subset. The power spectra of the difference map, $m_{\text{diff}} \equiv (m_1 - m_2)/2$, are analyzed for consistency with the hypothesis of zero signal. The null suite consists of 42 tests³⁵, each targeting a possible source of signal contamination or miscalibration. These are highly independent tests; the data divisions for different null tests are correlated at only 8.8% on average. Nine tests divide the data by detector diode based on susceptibility to instrumental effects, such as instrumental polarization. Ten tests target effects that depend on the telescope pointing such as data taken at high or low elevation. Five tests divide based on the proximity of the main or sidelobe beams to known sources such as the Sun and Moon. Eight tests target residual contamination in the TOD using statistics mentioned in Section 5.2. Ten tests divide the data by environmental conditions such as ambient temperature or humidity.

Each null test yields EE and BB power spectra in nine ℓ bins, calculated separately for each CMB patch. Figure 6 shows the power spectra from one null test. Although the EB spectra are also calculated for each null test, they are assigned lesser significance since sources of spurious EB power will also result in the failure of EE and BB null tests. Combining all EE and BB points for all patches and null tests in the null suite yields a total of 3006 null-spectrum points. For each power-spectrum bin b , we calculate the statistic $\chi_{\text{null}} \equiv C_b^{\text{null}}/\sigma_b$, where C_b^{null} is the null power and σ_b is the standard deviation of C_b^{null} in MC simulations. We evaluate both χ_{null} and its square; χ_{null} is sensitive to systematic biases in the null spectra while χ_{null}^2 is more responsive to outliers. We run MC simulations of the full null suite to take into account the small correlation among the null tests and the slight non-

Gaussianity of the χ_{null} distribution. Non-Gaussianity is caused by the small number of modes at low ℓ .




FIG. 6.— EE and BB power spectra for the patch CMB-1 null test between Q and U detector diodes. The inset shows the low- ℓ region in detail.

As we refine the data-selection criteria based on the results of the null suite, we use a second test to monitor changes in the non-null power spectra. Using a blind analysis framework, we compute the difference of the power spectra between any two iterations of the data selection without revealing the non-null spectra. Further, we randomize the sign of the difference to hide the direction of the change; knowledge of the direction could allow experimenter bias (e.g. a preference for low BB power). Figure 7 shows the differences in the power spectra between the final configuration and several intermediate iterations of the data selection, starting with data sets that showed significant failures for the null-test suite. Statistically significant differences indicate a change in the level of contamination in the selected data set. Our data-selection criteria are finalized when further iterations only result in statistically expected fluctuations. The sensitivity of this test is demonstrated by the fact that the expected fluctuations are much less than the statistical error of the final result.

Finally, the non-null power spectra are compared among the four CMB patches. A χ^2 statistic is computed from the deviation of each patch's non-null power spectra from the weighted average over all patches. The total χ^2 is compared to MC simulations to compute probabilities to exceed (PTE).

When all aspects of the analysis are finalized, the last round of null tests and CMB patch comparisons validates the non-null-power-spectra results. Figure 8 shows the distributions of the χ_{null} statistic and of the PTEs corresponding to all χ_{null}^2 values from the full null suite. In pipeline A, the distribution of χ_{null} is consistent with the expectation from MC simulations. The mean of the χ_{null} distribution is 0.02 ± 0.02 ; the mean of the MC-ensemble χ_{null} distribution is also consistent with zero. The distribution of the χ_{null}^2 PTEs is uniform as expected. Table 4 lists the PTEs for the sums of the χ_{null}^2 statistic over all bins in each patch. Examinations of various subsets of the null suite, such as EE or BB only, do not reveal any anomalies. The EB null spectra do not indicate any failure either. Patch comparison PTEs are 0.16, 0.93, and 0.40 for EE, BB, and EB, respectively, demonstrating no statistically significant difference among the patches.

A similar, but smaller, null suite is run by pipeline B.

³⁴ Some systematic effects, such as a uniform responsivity-calibration error, cannot be detected by these techniques, and are addressed in Section 8.

³⁵ Only 41 null tests are performed for patch CMB-4; one test is dropped because there are no data in one of the subsets.




FIG. 7.— Power-spectra differences between the final data selection and six of the 32 earlier data-selection iterations, ordered by date. The lowest- ℓ bin of patch CMB-1 is shown. The error bars correspond to the expected fluctuations due to the differences in data selected, which are much smaller than the final statistical errors in this bin ($\approx 0.10 \mu\text{K}^2$ for BB). Iterations that are closer to the final data selection have smaller errors. The expected EE power in this bin from the ΛCDM model is also shown for comparison.

Specifically, 21 null tests are made at a HEALPix resolution of $N_{\text{side}} = 128$. The results obtained in these calculations are summarized in the bottom panel of Figure 8, and total PTEs for each patch are listed in Table 4. As in pipeline A, no anomalous values are found.

Finally, we make a comment on the usefulness of the χ_{null} distribution (as opposed to the χ_{null}^2 distribution) for identifying and quantifying potential contaminants. During the blind stage of the analysis, a positive bias in the χ_{null} distribution of 0.21 (0.19) was identified using pipeline A (B) (corresponding to 21% (19%) of the statistical errors). The number from pipeline A was obtained when including auto correlations in its power-spectra estimator. When excluding auto correlations, and cross-correlating maps made from data divided by time (day by day), the bias decreased to 0.10. Further detailed studies lead to the division of data into maps based on the telescope pointing, as described in Section 5.3; the result is an elimination of the observed bias.

The maximum-likelihood technique employed by pipeline B intrinsically uses auto correlations, and a corresponding shift in the χ_{null} distribution is seen in Figure 8. However, as will be seen in Section 7, the power spectra from the two pipelines are in excellent agreement, thereby confirming that any systematic bias coming from including auto correlations is well below the level of the statistical errors. We close this section by mentioning that we know of no other CMB experiment reporting an examination of the χ_{null} distribution, which is sensitive to problems not detected by examining the χ_{null}^2 distribution only.

6.1. Validation of the Temperature Analysis

A smaller number of null tests is used for the temperature analysis. Several are not applicable and others are discarded due to lack of data with sufficient crosslinking. Even so, we are able to run suites of 29, 27, and 23 TT null tests on patches CMB-1, CMB-2, and CMB-3, respectively. We calculate the sums of χ_{null}^2 statistics, yielding PTEs of 0.26 and 0.11 for patches CMB-1 and CMB-2, respectively. No significant outliers are found for

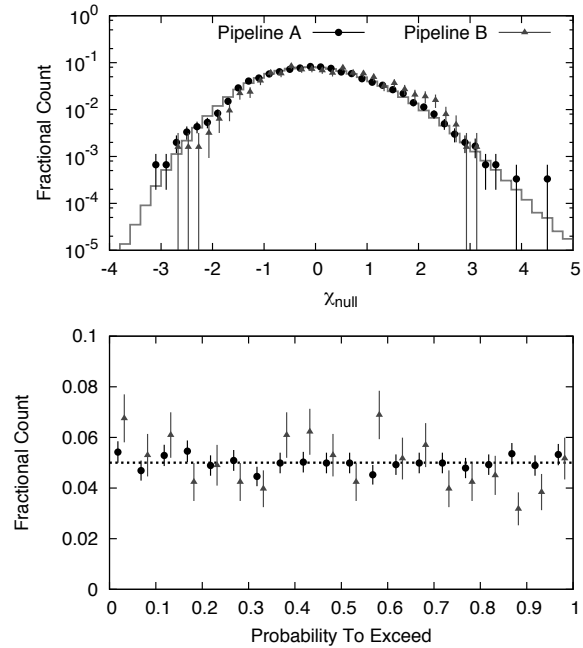


FIG. 8.— Null-Suite Statistics. The upper panel shows a histogram of the χ_{null} values for the pipeline-A null suite (circles), pipeline-B null suite (triangles), and the average of 1024 MC realizations of the pipeline-A null suite (gray histogram). Both data and MC distributions show similar non-Gaussianity in the χ_{null} statistic. The shift in χ_{null} seen for pipeline B, also seen in earlier iterations of pipeline A, is discussed in the text. The lower panel shows a histogram of PTEs calculated from the χ_{null}^2 statistic (outliers from either side of the upper distribution manifest as low PTEs).

TABLE 4
NULL SUITE PROBABILITY TO EXCEED BY PATCH

Patch	Pipeline A %	Pipeline B %
CMB-1	44	7
CMB-2	19	43
CMB-3	16	23
CMB-4	68	28

NOTE. — PTEs calculated from the sums of the χ_{null}^2 statistics, for EE and BB spectra points, over the null tests for each patch.

these patches. However, a $5\text{-}\sigma$ outlier in a single test³⁶ is found in patch CMB-3, implying contamination in its temperature map. CMB-3 is therefore excluded from further analysis. We confirm consistency between the patches CMB-1 and CMB-2 with a PTE of 0.26.

With no significant contamination in TT, EE, or BB spectra, one may be confident that the TE and TB spectra are similarly clean. For confirmation, we calculate TE and TB null spectra for the five null tests that are common to the temperature and polarization analyses. These yield PTEs of 0.61 and 0.82 for TE, and 0.16 and 0.55 for TB, for patches CMB-1 and CMB-2, respectively, with no significant outliers. Patch consistency checks give PTEs of 0.48 for TE and 0.26 for TB. Thus, the TE and TB power spectra, as well as the TT, pass all validation tests that are performed.

³⁶ This null test divides the data based on array pointing.

7. RESULTS

We report results from the first season of QUIET Q-band observations: CMB power spectra, derived foreground estimates, and constraints on the tensor-to-scalar ratio, r .

7.1. Polarization Power Spectra

The CMB power spectra are reported in nine equally-spaced bands with $\Delta\ell = 50$, beginning at $\ell_{\min} = 25$. Given the patch size, modes with $\ell < \ell_{\min}$ cannot be measured reliably. The correlation between neighboring bins is typically -0.1 ; it becomes negligible for bins further apart.

The EE, BB, and EB polarization power spectra estimated by both pipelines are shown in Figure 9. The agreement between the results obtained by the two pipelines is excellent, and both are consistent with the Λ CDM concordance cosmology. Our findings and conclusions are thus fully supported by both pipelines. Only the statistical uncertainties are shown here; we treat systematic errors in Section 8. Because the systematic error analysis was only done for pipeline A, we adopt its power-spectra results (tabulated in Table 5) as the official QUIET results.

The bottom sub-panels in Figure 9 show the window and transfer functions for each bin computed by pipeline A. Figure 10 shows the maps for patch CMB-1 computed by pipeline B, and Figure 11 shows the QUIET power spectra in comparison with the most relevant experiments in our multipole range. Additional plots and data files are online³⁷.

Fitting only a free amplitude, q , to the EE spectrum³⁸ relative to the 7-year best-fit *WMAP* Λ CDM spectrum (Larson et al. 2010), we find $q = 0.87 \pm 0.10$ for pipeline A and $q = 0.94 \pm 0.09$ for pipeline B. Taking into account the full non-Gaussian shapes of the likelihood functions, both results correspond to more than a $10\text{-}\sigma$ detection of EE power. In particular, in the region of the first peak, $76 \leq \ell \leq 175$, we detect EE polarization with more than $6\text{-}\sigma$ significance, confirming the only other detection of this peak made by BICEP at higher frequencies. The χ^2 relative to the Λ CDM model, with $C_\ell^{EE} = C_\ell^{BB} = 0$, is 31.6 (24.3) with 24 degrees of freedom, corresponding to a PTE of 14% (45%) for pipeline A (B).

7.2. Foreground Analysis

In order to minimize possible foreground contamination, QUIET's four CMB patches were chosen to be far from the Galactic plane and known Galactic synchrotron spurs. In these regions, contributions from thermal dust emission are negligible in Q band. Spinning dust is expected to be polarized at no more than a few percent in Q band (Battistelli et al. 2006; Lopez-Caraballo et al. 2010), so we expect the contribution to polarized foreground emission in our patches to be small. We therefore consider only two dominant sources of possible foreground contamination, namely compact radio sources and Galactic diffuse synchrotron emission.

³⁷ <http://quiet.uchicago.edu/results/index.html>

³⁸ Only $\ell \geq 76$ are used in the EE fit and the χ^2 calculation relative to Λ CDM because the first EE bin has a significant foreground contribution; see Section 7.2.

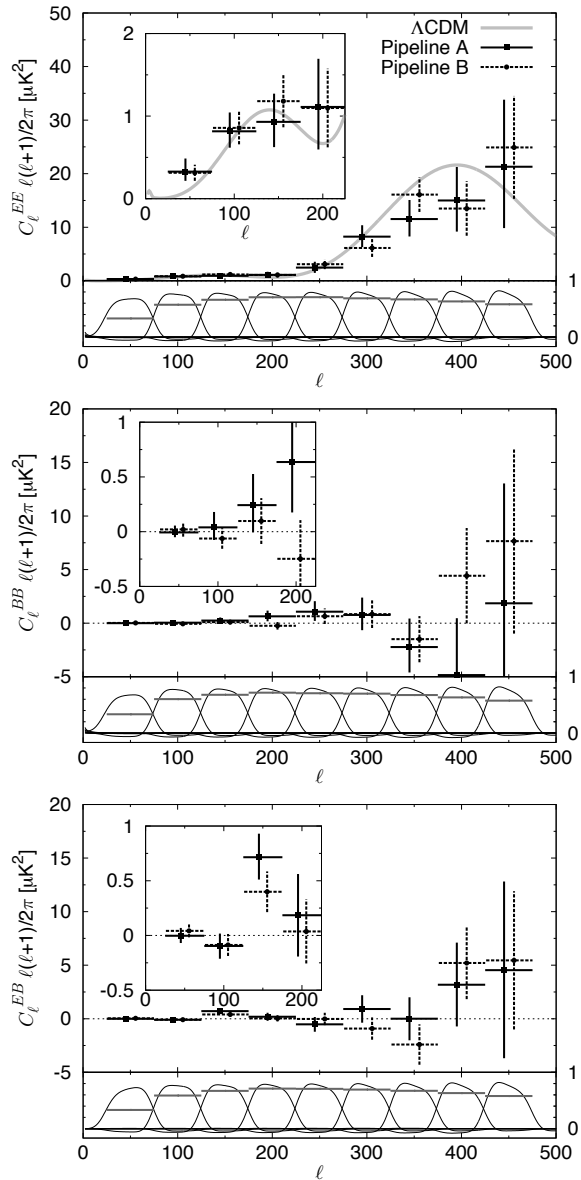


FIG. 9.— EE, BB, and EB power spectra from each QUIET pipeline, all four patches combined. The insets show the low- ℓ region in detail. Window and transfer functions for each ℓ bin are shown below the corresponding power spectra in black and gray, respectively. The window function combines the mode-mode-coupling kernel $M_{\ell\ell'}$ with the beam (B_ℓ) and represents, in combination with the transfer function (F_ℓ), the response in each band to the true C_ℓ spectrum. The EE point in the lowest- ℓ bin includes foreground contamination from patch CMB-1. For this display, pipeline A shows frequentist 68% confidence intervals while pipeline B uses the diagonal elements of the Fisher matrix; the difference is most pronounced in the lowest- ℓ bin where the likelihood is the most non-Gaussian.

To limit the effect of compact radio sources, we apply a compact-source mask to our maps before computing the power spectra, as described in Section 5. We also evaluate the CMB spectra both with and without the full *WMAP* temperature compact-source mask (Gold et al. 2010), and find no statistically significant changes. The possible contribution from compact radio sources with

TABLE 5
CMB-SPECTRA BAND POWERS FROM QUIET Q-BAND DATA

ℓ bin	EE	BB	EB
25-75	$^{+0.16}_{-0.11} 0.33^a$	$^{+0.06}_{-0.04} -0.01$	$^{+0.07}_{-0.07} 0.00$
76-125	$^{+0.23}_{-0.20} 0.82$	$^{+0.14}_{-0.12} 0.04$	$^{+0.11}_{-0.12} -0.10$
126-175	$^{+0.34}_{-0.31} 0.93$	$^{+0.28}_{-0.25} 0.24$	$^{+0.22}_{-0.20} 0.71$
176-225	$^{+0.58}_{-0.52} 1.11$	$^{+0.53}_{-0.46} 0.64$	$^{+0.38}_{-0.38} 0.18$
226-275	$^{+1.10}_{-0.99} 2.46$	$^{+0.98}_{-0.86} 1.07$	$^{+0.68}_{-0.69} -0.52$
276-325	$^{+2.1}_{-1.9} 8.2$	$^{+1.6}_{-1.4} 0.8$	$^{+1.3}_{-1.3} 0.9$
326-375	$^{+3.6}_{-3.3} 11.5$	$^{+2.7}_{-2.4} -2.2$	$^{+2.0}_{-2.0} 0.0$
376-425	$^{+6.2}_{-5.8} 15.0$	$^{+5.3}_{-4.9} -4.9$	$^{+3.9}_{-3.9} 3.2$
426-475	$^{+13}_{-11} 21$	$^{+11}_{-10} 2$	$^{+8.3}_{-8.2} 4.5$

NOTE. — Units are thermodynamic temperatures, μK^2 , scaled as $C_\ell \ell(\ell+1)/2\pi$.

^aPatch CMB-1 has significant foreground contamination in the first EE bin.

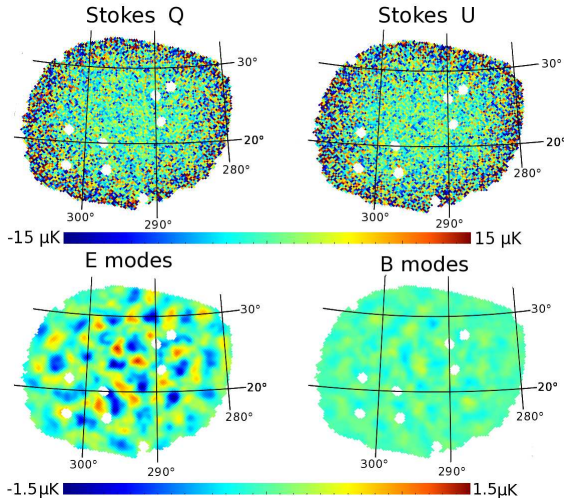


FIG. 10.— Maps of patch CMB-1 in Galactic coordinates. The top row shows our polarization maps with compact sources masked (white disks). The bottom row shows E and B modes decomposed using a generalized Wiener filter technique, implemented through Gibbs sampling (Eriksen et al. 2004; Larson et al. 2007), including only modes for $\ell \geq 76$ and smoothed to 1° FWHM; lower multipoles are removed due to a significant foreground contribution. Note the clear difference in amplitude: the E modes show a high-signal-to-noise cosmological signal while the B modes are consistent with noise.

fluxes below the *WMAP* detection level (1 Jy) is small: $0.003 \mu\text{K}^2$ at $\ell = 50$ and $0.01 \mu\text{K}^2$ at $\ell = 100$ (Battye et al. 2010). We therefore conclude that our results are robust with respect to contamination from compact radio sources and that the dominant foreground contribution comes from diffuse synchrotron emission.

In Figure 12 we show the power spectra measured from each patch. The CMB-1 EE band power for the first bin is $0.55 \pm 0.14 \mu\text{K}^2$, a $3\text{-}\sigma$ outlier relative to the expected ΛCDM band power of $0.13 \mu\text{K}^2$; while not significant enough to spoil the overall agreement among the patches as shown in Section 6, this is a candidate for a bin with foreground contamination.

To estimate the Q-band polarized synchrotron contamination in our CMB patches, we process the *WMAP* 7 K-band (23-GHz) map through pipeline A and estimate

its band power, \hat{C}_b^{KK} , as well as the cross spectra with the QUIET Q-band data, \hat{C}_b^{QK} . These results are shown for the first bin ($25 \leq \ell \leq 75$; $b = 1$) in Table 6, together with the corresponding QUIET band powers, \hat{C}_b^{QQ} . Since foregrounds do not contribute to the sample variance, the uncertainties for $\hat{C}_{b=1}^{\text{KK}}$ and $\hat{C}_{b=1}^{\text{QK}}$ are given by instrumental noise only, including contributions from both *WMAP* and QUIET. For $\hat{C}_{b=1}^{\text{QQ}}$, sample variance as predicted by the ΛCDM model is also included.

There is significant EE power in patch CMB-1 as measured by $\hat{C}_{b=1}^{\text{KK}}$. We also find a correspondingly significant cross correlation between the *WMAP* K band and the QUIET Q band, confirming that this excess power is not due to systematic effects in either experiment and is very likely a foreground. No significant power is found in any other case. The non-detection of foreground power at $\ell > 75$ is consistent with the expected foreground dependence: $\propto \ell^{-2.5}$ (Carretti et al. 2010), and the low power found in $\hat{C}_{b=1}^{\text{KK}}$.

The excess power observed in the first EE bin of CMB-1 is fully consistent with a typical synchrotron frequency spectrum. To see this, we extrapolate $\hat{C}_{b=1}^{\text{KK}}$ from K band to Q band, assuming a spectral index of $\beta = -3.1$ (Dunkley et al. 2009), and calculate the expected power in $C_{b=1}^{\text{QK}}$ and $C_{b=1}^{\text{QQ}}$,

$$C_{b=1}^{\text{QK}} = \frac{1.05}{1.01} \left(\frac{43.1}{23} \right)^\beta \hat{C}_{b=1}^{\text{KK}} = 2.57 \pm 0.69 \mu\text{K}^2, \quad (8)$$

$$C_{b=1}^{\text{QQ}} = \left[\frac{1.05}{1.01} \left(\frac{43.1}{23} \right)^\beta \right]^2 \hat{C}_{b=1}^{\text{KK}} = 0.38 \pm 0.10 \mu\text{K}^2, \quad (9)$$

where the prefactor accounts for the fact that β is defined in units of antenna temperature, and the uncertainties are scaled from that of $\hat{C}_{b=1}^{\text{KK}}$. These predictions are fully consistent with the observed values of $\hat{C}_{b=1}^{\text{QK}}$ and $\hat{C}_{b=1}^{\text{QQ}}$, when combined with the ΛCDM -expected power. We conclude that the excess power is indeed due to synchrotron emission.

7.3. Constraints on Primordial B modes

We constrain the tensor-to-scalar ratio, r , using the QUIET measurement of the BB power spectrum at low multipoles ($25 \leq \ell \leq 175$). Here r is defined as the ratio of the primordial-gravitational-wave amplitude to the curvature-perturbation amplitude at a scale $k_0 = 0.002 \text{Mpc}^{-1}$. We then fit our measurement to a BB-spectrum template computed from the ΛCDM concordance parameters with r allowed to vary. For simplicity, we fix the tensor spectral index at $n_t = 0$ in computing the template³⁹. This choice makes the BB-power-spectrum amplitude directly proportional to r .

For pipeline A, we find $r = 0.35^{+1.06}_{-0.87}$, corresponding to $r < 2.2$ at 95% confidence. Pipeline B obtains $r = 0.52^{+0.97}_{-0.81}$. The results are consistent; the lower panel of Figure 11 shows our limits on BB power in comparison with those from BICEP, QUaD, and *WMAP*. QUIET lies between BICEP and *WMAP* in significantly limiting r

³⁹ Our definition of r agrees with Chiang et al. (2010)




FIG. 11.— The top panel shows EE results with 68% C.L. error bars; the bottom panel shows BB 95% C.L. upper limits. For comparison, we also plot results from previous experiments (Brown et al. 2009; Chiang et al. 2010; Larson et al. 2010) and the Λ CDM model (the value $r = 0.2$ is currently the best 95% C.L. limit on tensor modes).

from measurements of CMB-B-mode power in our multipole range. Although we neither expected nor detected any BB foreground power, the detection of an EE foreground in patch CMB-1 suggests that BB foregrounds might be present at a smaller level. We emphasize that the upper limit we report is therefore conservative.

7.4. Temperature Power Spectra

Figure 13 compares the QUIET and WMAP Q-band temperature maps and TT, TE, and TB power spectra. Agreement with the Λ CDM model is good. This is a strong demonstration of the raw sensitivity of the QUIET detectors; the single QUIET differential-temperature assembly produces a high-signal-to-noise map using only 189 hours (after selection) of observations. The high sensitivity of these modules makes them very useful for calibration, pointing estimation, and consistency checks (see Section 4).

8. SYSTEMATIC ERRORS

The passing of the null suite itself limits systematic uncertainty, but to get well below the statistical errors, dedicated studies are needed. They are important in gaining confidence in the result and also in evaluating the potential of the methods and techniques we use for future efforts. We pay special attention to effects that can generate false B-mode signals. Our methodology is to simulate and then propagate calibration uncertainties (see Section 4) and other systematic effects through the entire pipeline. The systematic errors in the power spectra are shown in Figure 14. The possible contaminations are well below the statistical errors; in particular, the levels of spurious B modes are less than the signal of $r = 0.1$. This is the lowest level of BB contamination yet reported by any CMB experiment. This section describes how each effect in Figure 14 is determined and considers three additional possible sources of contamination.

An uncertainty not shown in Figure 14 is that aris-

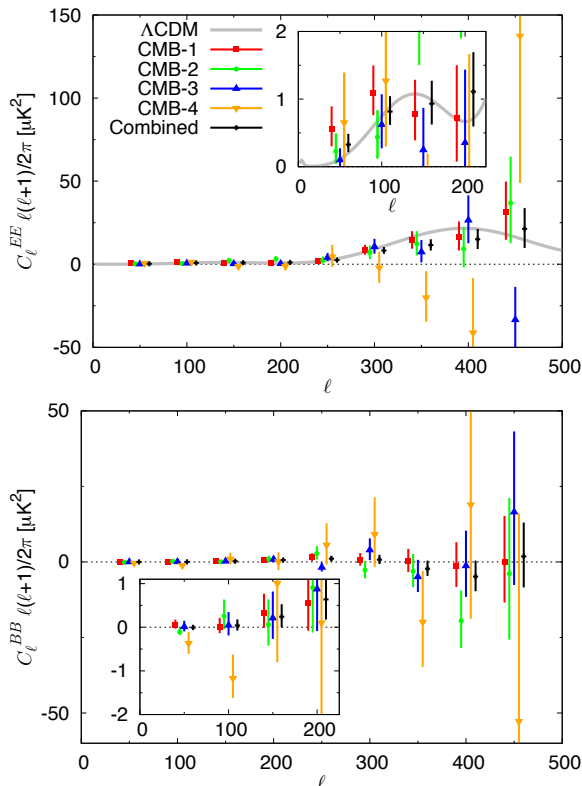


FIG. 12.— CMB power spectra are shown for each patch individually. The top and bottom panels show the EE and BB spectra, respectively. The different error bars for each patch mainly reflect the amounts of time each was observed.

TABLE 6
BAND AND CROSS POWERS FOR $\ell = 25\text{--}75$

Patch	Spectrum	$\hat{C}_{b=1}^{\text{KK}}$	$\hat{C}_{b=1}^{\text{QK}}$	$\hat{C}_{b=1}^{\text{QQ}}$
CMB-1	EE	17.4 ± 4.7	3.30 ± 0.55	0.55 ± 0.14
	BB	4.8 ± 4.5	0.40 ± 0.41	0.06 ± 0.08
	EB	-6.2 ± 3.2	0.27 ± 0.38	0.10 ± 0.08
CMB-2	EE	5.5 ± 3.7	0.01 ± 0.56	0.23 ± 0.19
	BB	4.6 ± 3.4	0.18 ± 0.48	-0.11 ± 0.13
	EB	-5.5 ± 2.8	-0.39 ± 0.41	-0.20 ± 0.12
CMB-3	EE	0.2 ± 1.9	0.64 ± 0.43	0.10 ± 0.18
	BB	-0.3 ± 2.6	0.33 ± 0.35	0.01 ± 0.13
	EB	1.4 ± 1.7	-0.34 ± 0.30	-0.27 ± 0.11
CMB-4	EE	-5.2 ± 5.1	0.7 ± 1.2	0.65 ± 0.58
	BB	-2.6 ± 5.2	-0.1 ± 1.1	-0.37 ± 0.52
	EB	-1.0 ± 3.9	0.0 ± 0.9	-0.15 ± 0.47

NOTE. — Power-spectra estimates for the first multipole bin for each patch, computed from the *WMAP7* K-band data and the QUIET Q-band data. The units are $\ell(\ell+1)C_\ell/2\pi$ (μK^2) in thermodynamic temperature. Uncertainties for $\hat{C}_{b=1}^{\text{KK}}$ and $\hat{C}_{b=1}^{\text{QK}}$ include noise only. For $\hat{C}_{b=1}^{\text{QQ}}$ they additionally include CMB sample variance as predicted by ΛCDM . Values in bold are more than 2σ away from zero.

ing from the overall responsivity error estimate of 6% (12% in power-spectra units). After including the effect of possible time-dependent responsivity variations (4%, see below), the power-spectra uncertainty is 13%. It is multiplicative, affecting all power-spectra results independent of multipole.




FIG. 13.— The top row compares our temperature map to the *WMAP7*-year Q-band map (Jarosik et al. 2010) for patch CMB-1 in Galactic coordinates. Lower panels show the CMB temperature power spectra: TT, TE, and TB.

8.1. Beam Window Function and Pointing

The uncertainty in the beam window function is another multiplicative factor, one which increases with multipole. We estimate this uncertainty using the difference of the beam window functions measured for the central module and the modules of the differential-temperature assembly, which are at the edge of the array. The difference is statistically significant, coming from the different locations (with respect to the optics) in the focal plane; it is expected from the pre-season antenna range measurements.

Uncertainties in pointing lead to distortions in polarization maps. E power will be underestimated and spurious B power (if the distortions are non-linear) generated (Hu et al. 2003). We quantify these effects by using the differences in pointing solutions from two independent models: the fiducial model used for the analysis and an alternative model based on a different set of calibrating observations. We also modeled and included the effects of the deck-angle-encoder shift which occurred for a portion of the season (Section 4.3).




FIG. 14.— Systematic uncertainty estimates for EE, BB, and EB power spectra. Estimates for a variety of effects (see text) are shown for the three power spectra. In all cases, they are well below the statistical errors, which are also shown. In particular, the contaminations to the primordial-B-mode signal, at multipoles below 100, are below the level of $r = 0.1$, even though we do not make a correction for the largest contaminant, the monopole leakage.

8.2. Responsivity and Polarization Angle

Responsivity shifts, particularly within CESes, lead to distortions in the maps. Full-pipeline simulations quantify the shifts caused by variations in the cryostat or electronics temperatures. Similarly shifts from using responsivities determined from the Moon data, Tau A data, or from the sparse wire grid, rather than those from the sky dipoles, are determined. We also incorporate the uncertainty in the atmospheric-temperature model used in analyzing the sky-dip data. The largest possible effects on the power spectra are shown in Figure 14.

Uncertainties in the orientation of the polarization axes of the modules can lead to leakage between E and B modes. To quantify this leakage, we use the differences in power spectra where these angles are determined from Moon data, Tau A data, and the sparse-wire-grid data. As expected, the largest effects show up in EB power.

8.3. Instrumental Polarization

As described in Section 4.5, the I to Q (U) leakage coefficients for the QUIET detector diodes are small: 1% (0.2%). Except in the case of patch CMB-4, our scanning strategy significantly reduces this effect with the combination of sky and deck-angle rotation.

We estimate spurious Q and U in the maps for each CES-diode using the *WMAP* temperature map and our known leakages. Shown in Figure 14 are the estimates of spurious EE, BB, and EB powers from full-pipeline simulations, where for each realization the spurious Q and U are added to the Q and U from simulated Λ CDM E modes. While this method has an advantage of being able to use the real (not simulated) temperature map, it does not incorporate TE correlation, which only affects the spurious EE power. As a complement, we repeat the

study, but using simulated Λ CDM maps for both temperature and polarization; this only changes the estimate of spurious EE power by 30% at most. Because the spurious power is as small as it is, we have treated it as a systematic rather than correcting for it. Doing so would give us a further order of magnitude suppression.

Differing beam ellipticities can also induce higher multipole polarization signals. We measure these leakages from Tau A and Jupiter observations and find that the higher-order multipoles are at most 0.1% of the main-beam peak amplitude. The corresponding effects on the power spectra, which are seen in Figure 14, are of little concern.

8.4. Far Sidelobes Seeing the Sun

While we make cuts to reduce the effects of far sidelobes seeing the Sun (Sections 2 and 5.1.3), small contaminations could remain. We make full-season maps for each diode in Sun-centered coordinates and then use these maps to add contamination to full-pipeline CMB simulations. The excess power found in the simulations is taken as the systematic uncertainty.

8.5. Other Possible Sources of Systematic Uncertainty

Here we discuss a few additional potential sources of systematic uncertainty, which are found to be subdominant.

Ground-Synchronous Signals. QUIET’s far sidelobes do see the ground for some diodes at particular elevations and deck angles. Ground pickup that is constant throughout a CES is removed by our TOD filters; the net effect of this filtering in the full-season maps is a correction of $\approx 1 \mu\text{K}$.

The only concern is ground pickup that changes over

the short span of a single CES. We find little evidence for changes even over the entire season, let alone over a single CES. We therefore conservatively place an upper limit on such changes using the statistical errors on the ground-synchronous signal. We start with the CES and module with the largest ground pickup. We then simulate one day’s worth of data, inserting a ground-synchronous signal that changes by its statistical error. Given the distribution in the magnitude of the ground-synchronous signal and assuming that changes in this signal are proportional to the size of the signal itself, by considering that the signals from changing pickup add incoherently into the maps made from multiple CES-diodes at a variety of elevations and deck angles, we estimate an upper limit on residual B power from possible changing ground-pickup signals. The result is $\lesssim 10^{-4} \mu\text{K}^2$ at multipoles below 100.

ADC Non-linearities. The possible residual after the correction for the non-linearity in the ADC system results in effects similar to the I to Q (or U) leakage and the variation of the responsivity during the CES. We estimate such effects based on the uncertainty in the correction parameters, confirming that there is at most a 3% additional effect for the leakage bias, and that the responsivity effect is also small, less than half of the systematic error shown for the responsivity in Figure 14.

Data-Selection Biases. Cuts can cause biases if they are, for example, too stringent. We expect none but to be sure we apply our selection criteria to 144 CMB + noise simulations. No bias is seen, and in particular we limit any possible spurious B modes from this source to $\lesssim 10^{-3} \mu\text{K}^2$ at multipoles below 100.

9. CONCLUSIONS

QUIET detects polarization in the EE power spectrum at 43 GHz. We confirm with high significance the detection of polarization in the region of the first acoustic peak (Chiang et al. 2010) in the multipole region $\ell = 76$ –175. We find no significant power in either BB or EB between $\ell = 25$ and $\ell = 475$. We measure the tensor-to-scalar ratio to be $r = 0.35^{+1.06}_{-0.87}$.

These results are supported by a very extensive suite of null tests in which 42 divisions of data were used for each of 33 different cut configurations. The selection criteria and systematic errors were determined before the power spectra themselves were examined. Biases were revealed during this process, the last of which was a contamination present in the null spectra at the level of about 20% of the statistical errors, but eliminated when cross-correlating maps with differing telescope pointings. The robustness of the final results is further supported by having two pipelines with results in excellent agreement, even though one uses only cross correlations while the other also uses auto correlations.

Several possible systematic effects are studied with full end-to-end simulations. The possible contaminations in the B-mode power are thereby limited to a level smaller than for any other published experiment: below the level of $r = 0.1$ for the primordial B modes; simply correcting for the known level of instrumental polarization would reduce this to $r < 0.03$. This very low level of systematic uncertainty comes from the combination of several important design features, including a new time-stream “double-demodulation” technique, Mizuguchi–Dragone

optics, natural sky rotation, and frequent deck rotation.

The correlation modules we use have a polarization sensitivity (Q and U combined) of $280 \mu\text{K}\sqrt{\text{s}}$, leading to an array sensitivity of $69 \mu\text{K}\sqrt{\text{s}}$. Further, the $1/f$ noise observed in our detectors is small: the median knee frequency is just 5.5 mHz. One important outcome of this work, then, is the demonstration that our detectors, observing from a mid-latitude site, give excellent sensitivity and systematic immunity.

Because of our mid-latitude site, we are driven to collect data in four separate patches. While we lose some sensitivity (compared to going deeper on a single patch), there are a few advantages that we have exploited. The patches are scanned differently, in terms of time of day and the degree of crosslinking, and these differences allow some important systematic checks. Another advantage concerns foregrounds.

Foreground contamination is expected to be one of the main limiting factors in the search for primordial B modes. Indeed we report a $3\text{-}\sigma$ detection of synchrotron emission in one of our four CMB patches, originally chosen for their expected low foreground levels. Our detection is only in EE but assuming a similar BB level and extrapolating to the foreground minimum of about 95 GHz, we would have synchrotron contamination at the level of $r = 0.02$. Neither *WMAP* nor *Planck* will have enough sensitivity (Tauber et al. 2010) to sufficiently constrain the polarized synchrotron amplitude at this level. In fact, our Q-band polarization maps are already as deep or deeper than what *Planck* will achieve at the same frequency. Dedicated low-frequency observations are clearly needed to achieve such constraints. When foreground cleaning becomes important, consistency among separate patches will be an important handle on our understanding.

Further progress must be made through larger arrays and longer integration times. In hand we have data collected by the 90-element W-band array with similar sensitivity to our Q-band array and more than twice the number of observing hours. Results from the analysis of that data set will be reported in future publications. A W-band receiver with the sensitivity to reach below the level of $r = 0.01$ is under development.

Support for the QUIET instrument and operation comes through the NSF cooperative agreement AST-0506648. Support was also provided by NSF awards PHY-0355328, AST-0448909, AST-1010016, and PHY-0551142; KAKENHI 20244041, 20740158, and 21111002; PRODEX C90284; a KIPAC Enterprise grant; and by the Strategic Alliance for the Implementation of New Technologies (SAINT).

Some work was performed on the Joint Fermilab-KICP Supercomputing Cluster, supported by grants from Fermilab, the Kavli Institute for Cosmological Physics, and the University of Chicago. Some work was performed on the Titan Cluster, owned and maintained by the University of Oslo and NOTUR (the Norwegian High Performance Computing Consortium), and on the Central Computing System, owned and operated by the Computing Research Center at KEK. Portions of this work were performed at the Jet Propulsion Laboratory (JPL) and California Institute of Technology, operating under

a contract with the National Aeronautics and Space Administration. The Q-band polarimeter modules were developed using funding from the JPL R&TD program.

C.D. acknowledges an STFC Advanced Fellowship and an ERC IRG grant under FP7. P.G.F. and J.A.Z. gratefully acknowledge the support of the Beecroft Institute of Particle Astrophysics and Cosmology, the Oxford Martin School, and the Science and Technology Facilities Council. L.B., R.B., and J.M. acknowledge support from CONICYT project Basal PFB-06. A.D.M. acknowledges a Sloan foundation fellowship.

PWV measurements were provided by the Atacama Pathfinder Experiment (APEX). We thank CONICYT for granting permission to operate within the Chajnantor Scientific Preserve in Chile, and ALMA for providing site infrastructure support. We are particularly indebted to the engineers and technician who maintained and operated the telescope: José Cortés, Cristobal Jara, Freddy Muñoz, and Carlos Verdugo.

In addition, we would like to acknowledge the following people for their assistance in the instrument design, construction, commissioning, operation, and in data analysis: Augusto Gutierrez Aitken, Colin Baines, Phil Bannister, Hannah Barker, Matthew R. Becker, Alex Blein, Mircea Bogdan, April Campbell, Anushya Chandra, Sea Moon Cho, Emma Curry, Maire Daly, Richard Davis, Fritz Dejongh, Joy Didier, Greg Dooley, Hans Eide, Will Grainger, Jonathon Goh, Peter Hamlington, Takeo Higuchi, Seth Hillbrand, Christian Holler, Ben Hooberman, Kathryn D. Huff, William Imbriale, Koji Ishidoshiro, Eiichiro Komatsu, Jostein Kristiansen, Richard Lai, Erik Leitch, Kelly Lepo, Martha Malin, Mark McCulloch, Oliver Montes, David Moore, Makoto Nagai, Ian O'Dwyer, Stephen Osborne, Stephen Padin, Felipe Pedreros, Ashley Perko, Alan Robinson, Jacklyn Sanders, Dale Sanford, Mary Soria, Alex Sugarbaker, David Sutton, Matias Vidal, Liza Volkova, Edward Wollack, Stephanie Xenos, and Mark Zaskowski.

REFERENCES

- Battistelli, E. S., Rebolo, R., Rubiño-Martín, J. A., Hildebrandt, S. R., Watson, R. A., Gutiérrez, C., & Hoyland, R. J. 2006, *ApJ*, 645, L141
- Battye, R. A., Browne, I. W. A., Peel, M. W., Jackson, N. J., & Dickinson, C. 2010, arXiv:1003.5846v1 [astro-ph.CO]
- Bischoff, C. et al. 2008, *ApJ*, 684, 771
- Bond, J. R., Jaffe, A. H., & Knox, L. 1998, *Phys. Rev. D*, 57, 2117
- Bornemann, J., & Labay, V. A. 1995, *IEEE Trans. MTT*, 43, 95
- Brown, M. L., et al. 2009, *ApJ*, 705, 978
- Buder, I. 2010, *Proc. SPIE*, 7741, 77411D
- Carretti, E., et al. 2010, *MNRAS*, 405, 1670
- Challinor, A., & Chon, G. 2005, *MNRAS*, 360, 509
- Chiang, H. C. et al. 2010, *ApJ*, 711, 1123
- Dragone, C. 1978, *Bell Syst. Tech. J.*, 57, 2663
- Dumoulin, R. N. 2010, *Proc. SPIE*, 7741, 77412N
- Dunkley, J. et al. 2009, *ApJ*, 701, 1804
- Eriksen, H. K. et al. 2004, *ApJS*, 155, 227
- Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, *ApJ*, 426, 23
- Gold, B. et al. 2010, arXiv:1001.4555v2 [astro-ph.GA]
- Gorski, K. M., et al. 2005, *ApJ*, 622, 759, <http://healpix.jpl.nasa.gov/>
- Gundersen, J., & Wollack, E. 2009, *Journal of Physics: Conf. Ser.*, 155, 012005
- Güsten, R., Nyman, L. Å., Schilke, P., Menten, K., Cesarsky, C., & Booth, R. 2006, *A&A*, 454, L13
- Hafez, Y., et al. 2008, *MNRAS*, 388, 1775
- Hamimeche, S., & Lewis, A. 2008, *Phys. Rev. D*, 77, 103013
- Hansen, F. K., & Gorski, K. M. 2003, *MNRAS*, 343, 559
- Hivon, E., Gorski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Hu, W., Hedman, M., & Zaldarriaga, M. 2003, *Phys. Rev. D*, 67, 043004
- Imbriale, W. A., Gundersen, J., & Thompson, K. L. 2010, To appear in the Special Issue on Radio Astronomy in the *IEEE Transactions on Antennas and Propagation*
- Jarosik, N. et al. 2003, *ApJS*, 145, 413
- Jarosik, N., et al. 2010, arXiv:1001.4744v1 [astro-ph.CO]
- Kangaslahti, P. et al. 2006, in *Microwave Symposium Digest*, 2006. *IEEE MTT-S International*, 89–92
- Keskitalo, R. et al. 2010, *A&A*, 522, A94
- Komatsu, E., et al. 2010, arXiv:1001.4538v3 [astro-ph.CO]
- Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzappel, W. L. 2002, *Nature*, 420, 772
- Kusaka, A., et al. 2010, To appear in *Proceedings of the 45th Rencontre de Moriond*
- Larson, D. et al. 2010, arXiv:1001.4635v2 [astro-ph.CO]
- Larson, D. L., Eriksen, H. K., Wandelt, B. D., Górski, K. M., Huey, G., Jewell, J. B., & O'Dwyer, I. J. 2007, *ApJ*, 656, 653
- Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, *ApJ*, 624, 10
- Liddle, A. R., & Lyth, D. H. 2000, *Cosmological Inflation and Large-Scale Structure*
- Lopez-Caraballo, C. H., Rubino-Martin, J. A., Rebolo, R., & Genova-Santos, R. 2010, arXiv:1011.1242v1 [astro-ph.CO]
- Lueker, M., et al. 2010, *ApJ*, 719, 1045
- Mauskopf, P. D. et al. 2000, *ApJ*, 536, L59
- Mizugutch, Y., Akagawa, M., & Yokoi, H. 1976, *IEEE Antennas and Propagation Society International Symposium*, 14, 2
- Monsalve, R. A. 2010, *Proc. SPIE*, 7741, 77412M
- Montroy, T. E. et al. 2006, *ApJ*, 647, 813
- Netterfield, C. B., et al. 2002, *ApJ*, 571, 604
- Newburgh, L., et al. 2010, To appear in *Proceedings of the Twelfth Marcel Grossmann Meeting on General Relativity*
- Padin, S., et al. 2002, *PASP*, 114, 83
- Page, L. et al. 2007, *ApJS*, 170, 335
- Pardo, J., Cernicharo, J., & Serabyn, E. 2001, *IEEE Transactions on Antennas and Propagation*, 49, 1683
- Sievers, J. L. et al. 2007, *ApJ*, 660, 976
- Smith, K. M., & Zaldarriaga, M. 2007, *Phys. Rev. D*, 76, 043001
- Tauber, J. A., et al. 2010, *A&A*, 520, A1
- Tegmark, M. 1997, *ApJ*, 480, L87
- Wandelt, B. D., Hivon, E., & Górski, K. M. 2001, *Phys. Rev. D*, 64, 083003
- Weiland, J. L. et al. 2010, arXiv:1001.4731v1 [astro-ph.CO]
- Wright, E. L., Hinshaw, G., & Bennett, C. L. 1996, *ApJ*, 458, L53
- Wu, J. H. P. et al. 2007, *ApJ*, 665, 55