SLAC-PUB-11713
February 2006

The Silicon Tracker Readout Electronics of the Gamma-ray Large Area Space Telescope

Luca Baldini, Alessandro Brez, Thomas Himel, Masaharu Hirayama, R. P. Johnson, Wilko Kroeger, Luca Latronico, Massimo Minuti, David Nelson, Riccardo Rando, H. F.-W. Sadrozinski, *Senior Member, IEEE,* Carmelo Sgro', Gloria Spandre, E. N. Spencer, Mutsumi Sugizaki, Hiro Tajima, Johann Cohen-Tanugi, Marcus Ziegler

***Abstract*—A unique electronics system has been built and tested for reading signals from the silicon-strip detectors of the Gamma-ray Large Area Space Telescope mission. The system amplifies and processes signals from 884,736 36-cm strips using only 160 W of power, and it achieves close to 100% detection efficiency with noise occupancy sufficiently low to allow it to self trigger. The design of the readout system is described, and results are presented from ground-based testing of the completed detector system.**

***Index Terms*— Application specific integrated circuits, Data acquisition, Gamma-ray astronomy detectors, Multichip modules, Silicon radiation detectors**

I. INTRODUCTION

The Large Area Telescope (LAT) of the Gamma-ray Large-Area Space Telescope (GLAST) mission [1]–[2] is a pair-conversion gamma-ray detector similar in concept to the previous NASA high-energy gamma-ray mission EGRET on the Compton Gamma-Ray Observatory [3]. High energy (>20 MeV) gamma rays convert into electron-positron pairs in one of 16 layers of tungsten foils. The charged particles pass through up to 36 layers of position-sensitive detectors interleaved with the tungsten, the "tracker," leaving behind tracks pointing back toward the origin of the gamma ray [4]. After passing through the last tracking layer they enter a calorimeter composed of bars of cesium-iodide crystals read out by PIN diodes. The calorimeter furnishes the energy measurement of the incident gamma ray. A third detector system, the anticoincidence detector (ACD), surrounds the top and sides of the tracking instrument. It consists of panels of plastic scintillator read out by wave-shifting fibers and photo-multiplier tubes and is used to veto charged cosmic-ray events such as electrons, protons or heavier nuclei.

In the LAT the tracker and calorimeter are segmented into 16 "towers," as illustrated in Fig. 1, which are covered by the ACD and a thermal blanket and meteor shield. An aluminum grid supports the detector modules and the data acquisition system and computers, which are located below the calorimeter modules. The LAT is designed to improve upon EGRET's sensitivity to astrophysical gamma-ray sources by well over a factor of 10. That is accomplished partly by sheer size, but also by use of state-of-the-art particle detection technology, such as the silicon-strip detectors [5] used in the tracker system.

Each of the 16 tracker modules is composed of a stack of 19 "trays," as can be seen Fig. 2. A tray is a stiff, lightweight carbon-composite panel with silicon-strip detectors (SSDs) bonded on both sides, with the strips on top parallel to those on the bottom. Also bonded to the bottom surface of all but

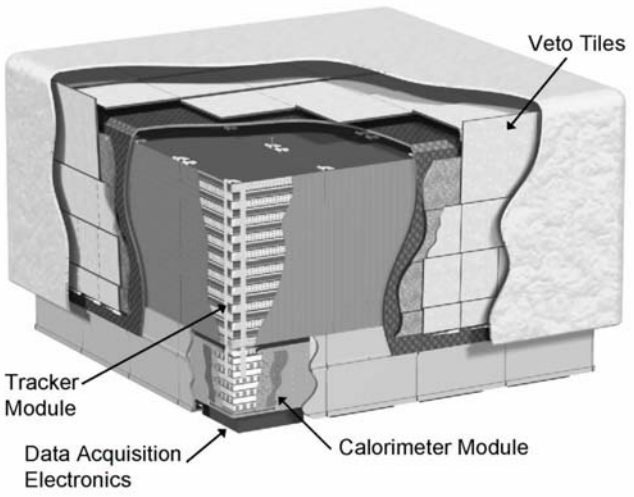

Fig. 1. Cutaway view of the LAT instrument. Each tower in the 4×4 array includes a tracker module and a calorimeter module.

Manuscript received January 26, 2006. This work was supported in part by the U.S. Department of Energy under Grant 22428-443410 and in part by the Agenzia Spaziale Italiana (ASI).

Luca Baldini, Alessandro Brez, Luca Latronico, Massimo Minuti, Carmelo Sgro', and Gloria Spandre are with the Istituto Nazionale di Fisica Nucleare and the Dipartimento di Fisica, Università di Pisa (INFN Pisa), Largo B. Pontecorvo, 3 PISA, Italy.

Thomas Himel, Wilko Kroeger, David Nelson, Mutsumi Sugizaki, and Hiro Tajima are with the Stanford Linear Accelerator Center (SLAC), Menlo Park, CA 94025.

Yohann Cohen Tanugi is currently with SLAC but worked at INFN Pisa during much of the time that the LAT Tracker was in development.

R. P. Johnson (phone: 831-459-2125, fax: 831-459-5777, email: rjohnson@scipp.ucsc.edu), H. F.-W. Sadrozinski, E. N. Spencer, and Marcus Ziegler are with the Santa Cruz Institute for Particle Physics (SCIPP), University of California, Santa Cruz, CA 95064.

Masaharu Hirayama is currently with the Joint Center for Astrophysics, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250. He worked on the LAT Tracker electronics while a member of SCIPP.

Riccardo Rando is with the Istituto Nazionale di Fisica Nucleare and the Dipartimento di Fisica, Università di Padova, I-35131 Padova, Italy.

Work supported in part by Department of Energy contract DE-AC02-76SF00515

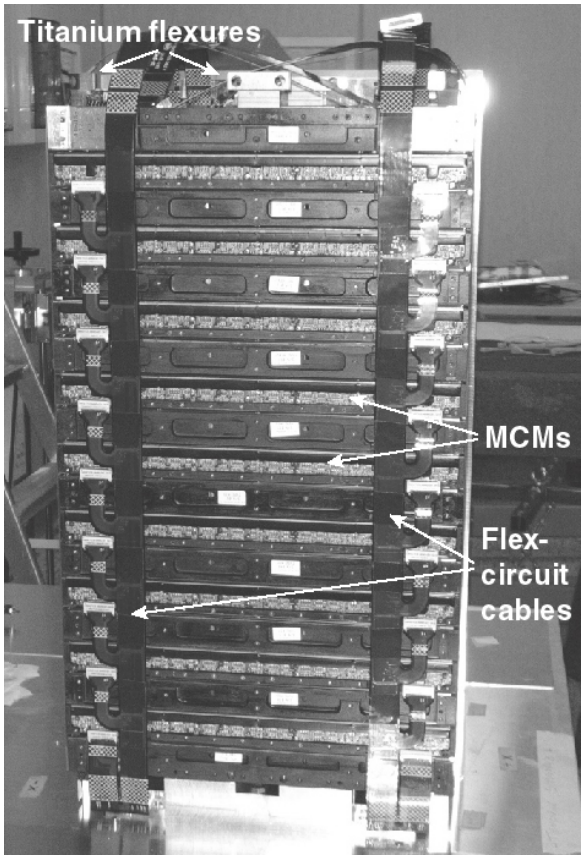

Fig. 2. Inverted view of one tracker module, with a sidewall removed. Nine MCMs and 2 flex-circuit cables are visible.

the 3 lowest trays, between the panel and the detectors, is an array of tungsten foils, one to match the active area of each detector wafer. Each tray is rotated 90° with respect to the one above or below. The detectors on the bottom of a tray combine with those on the top of the tray below to form a 90° stereo *x,y* pair with a 2 mm gap between them, and with the tungsten converter foils located just above.

The gaps and amount of material between the 16 tracker modules must be minimized to achieve optimal performance of the detector system. Therefore, the front-end electronics are mounted on the sides of the panels. A special "right-angle interconnect," described in Section IV.A, brings the signals and bias currents around the corner of the tray between the silicon strips and the amplifier-discriminator integrated circuit chips.

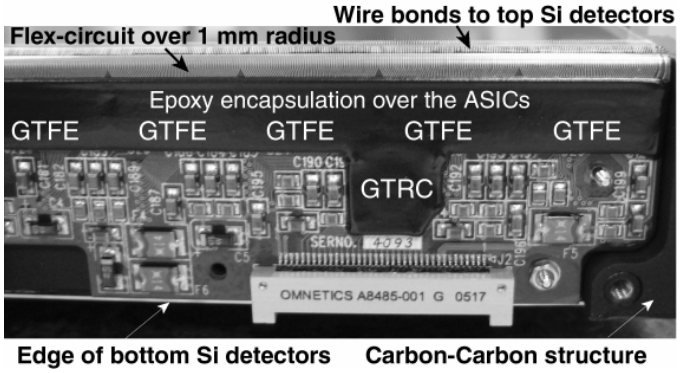

Fig. 3. View of almost ¼ of an MCM, mounted on a tray. This MCM reads out the upper layer of silicon detectors.

Each front-end electronics multi-chip module (MCM) supports the readout of 1536 silicon strips. It consists of a single printed wiring board (PWB) upon which are mounted 24 64-channel amplifier-discriminator ASICs (GTFE), two digital readout-controller ASICs (GTRC), the right-angle interconnect, bias and termination resistors, decoupling capacitors, resettable fuses, and two nano-connectors. See Fig. 3 for a photograph of one end of an MCM mounted on a tray. Each nano-connector plugs into a long flex-circuit cable, each of which interfaces 9 MCMs to the Tower Electronics Module (TEM), a custom-design data acquisition module located below the calorimeter [6]. Thus on each of the 4 sides of a tracker module one finds 9 readout boards to support 9 layers of silicon-strip detectors, which send their data to the TEM via two flex-circuit cables (see Fig. 2).

II. Requirements

The tracker electronics were designed with a goal of operating with under 200 microwatts of conditioned power per channel, in order to allow us to launch a detector with close to a million readout channels. Of course, low power has to be balanced against noise and efficiency requirements.

Achieving optimal angular resolution requires highly efficient detector layers placed as close as possible to the converter foils, because a high penalty is paid in multiple scattering if the first or second measurements after the conversion are missed in either projection. Our goal was to minimize dead regions between the SSDs (and between tracker modules) and to have an efficiency for detecting a minimum-ionizing particle of >98% within the active region of each SSD.

In contrast to EGRET, in which the tracking detector played no part in the trigger, the LAT tracker must provide the principal trigger. A practical trigger can only be formed if the noise rate from a single layer is not too high. Furthermore, the noise occupancy for a given trigger should not be too high ($<5\times10^{-5}$), or else the data volume will be prohibitive.

The readout system should have sufficient speed and buffering such that the dead time is negligible at trigger rates as high as 10 kHz.

The system should be designed to minimize the impact of single point failures. The 16 independent tracker modules already go a long ways toward achieving that goal. However, even within a single tracker module we have built in enough redundancy that in nearly all cases failure of a single component will cause a loss of no more than 64 channels out of 55,296.

III. Architecture

Fig. 4 partially illustrates the architecture of the tracker readout system, which originally evolved from experience with the BaBar Silicon Vertex Tracker readout [7]. The figure represents one of the four sides of each of the 16 tracker modules. Each module side has 9 readout boards, not all of

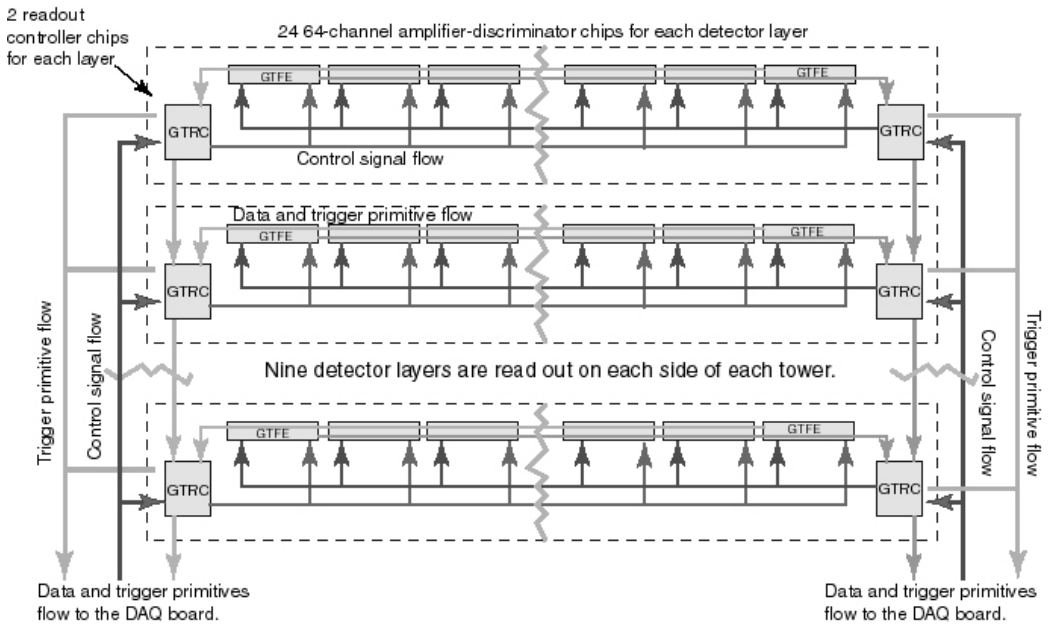

Fig. 4. Architecture of the tracker readout system, depicting one side of one tracker module. For brevity, only 3 of 9 layers are shown, and only 6 of 24 GTFE chips are shown within each layer. The arrows from GTRC to GTRC indicate the flow of data packets. The opposing flow of the readout token is not shown.

which are shown in the figure, and each board supports 24 GTFE chips, for a total of 1536 amplifier channels, and 2 GTRC chips. Each channel has a preamplifier, shaping amplifier, and discriminator similar, although not identical, to the prototype circuits described in [8]. The amplified detector signals are discriminated by a single threshold per GTFE chip; no other measurement of the signal size is made within the GTFE.

The GTFE chips are arranged on the MCM in 4 groups of 6. Each group reads out one SSD "ladder," which consists of 4 SSDs connected in series to yield strips of about 36 cm effective length.

All communication with the TEM passes through the GTRC chips, which in turn relay commands and data to and from the GTFE chips. Event data and trigger primitives flow from the GTFE chips into one or the other of the GTRC chips by passing through one GTFE chip after another. This scheme was chosen over the use of a common bus in order to avoid the possibility of a single malfunctioning chip pulling down the entire bus. Concern that in the chosen scheme a single bad chip could block the flow of data is mitigated by the left-right redundancy described below.

Each GTFE can be programmed at any time by either GTRC to send data and trigger signals to either the left or the right and to receive commands from only either the left or right GTRC (except that the command to set the direction can be received at any time from either GTRC). This architecture establishes a redundancy in the control and readout that allows the rest of the system to continue to function even in the event of loss of any single chip or readout cable. For example, if a GTFE chip in a readout board fails, then all chips to its left in the same board can be programmed to read to the left, while those to the right can be programmed to read to the right.

Each GTFE chip has two command decoders, one that listens to the left-hand GTRC, and a second that listens to the right-hand GTRC. Each GTFE also has two output data shift registers, one that moves data to the left, and a second that moves data to the right. Trigger information is formed within each GTFE chip from a logical OR of the 64 channels, of which any arbitrary set can be masked. The OR signal is passed to the left or right, depending on the setting of the chip, and combined with the OR of the neighbor, and so on down the line, until the GTRC receives a logical OR of all non-masked channels in those chips that it controls. This "layer-OR" trigger primitive initiates in the GTRC a one-shot pulse of adjustable length, which is sent down as a "trigger request" to the TEM for trigger processing. In addition, a counter in the GTRC measures the length of the layer-OR signal (time-over-threshold) and buffers the result for inclusion in the event data stream.

The usual tracker trigger is formed from a coincidence of trigger requests from 3 consecutive *x,y* pairs of tracker layers. Triggers can also be formed by the calorimeter, and when any of the 16 tracker modules triggers, a "trigger-acknowledge" signal (level-1 trigger) is sent to all 16 tracker modules. The trigger acknowledge is sent to a tracker module as a serial signal that includes a 2-bit trigger code. Upon receipt of a trigger acknowledge, each GTFE chip latches the status of all 64 channels into one of 4 internal event buffers, as specified by the 2-bit trigger code. A 64-bit mask, which is separate from the trigger mask mentioned above, can be used to mask any subset of channels from contributing data, as may be necessary in case of noisy channels. In addition to the discriminator data, the 2-bit event code is also written into the event register.

When a read-event command is sent to the GTRC chips, and relayed to the GTFE chips, the event data and trigger codes are read from the event buffer addressed by a 2-bit code

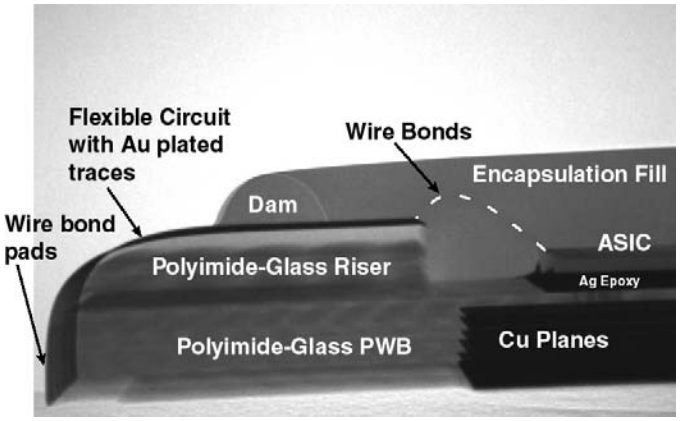

Fig. 5. X-ray cross section of the edge of the MCM with the right-angle interconnect. The pads on the flexible circuit at the left-hand edge of the photograph are for the wire bonds that go to the plane of SSDs.

in the read-event command and written into the output register. From there the data flow to one or the other of the GTRC chips. This data flow includes a partial zero suppression scheme to speed up the readout in the typically sparse events found in GLAST. Each GTFE chips sends a control bit preceding its channel data. If that bit is zero, then no channel data follow. If it is unity, then all 64 bits of channel data follow. In either case the two trigger bits are sent. These data from all GTFE chips flow serially into the GTRC, which formats a list of addresses of hit channels by counting the incoming bits. The GTRC also verifies that all of the 2-bit trigger codes match across the MCM.

Readout of the hit lists from the GTRC chips is initiated by tokens sent nearly simultaneously up the two cables. When a GTRC receives a token, it waits, if necessary, for completion of its own hit list before sending its data to the TEM and then passing the token to the layer above. Each GTRC has two event buffers, so they can begin reading a new event from the GTFE chips while the process of sending the hit lists to the TEM for the previous events is still in progress.

All communication between the TEM and GTRC chips is monitored by parity bits. No parity checks exist for the communication internal to the MCM (between GTFE chips and GTRC chips).

IV. Mechanical Integration

A. Right-Angle Interconnect

The MCMs are mounted on the edges of the trays to minimize dead space between tracker modules, which requires a method to carry 1536 detector strip signals plus 16 bias connections around the 90º corner to the SSDs. That is accomplished by a 1-layer Kapton flexible circuit that is bonded over a 1-mm radius machined into the edge of the polyimide-glass PWB. That edge of the board is roughly doubled in thickness to provide space for the radius plus additional space on the edge for wire bonding between the MCM and the SSDs. See the *x*-ray image in Fig. 5.

In the original design of the flexible circuit a cover layer or polymer mask was used to confine the plating to just the ends of the traces, where wire bonds are made. However, it proved to be difficult to position the circuit accurately enough during bonding to ensure that the stress riser at the edge of the cover or mask was not on the curve and that all of the wire bond regions were uncovered. Therefore we resorted to plating the full lengths of the traces, which caused some problems with cracking, due to the brittle nickel. The yields were acceptable as long as the plating was electrolytic (non-electrolytic plating results in cracking of nearly 100% of the traces). Nevertheless, cracked traces made up the dominant contribution to the count of dead channels in the final system.

B. Connectors and Cables

Minimizing the dead space between tracker modules also calls for very low profile connectors on the MCM. We chose 37-pin, single-row, surface mount nano-connectors with 25-mil pin spacing, manufactured by Omnetics. The connectors have an aluminum shell and use 080 jack screws. Countersunk screws hold the connectors to the board.

Two cables connect a set of 9 MCMs to the TEM. Each cable is a 4-layer Kapton flexible circuit. Two layers are used for power and ground, and the other two layers are for signal traces. One signal layer contains the busses running up the length of the cable, while the other layer holds traces to connect the busses to the 9 MCMs. All signals on the cables are low-voltage differential. Several surface-mount termination and bias resistors are soldered onto each cable. Two thermistors for monitoring the tracker temperature are also soldered onto each cable. The 9 Omnetics connectors are bonded to the cable by a film adhesive, and the surface-mount pins are soldered and then potted with epoxy. The connector at the TEM end of the cable is a 51-pin Micro-D connector with through solder pins. That end of the cable is also reinforced by bonding on a layer of fiberglass. Kapton cover layers protect the conductive traces, and conformal coating is applied over the exposed conductors on the soldered components.

V. Front-End Readout ASIC

The GTFE design achieves low power in large measure by keeping the amplifiers and digitization schemes very simple. The first stage is a folded cascode, with the input transistor bias current supplied at 1.5 V, and an output source follower. It is AC coupled to the second stage (shaping amplifier), which has only a single integration plus a source follower that is DC coupled to the discriminator. The main supply voltage is nominally 2.65 V. Good noise performance is achieved using a 1490 μm by 1.2 μm input transistor, biased at 38 μA, and a shaper output peaking time of about 1.5 μs. For the 36-cm strips (about 41 pF load) the equivalent noise charge (ENC) is about 1500 electrons, compared with a most probable signal of 32,000 electrons for a minimum-ionizing particle (MIP) passing perpendicular through the 400 μm thick silicon.

The discriminator, a simple comparator, sits very close to the amplifier output, and as a result, the system has never had

any problems with coherent noise causing the pedestal (or effective threshold) to wander, as has often been seen in systems in which the front-end chip outputs analog levels to be digitized elsewhere [9]. Since the threshold can only be adjusted per set of 64 channels, using one of the two 7-bit DACs in the GTFE, it is important to minimize the threshold variation from channel to channel. That was accomplished for the most part by the feedback system on the shaper, in which a differential amplifier stabilizes the DC output level [7], and by careful design of the comparator.

The GTFE chip has a built-in charge injection system controlled by a 64-bit calibration mask and the second DAC. Each DAC has two 6-bit linear ranges, and the 7th bit is used to select the high or low range. The mask is used to select any subset of the 64 channels for injection of charge. The calibration command causes a step voltage, set by the DAC, to be applied to each of the selected channels for a duration of 512 clock cycles.

Two other 64-bit masks control which channels contribute to the trigger and the data flow, as described in Section III. All of a chip's masks and control registers can be read back nondestructively by commands addressed to the chip.

The tracker's pipelined, buffered readout system allows the detector trigger to be active while data are being read from the tracker. For this to work properly, it is crucial for the digital readout system to operate quietly enough not to disrupt the sensitive amplifiers. That was achieved by careful attention to several design details, including the following:

- All digital communication between chips takes place by low-voltage differential signaling (LVDS), with the exception of the hard reset line and the bus used to read register contents from the GTFE chips back to the GTRC chips. Since the trigger is never active during the setup verification process, noise picked up by the amplifiers during that time does not matter. (However, the single-ended CMOS tri-state bus did cause problems of digital interference with the MCM clock bus, resulting in bit errors in the register readback process for a few chips. The problem was largely overcome by tuning the termination impedance of the clock bus, but the register readback would be more reliable if designed to use only differential signals.)
- The 20 MHz digital clock runs continuously throughout the system. Furthermore, all shift registers in the command decoders and the event readout system shift continuously, whether in use or not. Through prototype studies we found this to be crucial. If the power load in the digital part of the system changes significantly, the resulting change in the ground potential appears at the input of the amplifiers and can cause the system to trigger erroneously.
- The digital activity on the MCM is kept well separated from the analog supplies, ground, and bias points. The analog bias and filter connections never form loops around the digital busses, which are restricted to the top two layers of the 8-layer board.
- The analog and digital parts of the GTFE chips operate on separate 2.65 V supplies. Furthermore, the analog portion has its ground bus locally tied to the chip substrate throughout, while the digital return current flows on metal that ties to ground off of the chip. This scheme did not cause any problems with latch-up susceptibility (see Section VIII.H). Analog and digital sections of the chip are separated by a barrier consisting of two wells biased to the corresponding supply voltage, with a series of ground contacts in between.

Both ASICs were implemented in the Agilent 0.5-μm 3-metal CMOS process (AMOS14). The GTFE amplifier, memory, I/O drivers and receivers, and the output register layouts are full custom layouts, while the remaining digital logic and the I/O pads are composed of SCMOS standard cells from Tanner EDA, laid out by automatic place and route. All ASICs were probe tested on the wafers to ensure that only good chips were used in MCM assembly [10].

VI. Readout Controller ASIC

The GTRC buffers all command, clock, data, trigger, and reset signals between the GTFE chips and the TEM. It has two event buffers for the data, each capable of holding the addresses of up to 64 hit strips. It also has a configuration register, in which several options may be set. The register can be read back nondestructively.

When it receives a read-event command from the TEM, the GTRC executes the readout sequence to move the data from the GTFEs into one of the GTRC buffers. However, prior to taking data the GTRC configuration register must be loaded with the number of GTFE chips to be read, and those GTFE chips must be configured to send data to that GTRC.

The GTRC also includes special logic for handling the layer-OR trigger primitive generated by the GTFE chips. It has the settable option either to send the trigger directly to the TEM (after aligning it with the clock) or to send a pulse to the TEM with a length exactly equal to a settable number of clock periods. The latter choice is the normal operating mode. Furthermore, the GTRC calculates and stores the length of the layer-OR for each event, that is, each time a trigger acknowledge signal is received from the TEM. The trigger acknowledge starts the counter. Hence the count corresponds to the time-over-threshold of the largest signal in the layer, minus the round-trip time from layer-OR to trigger acknowledge.

The GTRC logic and I/O pads are composed of SCMOS standard cells from Tanner EDA, with automatic place and route, but the event memory and the I/O drivers and receivers are custom designs. The design was initially done in VHDL, from which the logic was synthesized.

VII. Multi-Chip Module (MCM)

After the flexible circuit has been bonded to the PWB and trimmed, the small surface-mount components are reflow soldered, and then the connectors are attached by screws and

hand soldered. The 26 chips are glued directly to the PWB and wedge wire bonded to gold traces on the PWB and flexible circuit. Some wire bonds also go from chip to chip. After functional testing the wires and chips are potted with epoxy (Hysol FP4450/4451 dam and fill), and then the remainder of the board is conformal coated.

The potting is unusual in that it extends over a distance of 36 cm and also covers a 0.06 cm vertical step in the board. In the early production we experienced frequent problems with delamination of the epoxy from the flexible circuit, resulting in many broken wire bonds and great concern about thermal stability. The potting material matches the coefficient of thermal expansion of the polyimide-glass board in the lateral dimensions, but not in the direction of the board thickness. However, no more delaminations occurred after a source of silicone contamination was found and removed from the process.

In addition to the left-right redundancy in control and readout, some other fault protection features are designed into the MCM. All low-voltage power flowing into the MCM passes through resettable poly-switches, which heat up and open the circuit in case of a short on the MCM. For this to function properly, we found that we had to add resistors on the cables in series with the address lines to prevent power from flowing into the MCM through the GTRC input protection diodes when the power is shut off by the polyswitches. The MCM protects against shorts in the 100V bias circuit by means of a 270 kΩ resistor placed in series with the bias current for each SSD ladder, so that an individual ladder can fail without disrupting the others. The bias circuit upstream of the resistors, however, is a potential single point of failure for a tracker module.

VIII. System Performance

Performance of the tracker readout system is discussed in the following sections, and the metrics are summarized in Table I, based on the 2nd through 17th tracker modules manufactured. All of the 17 modules are nearly identical in performance, with the exception of the first one fabricated (excluded from Table I), which has a slightly lower efficiency and higher dead-channel count.

A. Power Consumption

Based on measurements made on 16 flight tracker modules, a tracker module consumes on average 10.0 W of power while taking data at a nominal level of activity. This corresponds to only 180 μW of power per channel. Note that it includes all power used in digital communications as well as that used by the channel amplifiers.

B. Noise Performance

The equivalent noise charge (ENC) of the SSD/amplifier system has been measured channel by channel by fitting threshold curves accumulated by using the internal calibration system to inject charge while scanning the threshold. The fitted ENC varies channel by channel roughly in the range from 1200 to 1800 electrons, with a mean of around 1500 electrons. The overall normalization of the ENC (and the amplifier gain) has some uncertainties arising from the calibration of the DACs and our knowledge of the capacitance of the charge injection capacitors.

TABLE I
Tracker Performance Metrics

Metric	Measurement
Power consumption per channel	180 μW
Layer hit efficiency within active area	>99.4%
Active area fraction within a tracker module	95.5%
Overall tracker active area fraction	89.4%
Tracker noise occupancy	$<5\times10^{-7}$
Threshold variation within a chip (rms)	<9% (typically 5.2%)
Time-over-Threshold resolution for a single hit	43% FWHM
Number of dead channels	0.20%
Number of noisy channels (occupancy $>10^{-4}$)	0.06%

What is much more relevant to the operation of the detector system is the noise occupancy, which can be directly measured by generating random triggers and reading out the resulting data. The noise occupancy represents the average fraction of channels above threshold at any random snapshot in time. The typical occupancy measured at the level of a single tray is less than 10^{-6} [11]. The average noise occupancy in an integrated tracker module is 4.7×10^{-7} for a threshold setting of 1.4 fC. Note that the threshold is set per GTFE chip, so the 64 channels in a chip vary somewhat around this value. Since the most probable signal of a MIP at normal incidence is 5.1 fC, this threshold results in an expected detection efficiency within the active area that is greater than 99%.

This particular occupancy measurement was based on all layers in two representative tracker modules and is really an upper limit, since it includes contamination from real hits produced by cosmic rays. For this measurement, one tracker module had 27 "hot" strips masked out (strips with occupancy $>10^{-4}$), while the other had 25 strips masked (0.05% of strips masked overall). With no masking the average occupancy was 2.6×10^{-6}.

C. Detection Efficiency

High layer-by-layer detection efficiency is critical to optimization of the angular resolution, and hence the gamma-ray-source point-spread function, or PSF. Within a plane of 16 SSDs, the fraction of area that is active is 95.5%, taking into account the small gaps (≤0.2 mm) between SSD wafers and the dead region around the perimeter of an SSD. Including the dead area between tracker modules, the active fraction of the overall tracker (16 tower modules) is 89.4%. However, the effects of the dead fraction are greatly reduced by the fact that each tungsten converter plane is divided into 16 squares that fit directly over the SSD active areas. Furthermore, the tracking code can reconstruct the photon vertex to determine whether it lies within a dead region, in which case at least the first measurement is expected to be missing and the resolution correspondingly reduced. Therefore, there is real benefit to keeping the efficiency within the live area as high as possible.

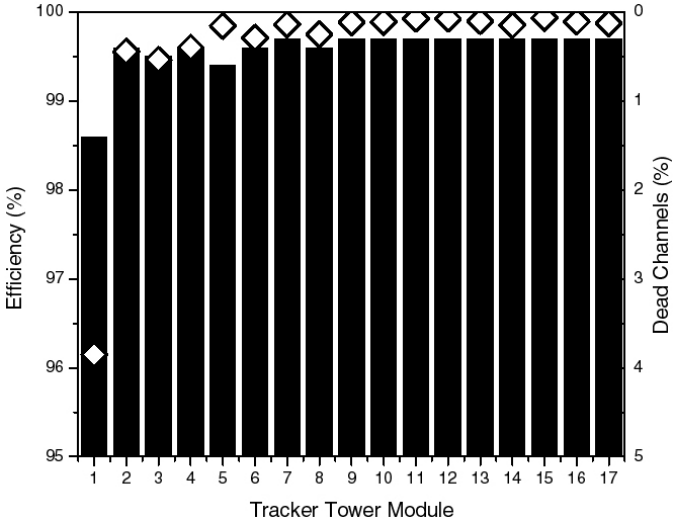

Fig. 6. The average single-layer MIP detection efficiency (bars) and percentage of dead channels (diamonds) for each tracker module.

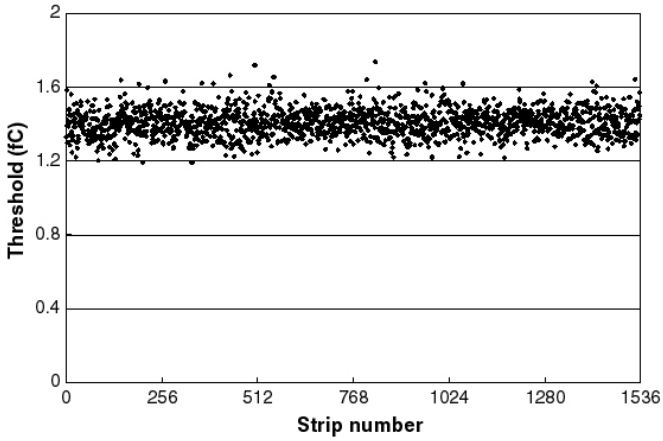

Fig. 7. Threshold measurements from a typical single MCM. The measurements were made by fitting threshold scans from an internal-charge-injection run. The injected charge corresponded to about 1.5 fC (0.29 MIPs).

Inefficiency comes from two sources: dead channels and low fluctuations in ionization, but in practice it is dominated by the former. Dead channels due to broken detector strips and to broken amplifiers number a few per ten thousand. Dead channels due to broken connections between the detector strips and the amplifiers are more common, although their number diminished greatly following experience with building the first tracker module.

The overall efficiency was measured for each layer of each tracker module using cosmic-ray tracks that pass through the active regions of the SSDs. Fig. 6 shows the average results for each of the 17 modules built. Also plotted are the percentage of dead channels. The first module had mechanical interconnect problems (including those mentioned in Section VII), resulting in a lower efficiency than was achieved in the following 16. These efficiencies were measured without any chip-to-chip tuning of the threshold, which was set at roughly 1.4 fC, the same value as used for measurements of the noise occupancy.

D. Threshold and Gain Uniformity

Since the threshold DAC has to be set per 64-channel chip, variations in effective threshold from channel to channel within a chip add directly into the noise budget. Effective threshold variations from channel to channel reflect variations in amplifier gain and variations in the pedestal set by the shaping amplifier DC feedback. Since we measure the effective threshold by use of the internal charge injection, the observed channel-to-channel variations also include variations among the 64 injection capacitors, giving an upper limit for the true threshold variation. Nevertheless, that measurement error is not expected to dominate.

Fig. 7 shows measurements of the thresholds across a single MCM. The chip-to-chip variations (rms of 3.8% over 14 tracker modules) are smaller than the variations within the 64 channels of a given chip (rms of 3.0% to 8.6%, with an average over 14 tracker modules of 5.2%), because chip-to-chip variations have been calibrated out by adjusting the threshold DACs. Note that the observed variances are not significantly affected by chip-to-chip variations (8.0% rms) of the calibration DACs, which were calibrated out by use of cosmic-ray muon data.

In all cases the rms variation of effective threshold within a chip is less than the design goal of 10%. That is good enough that for any chip there is no difficulty in setting the threshold DAC such that all 64 thresholds lie at least 4σ above the noise, but still low enough to ensure >99% efficiency for detection of MIPs.

E. Time-Over-Threshold Performance

Measurement of the time-over-threshold (TOT) of the signal is not strictly required for operation of the detector system, but it does provide information on the energy deposition in the SSDs that is useful for background rejection. For example, it can readily identify charged particles emerging from the calorimeter and ranging out in the tracker. It can also help distinguish a single background electron track from a high-energy photon conversion that results in electron and positron tracks nearly on top of each other. For simplicity and low power consumption, the tracker electronics measure the TOT only on the layer-OR trigger primitives, but that is sufficient in the low-occupancy environment of a GLAST gamma-ray event.

Fig. 8 shows a histogram of the single-layer TOT measured from high-energy, minimum-ionizing cosmic-ray muons. The preamplifier limited dynamic range causes the TOT to saturate at 150 μs (about 200 MIPs or 1000 fC), but the digitization in the GTRC truncates the measurement at 50 μs (about 6 MIPs or 31 fC), just to avoid delays in the event readout. Therefore, the TOT is not useful for studies of heavy ions, but it does have enough range and resolution to help with background rejection.

F. Readout Speed

The digital readout of the tracker system works as designed. Two levels of event buffering (4 buffers in the GTFE chips and 2 buffers in the GTRC chips) ensure that the dead time is

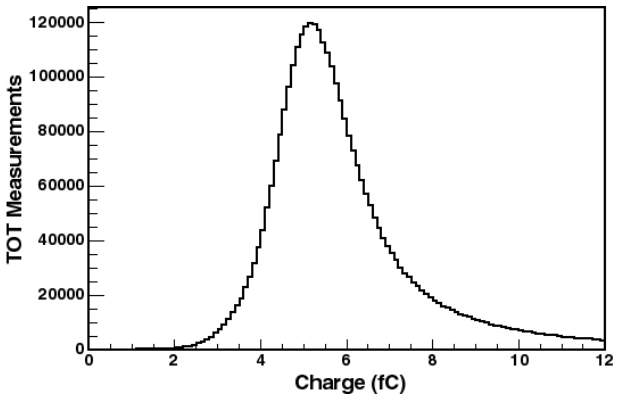

Fig. 8. Single-layer charge deposition from cosmic-ray muons, measured by time-over-threshold. The measurements were calibrated chip by chip to remove instrumental variations.

negligible until the data transmission bandwidth saturates. There is no significant increase in noise observed when the trigger is active during readout of previous events.

The ground level cosmic-ray flux is too low to test the speed of the system, so a test was conducted with an intense ^{241}Am *x*-ray source placed on top of the tracker. The occupancy of *x*-ray hits was high enough to produce a high rate of accidental triggers at various rates determined by placement of absorbers between the source and tracker. The readout rate was limited to 6.5 kHz by the writing of data to disk. With the trigger rate tuned to 6.5 kHz the tracker made no measurable contribution to the dead time, and the distribution of time between triggers was perfectly Poisson. To fully exercise the tracker buffering, other runs were successfully taken with instantaneous tracker trigger rates as high as 54 kHz, but with dead time imposed by the system elements downstream of the tracker readout.

G. Electromagnetic Interference

One of the tracker modules was put through the NASA specified qualification tests for electromagnetic interference and susceptibility (EMI/EMC). To limit emissions, the tracker carbon-composite sidewalls are coated with aluminum foil on each side. Furthermore, the top of the tracker module and all cracks on the corners are covered with aluminum tape. The bottom of the tracker module, however, cannot be covered, due to all of the interface hardware, cables, and ventilation holes in that region. The testing included conducted susceptibility and emissions, concerning the power cables connecting the outside world to the tracker plus the TEM and the tower power supply. It also included radiated susceptibility and radiative emissions from 20 Hz to 50 kHz magnetic and 10 kHz to 18 GHz electric.

Susceptibility was checked by measuring the noise occupancy while applying a conductive or radiative source. With radiated fields of 20 V/m the tracker passed with an occupancy of less than 10^{-4} in the region of the amplifier bandpass (around 100 kHz) as well as at all other test frequencies.

Emissions were measured while the tracker ran repeated readout loops. The electric field requirements are shown in Fig. 9 and were satisfied by the tracker except in the region 20 MHz to 1.5 GHz, where many narrow harmonics of the 20 MHz system clock exceeded the limits by up to 20 dB. Those were the only nonconformances found in all of the EMI/EMC testing, and they do not pose any risk to the mission. Note that the emission requirements were fully satisfied throughout the notch regions where the spacecraft GPS receiver and the S-band radio operate.

H. Radiation Hardness

Radiation hardness of the ASICs was verified by testing with ionizing radiation from a ^{60}Co source and by testing for single-event effects (SEE) in a heavy ion beam at the INFN SIRAD facility [12]. The SEE results were also crosschecked in a cyclotron beam at Texas A&M University (TAMU) with 4 times greater ion range [13], giving nearly identical results. A full report on these measurements is found in [14].

The effects of ionizing radiation were measured up to a dose of 10 kRad, more than 10 times the expected dose over a 5-year mission. That level of radiation was found not to have a significant effect on any aspect of the performance of the ASICs. The main effect of the radiation on the detector system will be increasing leakage current in the SSDs. The integration time of the amplifiers is short enough that this expected increase will have only a minor effect on the noise budget at end of life, even at the upper limit of the operating temperature range (35ºC).

The SIRAD testing for single-event effects yielded linear energy transfers from 8.5 to 82 MeV-cm^2/mg and range from 62 to 23 micrometers. No latch-up was ever observed in any of the experiments (SIRAD or TAMU), and the probability of encountering a latch-up in one of the 14,976 tracker chips during the 5-year mission is less than 1/2000 at 95% confidence level.

Single-event upset (SEU) is only an issue for the configuration registers in the chip. A rare upset in the event memory would be of no consequence, for example, as it would add negligibly to noise that is already present. The configuration register cells are specially designed to be resistant to upset [15], which reduces the SEU cross section by a factor of about 0.004 (our measurement). In summary, the expected number of upsets for a 5-year mission is 0.7 in the GTFE chips and 0.005 in the GTRC chips (the latter

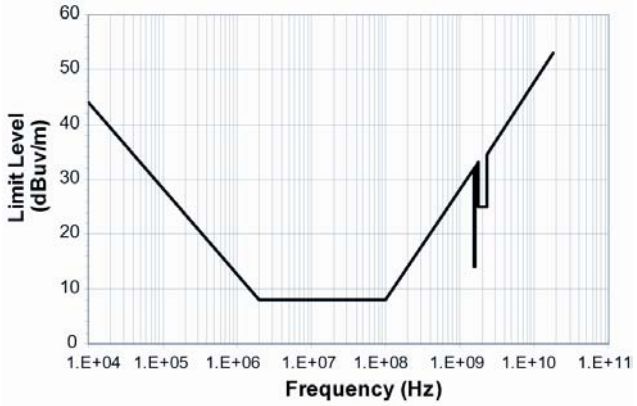

Fig. 9. Upper limits allowed for electric field emissions at a distance of 1 m.

number is so much smaller because there are 12 times fewer GTRC chips, and each GTRC has fewer register cells than the GTFE has). These rates are negligible, especially since the configuration registers will be routinely reloaded during the mission.

IX. Conclusion

With all of the tracker modules built and fully tested, the GLAST LAT tracker readout electronics have been demonstrated to meet all of the design goals. In particular, the detector system has been demonstrated to detect minimum ionizing particles with hit efficiencies >99% and with noise low enough such that the tracker can provide the primary trigger for the LAT instrument. Furthermore, that is accomplished with power consumption low enough (160 W) to allow the 880,000 channel instrument to operate continuously in space.

Acknowledgment

We would like to thank Thomas Borden, Richard Gobin, Albert Nguyen, David Rich, Jeff Tice, Roger Williams, and Charles Young of SLAC, Tim Graves of Sonoma State University, and Kamal Prasad for their dedicated work supporting the assembly and testing of the LAT tracker readout electronics boards. We thank Dieter Freytag, Gunther Haller, and Jeff Olsen of SLAC and Sergei Kashiguine of SCIPP for their contributions to the electronics design, and we thank the LAT Integration and Test group for measurements made with the ^{241}Am source. R.P. Johnson thanks William Atwood of SCIPP for many useful discussions and analysis support during the electronics development.

References

[1] W.B. Atwood, "GLAST: applying silicon strip detector technology to the detection of gamma rays in space," *Nucl. Instrum. Meth..A*, vol. 342, p. 302, 1994.

[2] N. Gehrels and P. Michelson, "GLAST: the next-generation high energy gamma-ray astronomy mission," *Astropart. Phys*., vol 11, p. 277, 1999.

[3] D.J. Thompson, *et al.* "Calibration of the Energetic Gamma-Ray Experiment (EGRET) for the Compton Gamma-ray Observatory," *ApJ Suppl*., vol. 86, pp. 629–656, June 1993.

[4] W.B. Atwood, *et al*., "The silicon tracker/converter for the Gamma-ray Large Area Space Telescope," *Nucl. Instrum. Meth. A*, vol. 435, p. 224, 1999.

[5] T. Ohsugi, *et al*., "Design and properties of the GLAST flight silicon micro-strip sensors," *Nucl. Instrum. Meth. A*, vol. 541, p. 29, 2005.

[6] Leonid Sapozhnikov, "Tower Electronics Module (TEM) Specification and ICD," Stanford Linear Accelerator Center GLAST/LAT internal document LAT-SS-00288, 2004.

[7] I. Kipnis, *et al.,* "A Time-Over-Threshold Machine: the readout integrated circuit for the BaBar Silicon Vertex Tracker," *IEEE Trans. Nucl. Sci.,* vol 44, p. 289, 1997.

[8] R.P. Johnson, P. Poplevin, H.F.-W. Sadrozinski, and E.N. Spencer, "An amplifier-discriminator chip for the GLAST silicon-strip tracker," *IEEE Trans. Nucl. Sci.*, vol. 45, pp. 927–930, June 1998.

[9] G. Batignani *et al.,* "Recent results and running experience of the new Aleph vertex detector," *IEEE Trans. Nucl. Sci,* vol 39, no. 4, pp. 701–706, Aug. 1992.

[10] Mutsumi Sugizaki, "Production and test of front-end electronics for the GLAST LAT silicon tracker," *Nucl. Instrum. Meth. A*, vol. 541, pp. 304–309, 2005.

[11] Luca Baldini, "The silicon strip tracker for the GLAST experiment," Ph.D. dissertation, Dept. Physics, University of Pisa, Pisa Italy, 2005.

[12] J. Wyss, D. Bisello, and D. Pantano, "SIRAD: an irradiation facility at the LNL tandem accelerator…", *Nucl. Instrum. Methods A*, vol. 462, no. 3, pp. 426–434, Aug. 2001.

[13] Texas A&M Univ. Cyclotron Institute, MS 3366, 77843-3366, College Station, TX, USA.

[14] R. Rando *et al.*, "Radiation testing of GLAST LAT tracker ASICs," IEEE Trans. Nucl. Sci, vol. 51, p 1067, 2004.

[15] L.R. Rockett, "An SEU-hardened CMOS data latch design," *IEEE Trans. Nucl. Sci.*, vol. 35, pp. 1682–1687, Dec. 1988.