The Silicon Strip Tracker of the *Fermi* Large Area Telescope: the first five years in orbit

Luca Baldini[*]

Università di Pisa and INFN-Sezione di Pisa
E-mail: luca.baldini@pi.infn.it

on behalf of the *Fermi*-LAT collaboration

The Large Area Telescope (LAT) is the main instrument onboard the *Fermi* Gamma-ray Space Telescope, a space observatory launched in low-Earth orbit on June 11 2008 to survey the high-energy gamma-ray sky. The LAT tracker/converter serves the twofold purpose of converting the incoming gamma-ray into an electron-positron pair and tracking the latter in order to measure the original photon direction. With its 73 m^2 of single-sided silicon-strip detectors, read out by almost 900,000 independent electronics channel, it is the largest solid-state tracker ever built for a space application. The tracker system operates on 160 W of conditioned power while achieving a single-plane hit efficiency in excess of 99% within the active area and a noise occupancy at the level of $\sim$ 1 channel per million.

We describe the basic tracker design and its performance throughout the first five years of operation in space.

---

[*]Speaker.

---

1. Introduction and mission status

Designed to survey the gamma-ray sky in the broad energy range from 20 MeV to more than 300 GeV, with the additional capability of studying transient phenomena at lower energies and charged species—notably leptons—from GeV to TeV energies, the *Fermi* gamma-ray Space Telescope is effectively the reference operating space-borne gamma-ray observatory of the current decade.

Fermi carries two instruments on board: the Large Area Telescope (LAT), a pair conversion telescope for photons above ~ 20 MeV, and the Gamma-ray Burst Monitor (GBM), designed to observe the full unoccluded sky between ~ 8 keV and ~ 40 MeV for the study of transient sources, particularly Gamma-Ray Bursts. The two instruments operate in synergy for a continuous monitor of the entire gamma-ray sky, achieving a fairly uniform exposure on time scales as short as a few hours.

More than five years into the mission, the Fermi LAT has been collecting science data at a steady rate of ~ 60 billion events read out per year (figure 1). About 12 billion events per year are actually down-linked to ground for further processing and, out of these, ~ 170 million gamma-ray candidates per year are distributed to the community within hours from the trigger time. The average data taking up-time, so far, amounts to a remarkable 99.7%[1], detector calibrations and a few minor hardware issues accounting for the remaining 0.3%.
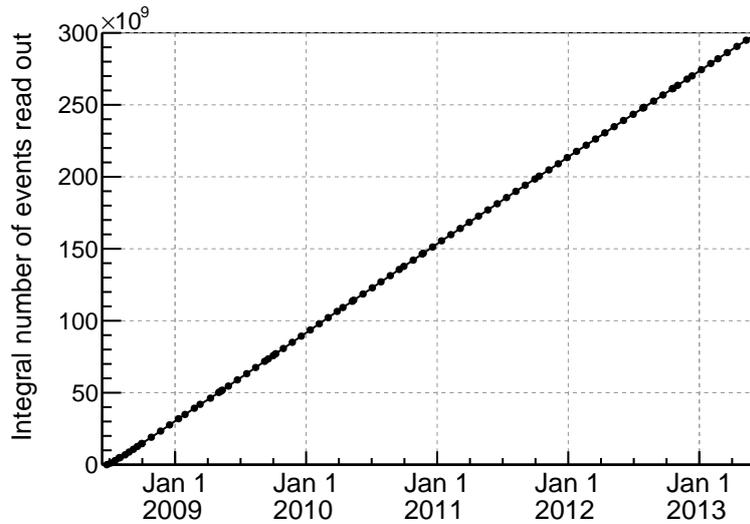

Figure 1: Integral number of events readout by the LAT, as a function of time, through the first five years of science operation.

Fermi completed the first five years of nominal science operation on August 4, 2013. This date marks the end of the so-called *prime phase* of the mission, after which successive extensions are negotiated with NASA every two years in a competitive process called *Senior Review*. *Fermi* underwent its first Senior Review at the beginning of 2012 and the decision was made to extend

---

[1]This figure does not take into account the time spent in the South Atlantic Anomaly ($\sim 14\%$ of the total) where the LAT does not acquire science data.

the mission for two years (i.e., through 2014), with the additional recommendation of a further extension, to be re-evaluated in 2014. The current baseline, therefore, is to operate through 2016 and the actual goal is still the original one of a 10-year mission.

2. The *Fermi* Large Area Telescope

The Large Area Telescope is designed to detect photons in the pair-conversion regime (i.e., above ~ 20 MeV). It is composed of three basic subsystems: a tracker/converter, where position sensitive detector planes are interleaved with tungsten radiators to promote photon conversions; a 8.6 X_0 calorimeter for the measurement of the photon energy; and an anti-coincidence detector to suppress the charged-particle background. The tracker and calorimeter subsystems are implemented as a 4×4 matrix of 16 identical modules called *towers*.
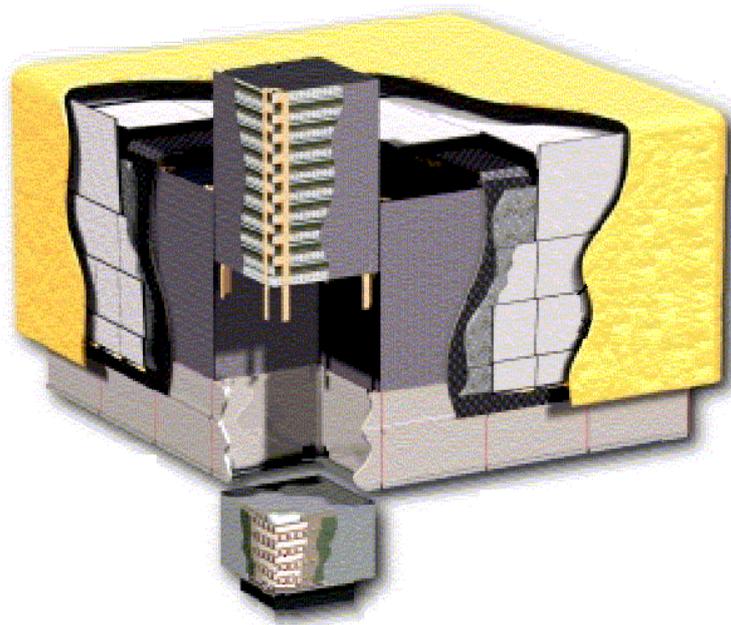

Figure 2: Schematic view of the Large Area Telescope on board *Fermi*. One of the corner tower modules is broken up into its tracker (top) and calorimeter (bottom) sections for illustrative purposes.

Figure 2 shows a schematic view of the LAT. The reader is referred to [1] for a more in-depth description of the instrument, to [2] for a detailed overview of its performance and to [3] for a summary of the on-orbit calibration.

3. The LAT silicon tracker

The primary function of the LAT tracker [4] is to convert the gamma ray into an electron-positron pair and determine its direction through the reconstruction of the tracks of the charged daughters. The Large Area telescope is a wide-field-of-view (~ 2.4 sr, i.e. with sensitivity up to $\sim 70°$ off-axis) detector designed to operate over nearly 5 decades in energy. Many of the

challenges connected with the design of the tracking stage come in fact from the need of a good tracking efficiency and pointing resolution over such a wide instrumental phase space. At low energy the angular resolution is dominated by multiple scattering, which in principle would call for thin converters. However, at the same time, the need of material to convert gamma rays effectively sets a trade off between the pointing resolution and the detection efficiency. At high energy (i.e., above ~ 10 GeV), on the other end, the angular resolution is mainly determined by the hit resolution and the lever arm between adjacent detection planes. This would ideally call for a fine segmentation of the position-sensitive detectors, but the power budget (160 W overall) poses strong constraints in this sense.
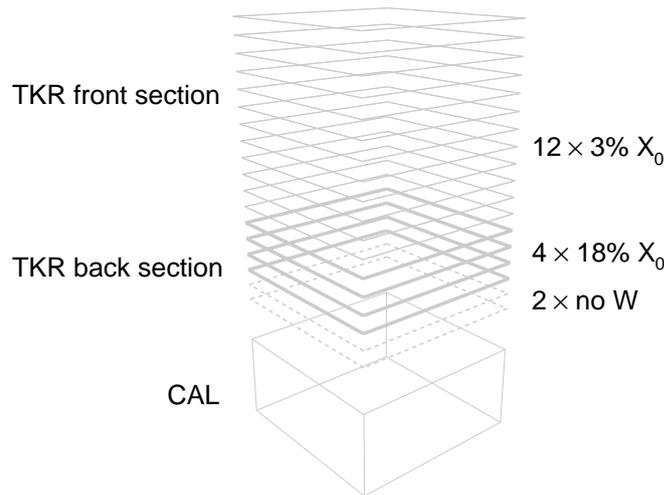

Figure 3: Basic layout of a LAT tower module. The thickness of the tungsten radiators for the front and back sections of the tracker are indicated.

The overall tracker design reflects these basics considerations. A tracker tower module consists of 19 tray structures supporting the 36 single-sided silicon planes (18 *x-y* layers) and the tungsten conversion foils. The tracker divides into two distinct sections, with very different performance: the first 12 conversion layers (constituting the so called *front* section, optimized for pointing resolution) are 3% X_0 each, while the subsequent 4 (or *back* section) are 18% X_0 thick. Since the basic primitive that the tracker contributes to the overall instrument trigger requires a sixfold coincidence of three *x-y* layers in a row, the last two bi-layers have no radiators.

One might naively think that the presence of the tungsten converters make the entire issue of the material budget essentially trivial. This is actually not the case, as simulations clearly indicate that missing the very first hit after the photon conversion can result in a worsening of up to a factor of two in the low-energy angular resolution. It is therefore highly desirable to promote the gamma-ray conversion in specific regions of the detector, immediately followed by the tracking detectors, and care was taken in the design phase to minimize the amount of all the rest of passive material.

3.1 The silicon strip detectors

The design of the LAT silicon strip detectors (SSD) is far from being aggressive by any high-

energy physics standard. The basic characteristics of the sensors are listed in Table 1: made from $6''$, 400 μm thick wafers, the SSDs are single-sided, AC-coupled and with a strip pitch of 228 μm. We stress that, unlike magnetic spectrometers such as AMS, where the hit resolution in the bending plane is one of the basic figures of merit determining the maximum detectable rigidity, the LAT is multiple-scattering dominated over most of the phase space due to the presence of the tungsten converters. Therefore, even if the power budget allowed for it, the gain in going to a finer pitch strip would be marginal for most of the science analyses.

Coupling	AC
Outer size	8.95×8.95 cm^2
Strip pitch	228 μm
Thickness	400 μm
Depletion voltage	< 120 V
Leakage current	~ 1 nA/cm^2 at 150 V
Breakdown voltage	> 175 V
Bad channels	$\sim 10^{-4}$ (of 900k)
Number of SSD tested	12500
Number of single strip tests	≈ 30M
Rejected SSDs	0.6%

Table 1: Summary table of the basic characteristics of the LAT silicon strip detectors. We refer the reader to [6] for more details. We note explicitly that defective SSD strips constitute a sub-dominant contribution to the actual inventory of bad channels, as we shall see in the following.

On a related note, the table makes it clear the excellent quality of the sensor production, which is reflected in the overall fraction of bad channels and in the sensor rejection rate—and is key to achieve a uniform, full detection efficiency.

3.2 Mechanical integration

From the standpoint of the mechanical integration, silicon strip detectors are glued head-to-head and wire bonded to form ~ 36 cm-long ladders. Four ladders are then integrated into a $\sim 36 \times 36$ cm^2 detection plane. As mentioned before, composite tray structures provide the mechanical framework and housing for both the silicon detectors and the tungsten converters.

The mechanical design is deliberately tailored at minimizing the inactive regions within the detector volume. Adjacent trays within a tower module are separated by less than 2 mm, so that each radiator is immediately followed by a silicon bi-layer (measuring both coordinates). The readout electronics is mounted on the tray sides, with the silicon strips connected to the front-end chips via a 90° pitch adapter and all the communication with the tower electronic module happening via flat cables. Finally, the separation between adjacent tower modules is of the order of 2 mm. Conjugating these design decision with the necessity of a a stiff enough structure to sustain the vibration loads at launch, the possible temperature gradients and operation in vacuum was highly non trivial. More details of the tracker environmental tests are given in [5].

3.3 The tracker electronics system

From the standpoint of the electronics system [7, 8], each silicon layer is handled by 24 front-end chips (each one serving 64 strips, for a total of 1536 strips per layer) and two readout controller chips, acting as an interface to the tower-level electronics. Each front-end chip can be configured to read commands and write data to either of the readout controllers, so that a defective chip can be effectively by-passed without loosing the entire layer. Two multi-layer flat cables connecting each readout controller to the corresponding tower electronics module complete the redundancy.

The basic tracker readout is digital, i.e., the signal from each preamplifier is discriminated and only the identifiers of the strips above threshold at the time of readout are included in the data stream, with the zero suppression taking place in the controller chips. In addition to that, the time over threshold of the logical OR of all the strip signals read by each controller is recorded and provide some indication of the amount of ionization in the event that can be effectively used at the background rejection stage.

The key features of the tracker electronics system are the low power consumption—160 W for the full tracker, or ~ 180 μW per channel—and the low noise occupancy—at the level of 10^{-7}). In addition to that the tracker provides the main trigger to the detector, in the form of a sixfold coincidence of three *x-y* planes in a row.

4. On-orbit tracker performance

In this section we shall briefly review the main outputs of the tracker performance monitoring applications. The input data include both dedicated calibration runs (e.g., tests using the internal charge injection system) and nominal science data taking runs (e.g., selections of minimum ionizing tracks for efficiency and alignment measurements).

4.1 Noise performance

The noise performance is periodically monitored through dedicated calibration runs using the tracker internal charge injection system. The effective hit threshold and the equivalent noise charge (ENC) are measured, on a channel-by-channel basis, by means of charge scan (typically 0.8–2.6 fC in steps of 0.07 fC) at the nominal discriminator threshold. For each charge injection setting the channel occupancy is recorded and the corresponding s-curve is constructed and fitted to an error function.

Figure 4 shows the time evolution of the ENC, averaged across all the LAT channels, through the first five years of the mission. The typical value of ~ 1300 electrons is mainly determined by the ~ 41 pF silicons-strip capacitance at the input of the preamplifier and corresponds to about 4% of the charge created by a MIP traversing 400 μm of silicon (the nominal threshold of 0.25 MIPs corresponds to more than 6σ over the noise floor). The slight increase of the average noise (at the level of $\sim 2\%$) through the prime phase of the mission is qualitatively in agreement with the effect of the increase of the leakage current due to radiation damage and translates into a negligible effect even when projected to the mission goal of ten years.

4.2 Defective channels

The basic inventory of defective channels roughly divides into four main groups:
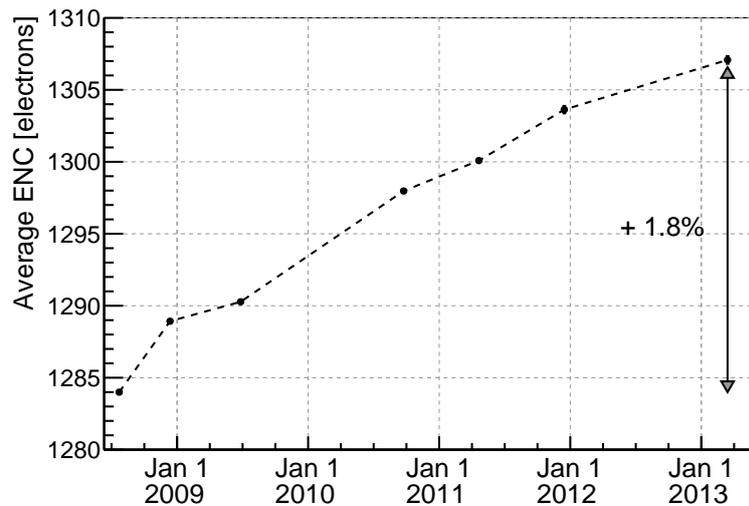
Figure 4: Time evolution of the average noise (across all the LAT channels) through the first five years of science operation. The slight increase is compatible with the increase in the leakage current due to radiation damage in the silicon detectors.

1. Noisy channels, i.e. channels whose noise occupancy is much larger than the typical value of $\sim 10^{-7}$. They are routinely monitored on a run by run basis and channels with an occupancy in excess of 1% are masked to both trigger and data.

2. Dead channels, i.e. channels whose preamplifier is dead. They are routinely monitored through the charge injection scan described in the previous section.

3. Disconnected channels, i.e. channels for which the silicon strip is not physically connected to the preamplifier input. In the vast majority of the cases this is due to manufacturing issues with the pitch adapter. The basic phenomenology of these channels is that they have a much lower (~ 250 electrons) ENC and no physical hits associated in nominal science data taking runs.

4. Partially disconnected channels, i.e. channels where one or more of the wire bonds along the ladder is defective. These channels feature intermediate noise levels and hit-map populations.

Figure 5 summarize the time evolution of the number of defective channels, broken up in the four basic types. The two most noticeable features are a slight increase at the beginning of the mission and one starting from the beginning of 2011. The first is due to some fabrication issues (mainly with the pitch adapters) that we encountered at the beginning of the flight production and is limited to the first flight module assembled. The second is due to a long-term noise increase that we observed in one of the ladders and for which we could not identify a root cause.

Overall, the total number of defective channels increased from 3661 to 4045 through the prime phase of the mission. When normalized to the total number of channels (884736) in the LAT tracker, these figures translate into 0.41% and 0.46%, respectively. Even if we lost the entire
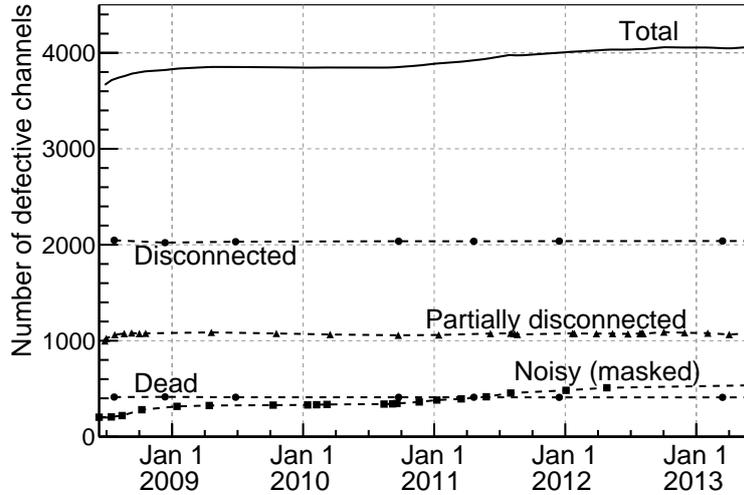
Figure 5: Time history of the number of defective channels in the LAT, grouped by defect type. For the dead and disconnected channels each data point corresponds to a charge injection run. For the noisy channels each data point corresponds to the upload of a new mask on the LAT. The number of partially disconnected channels is regularly measured by a dedicated analysis of standard science data taking runs.

noise ladder mentioned before, the corresponding average loss of efficiency for the LAT would be essentially negligible.

4.3 Hit and trigger efficiency

We use on-orbit non interacting, minimum ionizing proton tracks to monitor the hit and trigger efficiency on a run-by-run basis for each single silicon layers. This is achieved by extrapolating the reconstructed track direction through the detector searching for missing hits in the active silicon area[2] (for the trigger efficiency measurements, since the trigger is tower-based, events involving more than one tower modules are used).

Figure 6 shows the average LAT hit efficiency as a function of time through the first five years of on-orbit operation. The figure of merit averages at $\sim 99.7\%$ (for reference, the *Fermi* science requirement document called for 98%). The slight decreasing trend traces the small increase in the number of noisy channels that have been masked in the course of the mission. We note that the fraction of defective channels is actually an upper limit to the hit inefficiency, as for disconnected and partially disconnected channels the signal tends to couple capacitively to the neighbor strips providing a mechanism for a partial recovery of hits that would otherwise be lost.

4.4 Alignment

MIP tracks are also used to continuously monitor the alignment of the tracker. This breaks up in two main parts: the *inter-tower* alignment, where we measure six parameters (three shifts and three rotations) for each of the 16 tower modules, and the *intra-tower* alignment, where we

---

[2]For completeness: the efficiency figures quoted here refer to the *active* silicon area. When viewed from the top, the LAT tracker feature a fractional dead area of the order of $\sim 11\%$.
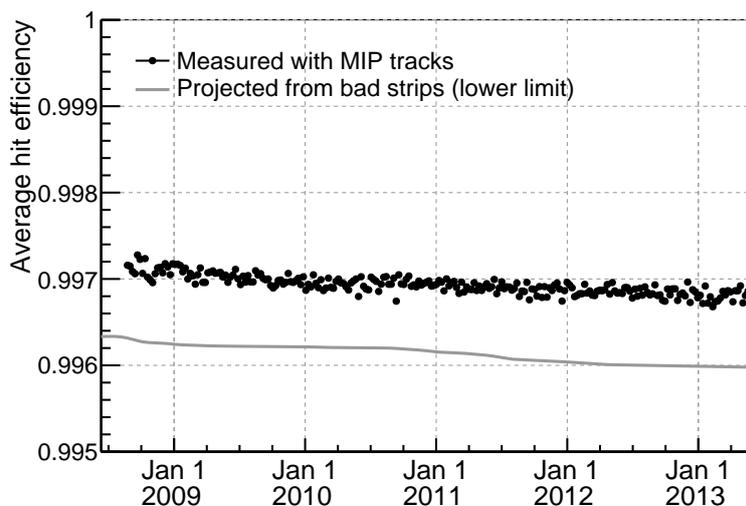
Figure 6: Hit efficiency, averaged over the active area of the LAT tracker. Each data point corresponds to the time average over a week. The gray line represents the level of efficiency that one expects, based on the total number of defective channels.

measure the same type of constants for each of the silicon planes within each module. In both cases we have no evidence of significant changes, neither due to vibrations at launch, nor during on-orbit operations.

5. Conclusions

The *Fermi* LAT silicon tracker has now operated for more than five years in orbit, showing no significant sign of degradation of performance. All the design goals were met with large margin—particularly an average hit efficiency well in excess of 99% at a corresponding noise occupancy of the order of 10^{-7} (i.e., less than one noise hit per event for the full LAT).

The tracker has served beautifully the science of the prime phase of the mission and, while the current baseline is to operate at least through 2016, there is in principle no reason—at least from the standpoint of the detector performance—why the original goal of a ten-year mission should not be met.

Acknowledgments

References

[1] W. B. Atwood et al., *The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission*, *The Astrophysical Journal* **697** (2009) 1071–1102 [`arXiv:0902.1089`].

[2] M. Ackermann et al., *The Fermi Large Area Telescope on Orbit: Event Classification, Instrument Response Functions, and Calibration*, *The Astrophysical Journal Supplement Series* **203** (2012) 4 [`arXiv1206.1896`].

[3] A. A. Abdo et al., *The on-orbit calibration of the Fermi Large Area Telescope*, *Astroparticle Physics* **32** (2009) 193–219 [`arXiv0904.2226`].

[4] W. B. Atwood et al., *Design and initial tests of the Tracker-converter of the Gamma-ray Large Area Space Telescope*, *Astroparticle Physics* **28** (2007) 422-434.

[5] R. Bagagli et al., *Environmental tests of the flight GLAST LAT tracker towers*, *Nuclear Instruments and Methods in Physics Research A* **584** (2008) 353–373.

[6] T. Ohsugi et al., *Design and properties of the GLAST flight silicon micro-strip sensors*, *Nuclear Instruments and Methods in Physics Research A* **541** (2005) 29–39.

[7] L. Baldini et al., *The Silicon Tracker Readout Electronics of the Gamma-Ray Large Area Space Telescope*, *IEEE Transactions on Nuclear Science* **53** (2006) 466–473.

[8] L. Baldini et al., *Fabrication of the GLAST Silicon Tracker Readout Electronics*, *IEEE Transactions on Nuclear Science* **53** (2006) 3013–3020.

PoS(Vertex 2013)039