



HAL
open science

Fermi Large Area Telescope Performance after 10 Years of Operation

M. Ajello, W. Atwood, M. Axelsson, R. Bagagli, M. Bagni, L. Baldini, D. Bastieri, F. Bellardi, R. Bellazzini, E. Bissaldi, et al.

► To cite this version:

M. Ajello, W. Atwood, M. Axelsson, R. Bagagli, M. Bagni, et al.. Fermi Large Area Telescope Performance after 10 Years of Operation. *The Astrophysical Journal Supplement*, 2021, 256 (1), pp.12. <10.3847/1538-4365/ac0ceb>. <insu-03349477>

HAL Id: insu-03349477

<https://insu.hal.science/insu-03349477v1>

Submitted on 27 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License



Fermi Large Area Telescope Performance after 10 Years of Operation

M. Ajello¹, W. B. Atwood², M. Axelsson^{3,4}, R. Bagagli^{5,7,6}, M. Bagni^{5,7,6}, L. Baldini⁵, D. Bastieri^{6,7}, F. Bellardi^{5,7,6}, R. Bellazzini⁸, E. Bissaldi^{9,10}, E. D. Bloom¹¹, R. Bonino^{12,13}, J. Bregeon¹⁴, A. Brez⁸, P. Bruel¹⁵, R. Buehler¹⁶, S. Buson¹⁷, R. A. Cameron¹¹, P. A. Caraveo¹⁸, E. Cavazzuti¹⁹, M. Ceccanti⁵, S. Chen^{6,20}, C. C. Cheung²¹, S. Ciprini^{22,23}, I. Cognard^{24,25}, J. Cohen-Tanugi²⁶, S. Cutini²⁷, F. D'Ammando²⁸, P. de la Torre Luque⁹, F. de Palma^{29,30}, S. W. Digel¹¹, F. Dirrsa³¹, N. Di Lalla¹¹, L. Di Venere^{9,10}, A. Domínguez³², D. Fabiani^{5,7,7}, E. C. Ferrara^{33,34}, A. Fiori³⁵, G. Foglia^{5,7,6}, Y. Fukazawa³⁶, P. Fusco^{9,10}, F. Gargano¹⁰, D. Gasparri^{22,23}, M. Giroletti²⁸, T. Glanzman¹¹, D. Green³⁷, S. Griffin³³, M.-H. Grondin³⁸, J. E. Grove²¹, L. Guillemot^{24,25}, S. Guiriec^{33,39}, M. Gustafsson⁴⁰, E. Hays³³, D. Horan¹⁵, G. Jóhannesson^{41,42}, T. J. Johnson⁴³, T. Kamae⁴⁴, M. Kerr²¹, M. Kuss⁸, S. Larsson^{4,45,46}, L. Latronico¹², M. Lemoine-Goumard³⁸, J. Li¹⁶, I. Liodakis⁴⁷, F. Longo^{48,49}, F. Loparco^{9,10}, M. N. Lovellette²¹, P. Lubrano²⁷, S. Maldera¹², A. Manfreda⁵, G. Martí-Devesa⁵⁰, M. N. Mazziotta¹⁰, N. Menon^{5,7,6}, I. Mereu^{27,51}, M. Meyer⁵², P. F. Michelson¹¹, M. Minuti⁵, W. Mitthumsiri⁵³, T. Mizuno⁵⁴, M. Mongelli⁵⁵, M. E. Monzani¹¹, I. V. Moskalenko¹¹, M. Negro^{56,57}, E. Nuss²⁶, R. Ojha³³, M. Orienti²⁸, E. Orlando^{11,58}, A. Paccagnella⁵⁹, V. S. Paliya¹⁶, D. Paneque³⁷, Z. Pei⁷, J. S. Perkins³³, M. Pesce-Rollins⁸, M. Pinchera⁵, F. Piron²⁶, H. Poon³⁶, T. A. Porter¹¹, R. Primavera²³, G. Principe^{28,48,49}, J. L. Racusin³³, S. Raino^{9,10}, R. Rando^{6,20,60}, B. Rani^{33,61,62}, E. Rapposelli^{5,7,6}, M. Razzano^{8,78}, S. Razzaque⁶³, A. Reimer^{11,50}, O. Reimer⁵⁰, J. J. Russell¹¹, N. Saggini^{5,7,6}, P. M. Saz Parkinson^{2,64,65}, N. Scolieri⁶⁶, D. Serini⁹, C. Sgrò⁸, E. J. Siskind⁶⁷, D. A. Smith^{38,68}, G. Spandre⁸, P. Spinelli^{9,10}, D. J. Suson⁶⁹, H. Tajima^{11,70}, J. G. Thayer¹¹, D. J. Thompson³³, L. Tibaldo⁷¹, D. F. Torres^{72,73,74}, G. Tosti^{27,51}, J. Valverde^{33,57}, L. Vigiani^{5,7,6}, and G. Zaharijas^{58,75}

¹ Department of Physics and Astronomy, Clemson University, Kinard Lab of Physics, Clemson, SC 29634-0978, USA

² Santa Cruz Institute for Particle Physics, Department of Physics and Department of Astronomy and Astrophysics, University of California at Santa Cruz, Santa Cruz, CA 95064, USA

³ Department of Physics, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden

⁴ Department of Physics, KTH Royal Institute of Technology, AlbaNova, SE-106 91 Stockholm, Sweden

⁵ Università di Pisa and Istituto Nazionale di Fisica Nucleare, Sezione di Pisa I-56127 Pisa, Italy

⁶ Istituto Nazionale di Fisica Nucleare, Sezione di Padova, I-35131 Padova, Italy; riccardo.rando@pd.infn.it

⁷ Dipartimento di Fisica e Astronomia "G. Galilei," Università di Padova, I-35131 Padova, Italy

⁸ Istituto Nazionale di Fisica Nucleare, Sezione di Pisa, I-56127 Pisa, Italy

⁹ Dipartimento di Fisica "M. Merlin" dell'Università e del Politecnico di Bari, via Amendola 173, I-70126 Bari, Italy

¹⁰ Istituto Nazionale di Fisica Nucleare, Sezione di Bari, I-70126 Bari, Italy

¹¹ W.W. Hansen Experimental Physics Laboratory, Kavli Institute for Particle Astrophysics and Cosmology, Department of Physics and SLAC National Accelerator Laboratory, Stanford University, Stanford, CA 94305, USA; rac@slac.stanford.edu

¹² Istituto Nazionale di Fisica Nucleare, Sezione di Torino, I-10125 Torino, Italy

¹³ Dipartimento di Fisica, Università degli Studi di Torino, I-10125 Torino, Italy

¹⁴ CNRS-IN2P3, Laboratoire de Physique Subatomique et de Cosmologie (LPSC), Grenoble, France

¹⁵ Laboratoire Leprince-Ringuet, École polytechnique, CNRS/IN2P3, F-91128 Palaiseau, France

¹⁶ Deutsches Elektronen Synchrotron DESY, D-15738 Zeuthen, Germany

¹⁷ Institut für Theoretische Physik and Astrophysik, Universität Würzburg, D-97074 Würzburg, Germany

¹⁸ INAF-Istituto di Astrofisica Spaziale e Fisica Cosmica Milano, via E. Bassini 15, I-20133 Milano, Italy

¹⁹ Italian Space Agency, Via del Politecnico snc, I-00133 Roma, Italy

²⁰ Department of Physics and Astronomy, University of Padova, Vicolo Osservatorio 3, I-35122 Padova, Italy

²¹ Space Science Division, Naval Research Laboratory, Washington, DC 20375-5352, USA

²² Istituto Nazionale di Fisica Nucleare, Sezione di Roma "Tor Vergata", I-00133 Roma, Italy

²³ Space Science Data Center—Agenzia Spaziale Italiana, Via del Politecnico, snc, I-00133, Roma, Italy

²⁴ Laboratoire de Physique et Chimie de l'Environnement et de l'Espace—Université d'Orléans/CNRS, F-45071 Orléans Cedex 02, France

²⁵ Station de radioastronomie de Nançay, Observatoire de Paris, CNRS/INSU, F-18330 Nançay, France

²⁶ Laboratoire Univers et Particules de Montpellier, Université Montpellier, CNRS/IN2P3, F-34095 Montpellier, France

²⁷ Istituto Nazionale di Fisica Nucleare, Sezione di Perugia, I-06123 Perugia, Italy

²⁸ INAF Istituto di Radioastronomia, I-40129 Bologna, Italy

²⁹ Dipartimento di Matematica e Fisica "E. De Giorgi," Università del Salento, Lecce, Italy

³⁰ Istituto Nazionale di Fisica Nucleare, Sezione di Lecce, I-73100 Lecce, Italy

³¹ Laboratoire d'Annecy-le-Vieux de Physique des Particules, Université de Savoie, CNRS/IN2P3, F-74941 Annecy-le-Vieux, France

³² Grupo de Altas Energías, Universidad Complutense de Madrid, E-28040 Madrid, Spain

³³ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

³⁴ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

³⁵ Dipartimento di Fisica "Enrico Fermi," Università di Pisa, Pisa I-56127, Italy

³⁶ Department of Physical Sciences, Hiroshima University, Higashi-Hiroshima, Hiroshima 739-8526, Japan

³⁷ Max-Planck-Institut für Physik, D-80805 München, Germany

³⁸ Centre d'Études Nucléaires de Bordeaux Gradignan, IN2P3/CNRS, Université Bordeaux 1, BP120, F-33175 Gradignan Cedex, France; smith@cenbg.in2p3.fr

³⁹ The George Washington University, Department of Physics, 725 21st St NW, Washington, DC 20052, USA

⁴⁰ Georg-August University Göttingen, Institute for theoretical Physics—Faculty of Physics, Friedrich-Hund-Platz 1, D-37077 Göttingen, Germany

⁴¹ Science Institute, University of Iceland, IS-107 Reykjavik, Iceland

⁴² Nordita, Royal Institute of Technology and Stockholm University, Roslagstullsbacken 23, SE-106 91 Stockholm, Sweden

⁴³ College of Science, George Mason University, Fairfax, VA 22030, resident at Naval Research Laboratory, Washington, DC 20375, USA

⁴⁴ Department of Physics, Graduate School of Science, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

⁴⁵ The Oskar Klein Centre for Cosmoparticle Physics, AlbaNova, SE-106 91 Stockholm, Sweden

⁴⁶ School of Education, Health and Social Studies, Natural Science, Dalarna University, SE-791 88 Falun, Sweden

⁴⁷ Finnish Centre for Astronomy with ESO (FINCA), University of Turku, FI-21500 Piikkiö, Finland

⁴⁸ Istituto Nazionale di Fisica Nucleare, Sezione di Trieste, I-34127 Trieste, Italy

⁴⁹ Dipartimento di Fisica, Università di Trieste, I-34127 Trieste, Italy

⁵⁰ Institut für Astro- und Teilchenphysik, Leopold-Franzens-Universität Innsbruck, A-6020 Innsbruck, Austria

⁵¹ Dipartimento di Fisica, Università degli Studi di Perugia, I-06123 Perugia, Italy

⁵² Friedrich-Alexander Universität Erlangen-Nürnberg, Erlangen Centre for Astroparticle Physics, Erwin-Rommel-Str. 1, D-91058 Erlangen, Germany

⁵³ Department of Physics, Faculty of Science, Mahidol University, Bangkok 10400, Thailand

⁵⁴ Hiroshima Astrophysical Science Center, Hiroshima University, Higashi-Hiroshima, Hiroshima 739-8526, Japan

⁵⁵ Istituto Nazionale di Fisica Nucleare, Sezione di Bari, I-70126 Bari, Italy

⁵⁶ Center for Research and Exploration in Space Science and Technology (CREST) and NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

⁵⁷ Department of Physics and Center for Space Sciences and Technology, University of Maryland Baltimore County, Baltimore, MD 21250, USA

⁵⁸ Istituto Nazionale di Fisica Nucleare, Sezione di Trieste, and Università di Trieste, I-34127 Trieste, Italy

⁵⁹ Department of Information Engineering, University of Padova, I-35131 Padova, Italy

⁶⁰ Center for Space Studies and Activities “G. Colombo,” University of Padova, Via Venezia 15, I-35131 Padova, Italy

⁶¹ Korea Astronomy and Space Science Institute, 776 Daedeokdae-ro, Yuseong-gu, Daejeon 30455, Republic of Korea

⁶² Department of Physics, American University, Washington, DC 20016, USA

⁶³ Centre for Astro-Particle Physics (CAPP) and Department of Physics, University of Johannesburg, PO Box 524, Auckland Park 2006, South Africa

⁶⁴ Department of Physics, The University of Hong Kong, Pokfulam Road, Hong Kong, People’s Republic of China

⁶⁵ Laboratory for Space Research, The University of Hong Kong, Hong Kong, People’s Republic of China

⁶⁶ INFN and University of Perugia, I-06123 Perugia, Italy

⁶⁷ NYCB Real-Time Computing Inc., Lattitown, NY 11560-1025, USA

⁶⁸ Laboratoire d’Astrophysique de Bordeaux, Université Bordeaux, B18N, allée Geoffroy Saint-Hilaire, F-33615 Pessac, France

⁶⁹ Purdue University Northwest, Hammond, IN 46323, USA

⁷⁰ Solar-Terrestrial Environment Laboratory, Nagoya University, Nagoya 464-8601, Japan

⁷¹ IRAP, Université de Toulouse, CNRS, UPS, CNES, F-31028 Toulouse, France

⁷² Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Magrans s/n, E-08193 Barcelona, Spain

⁷³ Institut d’Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain

⁷⁴ Institució Catalana de Recerca i Estudis Avançats (ICREA), E-08010 Barcelona, Spain

⁷⁵ Center for Astrophysics and Cosmology, University of Nova Gorica, Nova Gorica, Slovenia

Received 2021 February 22; revised 2021 June 17; accepted 2021 June 17; published 2021 September 6

Abstract

The Large Area Telescope (LAT), the primary instrument for the Fermi Gamma-ray Space Telescope (Fermi) mission, is an imaging, wide field-of-view, high-energy gamma-ray telescope, covering the energy range from 30 MeV to more than 300 GeV. We describe the performance of the instrument at the 10 yr milestone. LAT performance remains well within the specifications defined during the planning phase, validating the design choices and supporting the compelling case to extend the duration of the Fermi mission. The details provided here will be useful when designing the next generation of high-energy gamma-ray observatories.

Unified Astronomy Thesaurus concepts: [Gamma-ray telescopes \(634\)](#); [Calibration \(2179\)](#)

1. Introduction

The Fermi Gamma-ray Space Telescope (Fermi) was launched on 2008 June 11 at 16:05 UTC. Commissioning of the Fermi Large Area Telescope (LAT) (Atwood et al. 2009) began on 2008 June 24.

On 2008 August 4, at 15:43:36 UTC (MJD = 54682.655; DOY 217.655) during orbit 813 after launch, following an early operations checkout period, the LAT began normal full science operations, surveying the entire sky every 2 orbits, or about every 3 hr. Ten years after the start of the Fermi science mission, Fermi had completed 55,992 orbits since launch. The most recent catalog of gamma-ray sources detected by the LAT is the 4FGL catalog (Abdollahi et al. 2020), using 8 yr of data. It was updated to 10 yr, called Data Release 2 (DR2),⁷⁹ and in 2021 the LAT collaboration will update the catalog to DR3 with 12 yr of data.

Detector technologies were chosen for the LAT that have an extensive history of high-energy physics applications with demonstrated high reliability. These detector technologies significantly improved the LAT’s performance compared to

previous space-based detectors in the LAT’s energy range, and allowed the LAT to meet the mission lifetime goal of 10 yr. The LAT is still performing well after 12 yr in orbit, and we describe the performance of the instrument at the notable 10 yr milestone.

In Section 2 we describe the overall instrument performance during the first 10 yr. The on-orbit calibration procedure for the LAT was introduced in Abdo et al. (2009) where calibrations and configurations in the first year of operation were described in detail. Overall performance depends on the raw data from the instrument, but also on event reconstruction. Significant postlaunch improvements yielded the “Pass 8” analysis (Atwood et al. 2013) and subsequent refinements (Bruehl et al. 2018). In Sections 3–5 we detail the evolution of the performance of the three subsystems: tracker (TKR), calorimeter (CAL), and anticoincidence detector (ACD), respectively. Each detector performance description is divided into subsections:

1. *Calibrations* are quantities used to translate electronics signals into physical quantities (e.g., the energy associated with a value read from an electronic channel). Calibrations can be updated after the data are collected, in which case data would need to be reprocessed.
2. *Configurations* define quantities necessary for optimal data taking, such as lists of disabled channels. Configurations are set before data acquisition and changes apply

⁷⁶ Former affiliation.

⁷⁷ Deceased.

⁷⁸ Funded by contract FIRB-2012-RBFR12PMIF from the Italian Ministry of Education, University and Research (MIUR).

⁷⁹ https://fermi.gsfc.nasa.gov/ssc/data/access/lat/10yr_catalog/




Figure 1. Left: the LAT during integration, showing the 16 tracker towers. Right: the Fermi spacecraft with the LAT on top, integrated on top of the Delta IIIH rocket, on the launchpad at Cape Canaveral. Figure 1 of Atwood et al. (2009) provides an artist’s cutaway drawing of the LAT.

only to future data. Configurations of collected data cannot be changed afterward.

3. *Performance* includes *failures* and *trending*. A failure is irrecoverable damage, such as an unresponsive component having no redundancy. *Trending* of operational parameters and derived quantities monitors how the performance of the LAT as a scientific instrument evolves over time.

Section 6 summarizes the performance of the trigger and readout systems, the accuracy and stability of the clocks that generate the event timestamps, and management of passages through the South Atlantic Anomaly (SAA) region of intense background radiation.

2. The LAT Instrument after 10 Years

After 10 yr of on-orbit operation, plus thousands of hours of prelaunch testing, the LAT is performing very well. No major component failures have occurred since launch, with only three notable interruptions to LAT data taking. In addition, since launch, there have been multiple improvements to LAT performance, both on orbit through updates of the onboard flight software, and on the ground in the data analysis and processing. Twenty updates to the onboard LAT flight software added functionalities, fixed bugs, provided faster task execution, and simplified operations. Three major updates of the ground-based analysis software, along with smaller incremental improvements, increased LAT gamma-ray collection efficiency, improved the accuracy of the photon energy and direction measurements, and improved the separation of gamma-ray photons from the significantly larger flux of charged cosmic-ray particles triggering the LAT (Bruehl et al. 2018).

2.1. In Orbit

Fermi was launched and deployed into low Earth orbit on a Delta IIIH rocket from Cape Canaveral in Florida (Figure 1). During launch, available surplus fuel from the rocket was used to reduce the orbital inclination to $25^{\circ}6'$, from the $28^{\circ}5'$ latitude of the launch site. The orbit is very close to circular, with an initial eccentricity of 0.0014 that decreased to 0.0012 after

10 yr. Fermi’s orbit altitude has decreased slightly over the 10 yr and is plotted as a function of time in Figure 2. The orbital period changed from the initial value of 95.7 minutes to 95.3 minutes. Altitude loss has not been constant over time. The rate of decrease was greater from 2012 to 2015 during the period of maximum Solar activity, when heating of the Earth’s atmosphere is greater and the atmospheric scale height is larger, producing more atmospheric drag on Fermi.

LAT temperatures have remained quite stable over the long-term mission. Figure 3 shows plots of temperatures of various parts of the LAT as a function of time, selected from the large number of sensors. The temperatures vary with position, with the LAT thermal radiators being the coldest, and temperatures rising toward the top of the LAT, through the LAT mechanical support structure, then the LAT calorimeter modules and electronics units, and with the LAT tracker towers and the ACD being the warmest, and having the largest annual variations. LAT temperatures are maintained within moderate ranges, with the LAT radiators having the most extreme temperatures. Small annual variations in LAT temperatures can be seen, as can short-term variations of various magnitudes due to intervals when Fermi was held in fixed pointings and other reasons. The small step increase in LAT temperatures in mid-2009 comes from the increase in Fermi’s rocking angle: to continuously scan the entire sky, Fermi points at alternating sky hemispheres on each orbit. To improve the temperature of the spacecraft batteries and hence ensure their long-term performance while maintaining uniform sky coverage, Fermi rocked $\pm 35^{\circ}$ from the orbital plane in the early mission and $\pm 50^{\circ}$ thereafter.

The large short-term temperature drops and recoveries shown in 2009 and 2018 March, and a smaller short-term change in 2008, are due to LAT power outages during the science mission. Three notable interruptions to LAT on-orbit operations have occurred since launch: on 2008 July 31 (shortly before the start of the LAT science mission), 2009 March 11, and 2018 March 16. Each time, the LAT was automatically powered off by the Fermi spacecraft. The general autonomous safing action for the LAT is to turn off electrical power to the instrument and to allow its temperature to be kept in a survival temperature range, maintained by thermostat-controlled survival heaters on the LAT




Figure 2. The orbit altitude (kilometers) of Fermi as a function of time, in years.

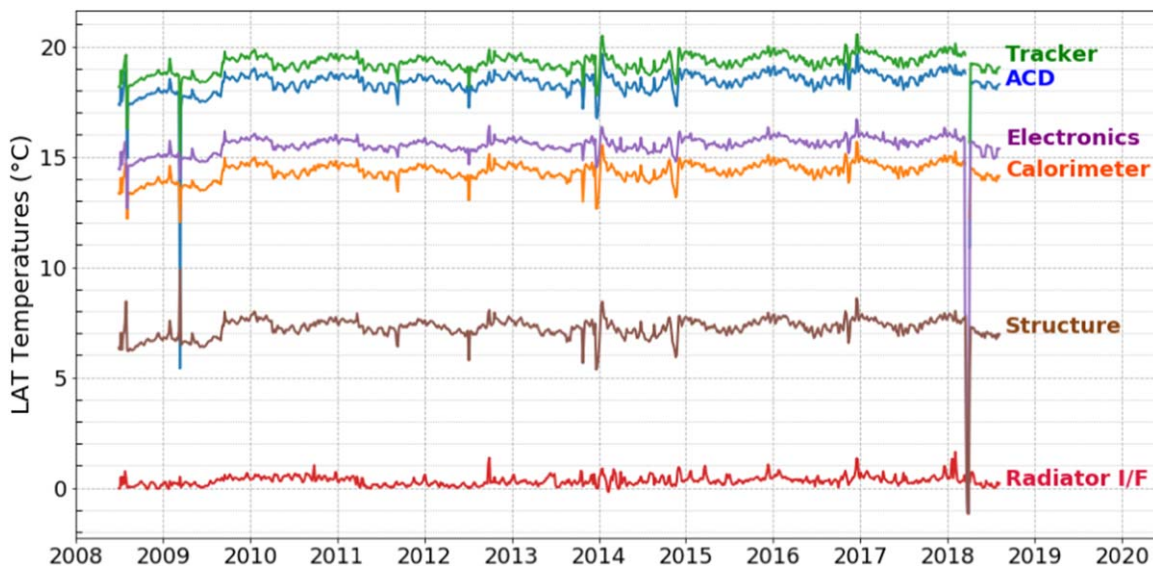


Figure 3. Temperatures ($^{\circ}\text{C}$) of various parts of the LAT as a function of time, in years. Long-term temperatures are quite stable with small annual variations and excursions due to fixed orientations and other reasons, with the larger changes and recoveries in 2008, 2009, and 2018 due to LAT power outages.

structure that are powered by the spacecraft. The LAT can survive for an indefinite period of time in that state. This LAT safe state provides time for the Fermi Flight Operations Team at the Fermi Mission Operations Center (MOC) at NASA's Goddard Space Flight Center (GSFC) and the LAT experts within the LAT Collaboration to diagnose the problem, and devise and implement a fix.

On 2008 July 31, an intermittent short first occurred in the wiring harness for the spacecraft Thermal Interface Board (TIB) module. The TIB module provides signal conditioning for several temperature sensors on the LAT to allow the spacecraft to monitor the temperature at various locations. This monitoring allows the spacecraft to perform the autonomous, automatic LAT safing if any monitored temperature is either too low or too high. The TIB harness failure caused several LAT temperature readouts to falsely appear too cold, outside

the safe range, and so the spacecraft turned off the LAT power. It was quickly realized that the abrupt, very large change in several measured LAT temperatures was not physically possible, and the problem was traced to the spacecraft TIB harness failure. The failure occurred very near the end of the initial on-orbit activation and checkout of the spacecraft and science instruments following launch. Assessment of the exact nature of the failure, followed by restarting the LAT and restoring it to readiness to perform science data collection, took 2 days. The onboard LAT temperature checks affected by the TIB harness failure have since been disabled to prevent further repeats of this problem. The disabling was done knowing that LAT temperatures will only change slowly because of the LAT's large thermal mass, which allows ground operators to see any real temperature changes and take corrective action before the temperatures change too much.

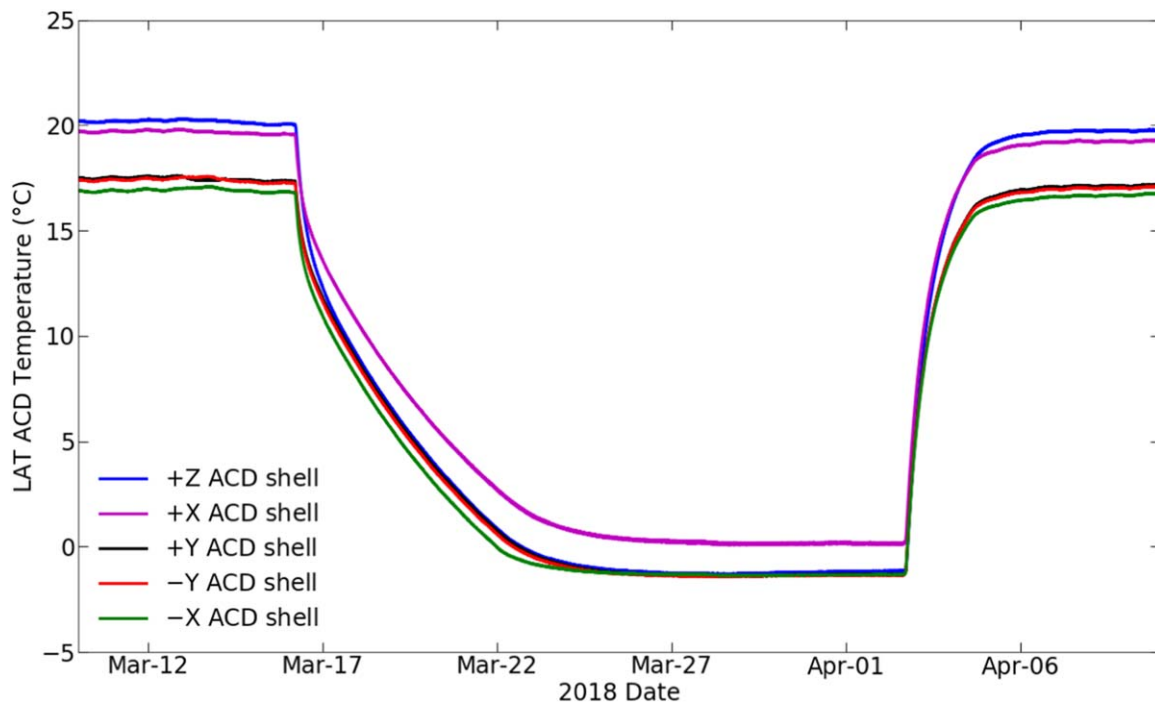


Figure 4. LAT ACD shell temperatures ($^{\circ}\text{C}$) as a function of time before, during, and after the LAT safing event in 2018 March. Because of the large thermal mass of the LAT, temperatures changed slowly over timescales of days after the LAT was powered down on March 16 and restarted on April 2. The two hotter temperatures are from the Sun-facing $+X$ side (magenta curve) and $+Z$ top (black curve) of the ACD shell, and the three colder temperatures are from the $-Y$ and $+Y$ sides (red and black curves) and anti-Sun-facing $-X$ side (green curve) of the ACD shell.

On 2009 March 11, the LAT computer that interfaces to the spacecraft suffered a software error and stopped operating normally. Several minutes later, this led the other two operating event processing computers in the LAT to also stop functioning correctly from a separate software error, as they attempted to process very high event rates during a transit of Fermi through the SAA (see Section 6.3), a region with a high density of geomagnetically trapped charged particles. During the LAT recovery, the spacecraft went into its safe mode because of an issue with the extended collection of diagnostic data, which then resulted in the LAT being automatically powered off by the spacecraft. Science data collection with the LAT was restarted by ground commanding on March 14. The errors in the LAT computers were determined through captured diagnostic data and fixed in a subsequent LAT flight software update to prevent a recurrence.

On 2018 March 16, the Fermi spacecraft went into safe mode and powered off the LAT, because the $-Y$ solar panel rotation drive on Fermi stopped moving. The LAT remained powered off for over 17 days while the solar panel problem was investigated. The LAT’s temperature was kept at survival levels by its spacecraft-powered heaters. Ground commands to power up proceeded without problems on April 2. The first science data were collected the same day, but the first science data flagged as suitable for routine science analysis were collected on April 8 because of the slow thermal response of the LAT calorimeter, which has temperature-dependent performance. Figure 4 shows the detailed time profile of LAT temperatures during the 2018 March LAT power outage, demonstrating the slow rate of change of LAT temperatures because of the slow thermal response of the 1.6 ton LAT calorimeter. The LAT missed a total of 23.3 days of routine science data collection, the longest interruption since launch.

Table 1
LAT Performance Measures After 10 Years of Fermi’s Science Mission

Parameter	Value
Science Mission start	2008-08-04 15:43:36 UTC
Days	3652
Orbits	55,179
LAT runtime fraction (including SAA transits)	99.8%
LAT triggered event readouts	598×10^9
Downlinked LAT events	120×10^9
Publicly released LAT photon-like events (transient class)	3.03×10^9
Publicly released Source Class photons	1.13×10^9

With one solar panel on Fermi stuck since 2018 March, the rocking profile enabling the LAT all-sky survey was replaced with periods of various alternating rocking angles. This keeps the Fermi power system operating nominally, but with minimal changes to the LAT’s long-term sky exposure.⁸⁰ Autonomous repointing of the LAT boresight (along the Fermi $+Z$ -axis) to temporarily search for high-energy “afterglow” emission from gamma-ray bursts has also been disabled.

Table 1 gives some numbers describing LAT performance over the first 10 yr of the mission. The overall uptime of the LAT for data taking was 99.8%, after allowing for time spent transiting the SAA (about 15% of the time). The 0.2% of time not used for data taking or in SAA transits includes the aforementioned three major outages, plus downtime for upgrades of the LAT flight software, regular (twice per year) charge injection calibrations of the LAT detector subsystems, and several short unplanned stops in LAT data taking for various reasons, including unanticipated short transits of

⁸⁰ https://fermi.gsfc.nasa.gov/ssc/observations/types/post_anomaly/




Figure 5. The daily average rate of triggers read out by the LAT over time, in years. About 20% of these event triggers pass through software filters in the LAT and are downlinked to the ground for analysis. Short-term variations in the trigger rate are caused by various operational issues, but variations corresponding to the 53 days precession period of the Fermi orbit can be seen, superimposed on the multiyear timescale variation related to the 11 yr solar cycle and the resultant variation in the density of trapped charge particles in the Fermi orbit. The spike of high LAT trigger rate in 2017 August was caused by a noisy trigger in CAL module 4, which caused periods of excessive trigger rates before it was disabled by command.

corners of the defined SAA region (see Section 6.3), configuration verification errors (see Section 2.2), and infrequent commanding errors. Almost 600 billion triggered readouts of the LAT detectors occurred in the first 10 yr, corresponding to a daily average trigger rate of about 1900 triggers per second (Figure 5). Each event readout is then passed through flight software event filters running in the two event processing unit (EPU) computers in the LAT. The EPUs perform simple filtering to discriminate between gamma-ray photons and charged particles (see Section 3.1.2 in Ackermann et al. 2012). The EPUs discard about 80% of the events, and losslessly compress the raw data for the remaining 20% of the events into a data stream with a rate of about 1.5 Mbps, corresponding to an average rate of about 400 triggers per second. These compressed event data are continuously transferred from the LAT to the Fermi spacecraft and stored in a solid-state data recorder. About once per orbit, a radio data link is opened by command between Fermi and the ground via the TDRSS geostationary data-relay satellite network to the TDRSS ground stations in White Sands, New Mexico. The data are transmitted at about 40 Mbps via TDRSS and then transferred to the Fermi MOC at the GSFC (see also Section 2.3). The MOC then transfers the raw LAT event data, plus engineering housekeeping data for the LAT and related spacecraft data, to the LAT Instrument Science Operations Center (LISOC) at the SLAC National Accelerator Laboratory (SLAC) at Stanford University. The downlinked data are checked for gaps at both the MOC and the LISOC, and any missing data are retransmitted from Fermi to the ground. There has been negligible loss of downlinked data during the mission. About 120 billion LAT events have been received at the LISOC and processed there during 10 yr.

After delivery to the LISOC, reconstruction algorithms process each LAT event to search for and reconstruct charged-particle tracks and electron-positron pair-production vertices in

the tracker, and to find associated energy signals in the calorimeter. Multiparameter selection filters applied to the reconstructed data distinguish gamma-ray photon events from charged-particle events. Multiple classes of assessed purity of photon selections are used to classify events. The summary data for the reconstructed filtered photons are then delivered back to NASA for immediate public release. LAT photon data are publicly available at NASA’s Fermi Science Support Center (FSSC)⁸¹ at the GSFC. About 1.13 billion “source” class gamma-ray photons have been detected by the LAT in 10 yr and are publicly available through the FSSC. This number of photons corresponds to a rate of about 4 photons per second. However, this photon count includes photons from interactions of Earth limb-grazing cosmic rays in the upper atmosphere: only about 40% of the detected LAT “source” photons have directions less than 110° from the local zenith and come from the sky clear of the Earth’s limb. About 3.03 billion events with slightly looser classification as photons (“transient” class events) have also been delivered to the FSSC for public release in the 10 yr. These events have slightly more contamination by residual charged-particle events, but are suitable for studies of brighter transient gamma-ray events, such as gamma-ray bursts where higher background is not a problem for analysis and source localization, and the sensitivity in these cases is photon limited rather than background limited.

2.2. Data Collection Runs

LAT data taking is split into *runs*, each one generally spanning one Fermi orbit. There were 55,259 science runs during the first 10 yr of the science mission, spanning 55,179 orbits. Occasionally, longer timespans between short SAA transits are split into two separate science runs, so the number

⁸¹ <https://fermi.gsfc.nasa.gov/ssc/data>

of science runs is slightly larger than the number of orbits. Not all of these science runs were used to observe the gamma-ray sky. During 2011 and 2012, 24 observation runs by the LAT were performed to search for Terrestrial Gamma-ray Flashes (Briggs et al. 2013) in which the LAT boresight was pointed at the orbit nadir, i.e., toward the Earth, accomplished by using a 180° rocking angle (angle of the Fermi +Z-axis from zenith pointing, perpendicular to the orbit plane). In addition, 55 non-routine pointings of Fermi have been done to observe Targets of Opportunity (ToO) with the LAT.⁸² During ToO observations, the Fermi spacecraft inertially points the LAT boresight toward the target when it is not occulted by the Earth, and resumes routine sky survey during target occultations. During 2014, Fermi's observing strategy was altered to provide more exposure toward the Galactic Center to increase sensitivity for a search for enhanced gamma-ray emission from the anticipated disruption of the object G2 during its closest approach to the supermassive black hole at Sgr A* (Schartmann et al. 2012).

Dividing LAT data taking into a sequence of runs also has the benefit of enabling a process of regular checking and reloading of the extensive LAT onboard electronic configuration. A large number of data registers in the LAT trigger and dataflow electronics and the detector subsystem electronics define and control data capture and movement in the LAT. This makes the LAT highly configurable, allowing detailed optimization of LAT performance and extensive reconfiguration ability in the event of detector or electronics component failures or performance degradation during the mission.

The LAT operates in a high-flux environment of charged particles (cosmic rays and terrestrial sourced particles) that can cause single event upsets in the electronics. To reliably operate in such an environment, the LAT electronic configuration is regularly automatically verified. Before each new science run, the configuration is loaded, then read back and checked. A check mismatch generates an error message from the LAT, and generally, data taking does not start for that run, although that can be overridden. At the end of each run the LAT electronic configuration is again read out, and if a mismatch is found, the configuration data are recorded for later analysis on the ground. Additionally, a command to check the electronic configuration is run on the LAT at the end of every transit through the SAA region of the high density of charged particles. In the first 10 yr of the mission, 11 configuration mismatches occurred: four in ACD electronics, three in Calorimeter electronics, two in tracker electronics, and two in dataflow electronics. All of these corruptions occurred during transits through the SAA region. The SAA transits also produce most of the corruptions of LAT computer memory (see Section 6.3).

In addition to the science runs, charge injection calibrations are periodically scheduled. Several predefined charge levels are injected into the detector electronics readouts to calibrate readout response (Baldini et al. 2006). Early in the mission, they were performed frequently. The LAT's stability having been established, we decreased the rate to twice per year. In all, 80 routine charge injection calibrations across the three ACD, CAL, and TKR subsystems were performed during the first 10 yr of the science mission.

2.3. Event Reconstruction, Analysis, and Storage

The LAT typically generates about 15 GB of compressed raw data every day. The data are downlinked to GSFC (Section 2.1) and delivered to SLAC for processing on the SLAC computer farm where over 3000 CPU cores are available for event processing, including event reconstruction and filtering. The incoming raw data are split into multiple parallel processing jobs to process the data quickly, so as to release it as soon as possible and to obtain time-critical results for any transient gamma-ray sources. Figure 6 shows histograms of data processing times for the LAT runs over the 10 mission years. The data processing expands the data volume by a factor of ~ 50 with respect to the raw data volume, and 3.7 PB of disk space and tape storage are used for storage and analysis. Legacy data are usually stored on tape, while a copy of the data products from the current version of the reconstruction and analysis pipeline is kept on disk for fast and easy access. However, in the future, to reduce data storage costs, only the most recent year (or so) of data will be kept on disk, with older data stored on tape only. By far the major part of the data consists of reconstructed events, which are more than 10 times more voluminous than the channel-by-channel digitized electronics data from which the reconstructed event data are derived, and 25 times larger than the tabular photon data in ROOT format⁸³ which are extracted from the reconstructed event data. About 200 MB of high-level photon data and ancillary data in FITS files are sent to the FSSC at GSFC each day.

The 1500 cores available for LAT needs at the IN2P3/CNRS facilities in Lyon, France, are commonly used for Monte Carlo simulations of LAT performance and results, thus assuring the availability of the SLAC computer farm for event reconstruction.

2.4. Data Quality Monitoring and Trending

LAT subsystem and ground-support systems performances are continuously monitored through the LAT Data Quality Monitoring (DQM) System to guarantee the quality of the LAT data delivered for science analysis. During each step of the ground processing pipeline, histograms of different quantities are generated and stored. The monitored quantities are related to the stability of the detectors (such as pedestals, gains, and rates of each channel), to the rates of events passing the different trigger conditions and onboard filters, and to high-level reconstruction outputs such as the average photon rate or the event energies. The performance of the timing system is also monitored (GPS lock loss and 20 MHz clock stability, see Section 6.2).

Weekly duty scientists use web-based tools to check all the processed science data from the LAT. Plots of all monitored quantities are automatically generated, and automatic alarms are generated if a monitored quantity deviates from its allowed range. Currently, about 12,000 parameters are monitored and the DQM system makes about 4100 checks on parameter ranges.

Many of these parameters vary significantly during an orbit as they depend on the geomagnetic cutoff (which influences the rates measurement, driven by the rate of charged cosmic rays) and on the spacecraft attitude (since it influences the arc length

⁸² Fermi Targets of Opportunity are listed at <https://fermi.gsfc.nasa.gov/ssc/observations/timeline/too/>.

⁸³ <https://root.cern.ch/>

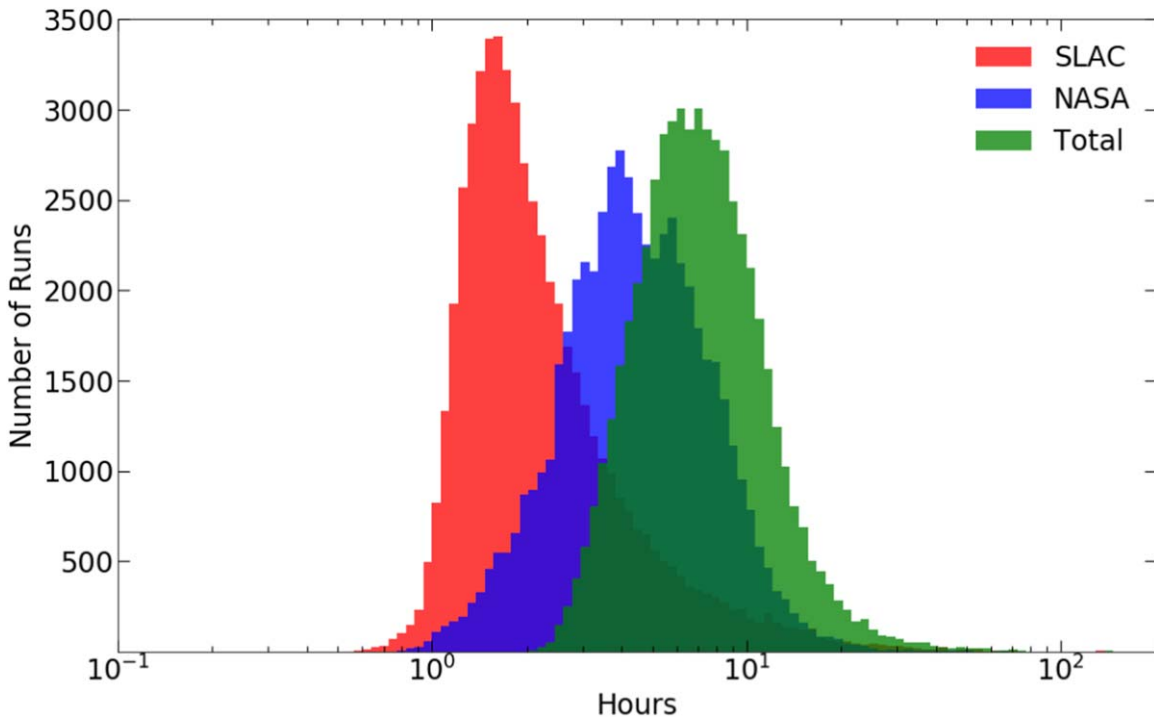


Figure 6. Histograms of data processing times, in hours, for the LAT runs over the 10 mission years, shown with logarithmic time bins. The overall total data processing time histogram is shown in green, which is the sum of data downlink and delivery time through NASA, shown in the blue histogram, plus the subsequent processing time at SLAC, shown in the red histogram.




Figure 7. Cumulative number of seconds flagged as bad. The red line shows the total bad time, while the blue line shows bad time associated with solar flares. The steady increase between 2011 and 2015 is due to increased solar flaring during the solar cycle maximum. The two jumps in 2009 March and 2018 April correspond to the LAT temperature stabilizing periods after LAT restarts. The recovery time was much longer in the second case because the LAT was powered off longer, and the temperature decreased to the lower allowed limit.

of the gamma-ray bright Earth limb in the LAT field of view). For some monitored quantities, these dependencies are parameterized as a function of the geomagnetic McIlwain-L and spacecraft rocking angle, leading to normalized quantities that ideally are independent (within 20%) of orbit position and spacecraft pointing, and therefore are easier to monitor.

If the DQM system identifies a problem that can potentially influence the data quality, a bad time interval (BTI) is flagged by setting the variable DATA_QUAL to a negative value for the corresponding time range in the delivered data file. While

the alarms are automatically generated by the DQM system, the final decision about a BTI is taken by the duty scientist and members of the LAT collaboration.

The operational environment of the LAT can influence the quality of the data. The main example is solar flares, which can cause excessive veto signaling from the ACD caused by X-ray absorption in the ACD tiles, reducing the LAT sensitivity to gamma-rays. In this case, the standard instrument response functions do not describe the detector accurately, and the inclusion of data from these time intervals in science analysis can potentially alter the results. These time intervals are identified by searching for high rates in the ACD and a deficit in the reconstructed photon rate, and are flagged as bad. BTIs due to solar flares are flagged with DATA_QUAL equal to -1 (see Appendix A of Ackermann et al. 2012 for a detailed description of the effect of solar flares). The Pass 8 event reconstruction uses the ACD veto signals to substantially mitigate the effect of X-rays in the ACD. Analyses of gamma-ray emission associated with solar flares have been developed to compensate for dead time effects in the LAT during periods of intense X-ray emission in the flares.

Figure 7 shows the cumulative sum of BTIs as a function of time, and the bad time due to solar flares. The period of increased solar flaring during the solar cycle maximum between 2011 and 2015 can be clearly seen in the steady rise of bad time during that period. After mid-2015, the combination of the switch to Pass 8 reconstruction and the solar cycle minimum greatly reduced BTIs due to solar flares. The two jumps in 2009 March and 2018 April correspond to the periods after LAT restarts, described in Section 2.1, while LAT temperatures were restabilizing. Since many instrument parameters depend on temperature, all runs before the instrument temperatures returned to the nominal range were marked as

Table 2
Sample of Types and Categories of Parameters Trended in Data Quality Monitoring (DQM) and for LAT Performance Trending

Spacecraft	LAT	Ground Processing
Orbit altitude	Power supply voltages	Data processing latency
Onboard data storage use	Power supply currents	Subsystem calibrations
Ground contact periods	Temperatures	Data storage volume
Data downlink latency	Active thermal control	Public data volume
Downlinked data fraction	CPU memory errors	Photon classes rates
Spacecraft attitude	Event trigger rate	
Spacecraft position	Data acquisition uptime	
	Channel and subsystem rates	
	Channel pedestals (ACD, CAL)	
	Channel calibrations	

bad. The fraction of bad data is about 0.2% of the total LAT data set.

Monitoring trends of the Fermi spacecraft and LAT performance and associated ground operations performance is done continuously to complement and extend the DQM described above, and to provide a long-term history of key parameters of the LAT and Fermi performance. Review and assessment of spacecraft and instrument performance and operations trends are done at weekly meetings of the spacecraft and instrument operations support teams, and performance and operations trends are also more comprehensively reviewed and assessed at three-month intervals. Trends of performance-related quantities are assessed within the reporting period, and also over the mission lifetime.

Both direct and derived measurements for the spacecraft, the instruments, and ground processing are monitored for trends. Table 2 lists some key performance parameters that DQM watches closely. “Data latency” is the time interval between raw data collection on the LAT and the resulting processed and filtered gamma-ray photons being delivered to NASA for public release through the FSSC (see Figure 6). “Public data volume” is the number of gamma-ray photon events released to the public through the FSSC.

3. Tracker

The tracker (TKR) is a pair-conversion detector. The TKR is a 4×4 matrix of towers, each a stack of 18 X - Y silicon (Si) strip detector modules, with interleaved tungsten (W) foils for converting gamma-rays to e^+e^- pairs and for initiating electromagnetic showers. The charged particles ionize the silicon as they pass through the layers, providing measurable tracks. The directions of incoming photons or charged particles (“primaries”) are evaluated and correlated with the location of the tiles and ribbons of the LAT anticoincidence detector (see Section 5) showing coincidental signals to enhance charged-particle background rejection.

The single-sided Si strip detectors, each $8.95 \text{ cm} \times 8.95 \text{ cm} \times 400 \mu\text{m}$, are arranged into planes of 4×4 detectors, with two planes having mutually orthogonal (X - Y) strips in each module. The top 12 modules include tungsten converter foils that are 0.03 radiation lengths thick. The W foils in the next four modules below these are ~ 6 times thicker to maximize conversion efficiency at the cost of additional Coulomb scattering and slightly worse angular resolution. Hits in consecutive TKR layers define the “primitive” signal used to build TKR triggers: the TKR generates the majority of LAT trigger requests for gamma-ray physics by triggering when three consecutive X - Y planes are hit




Figure 8. ToT distribution in one run for the whole TKR for a selection of MIP events, and fit to the MIP peak.

(“three-in-a-row”). The last two tracking modules closest to the LAT calorimeter have no conversion foils: hits in these layers can participate in but not initiate TKR triggers.

Each individual Si plane in the TKR contains 1536 strips, read out by 24 front-end analog ASICs in a daisy-chain configuration. Two digital readout controllers at both ends nominally deal with half of the front-end ASICs. Strip readout is digital: “hit” discriminator thresholds are set to one-fourth of the charge deposited by a single minimum ionizing particle (an “MIP”). Recording the time that the strip signal remains above threshold provides an estimate of the total charge deposit (i.e., of the total energy lost by ionizing particles traversing a strip). This Time Over Threshold (ToT) measurement is performed on the logical OR of all the strips on a TKR plane and is included in the data stream along with the digital position of the strips above threshold. The overall small energy loss in the TKR is especially important for primaries with energy $\lesssim 100 \text{ MeV}$, which deposit most of their energy before reaching the calorimeter. Atwood et al. (2007) describe the TKR subsystem in more detail.

3.1. Tracker Calibration

Detailed descriptions of the calibrated quantities and calibration procedure for the TKR are given in Abdo et al. (2009). Here we briefly review the calibration constants and discuss how some of the values have evolved during the mission.




Figure 9. Trending of the minimum ionizing particle (MIP) peak energy value during the first weeks after launch (left) and for the remainder of the mission (right).

The split point in the front-end daisy chain is nominally in the middle of the 24 ASICs, so that two equal groups of 12 ASIC outputs are directed to the two controllers. The split point can be changed, e.g., to isolate a faulty front end. A single readout controller ASIC failed before launch (see Section 3.3), but a full readout of the plane was guaranteed by applying a different split point. No failures have occurred on orbit, and all other plane readout split points are still set to the middle of the plane.

The deposited charge in any TKR channel is estimated from the ToT (see Section 3): time is measured in 20 MHz clock ticks and converted into charge (fC). Calibrations include one overall energy-scale constant and one scale factor per channel to allow for channel-by-channel scale adjustments. The overall scale factor affects the energy measurement directly as a multiplicative constant, so good accuracy is necessary. This global energy scale is monitored by observing the energy deposited by minimum ionizing particles, mostly protons. The ToT values in the MIP data set are corrected for the incidence angles on the Si planes (a 10% effect). The distribution is fitted with a Landau curve convolved with two Gaussians, and the location of the MIP peak is determined with a statistical uncertainty around 0.2%. Systematic effects include charge sharing and a residual effect of the incidence angle (of order 1%), evaluated by Monte Carlo simulations and corrected for. A varying amount of non-MIP contamination in different runs causes the largest remaining systematic effect. In Figure 8 the uncalibrated ToT is shown for one recent run, together with the fitted curve.

Figure 9 shows the location of the uncalibrated MIP peak as a function of time; each point is the average of several runs, and residual scatter, larger than the statistical uncertainties, is attributed to varying non-MIP contamination in different runs. Rare major changes in the event processing algorithms cause changes in the energy scale that would appear as abrupt offsets in the MIP peak trend. These are quantified and scaled away so they do not appear in the figure. A $\sim 10\%$ increase in the overall energy scale occurred in the first months (Figure 9, left). This is interpreted as initial radiation damage quickly coming to saturation, but a detailed study to understand the root causes of this phenomenon has not been performed. Since then, the additional increase is less than 1% (Figure 9, right). No channel-by-channel variations of more than a few percent have been seen.




Figure 10. Dead and noisy channels in the tracker as a function of time.

We maintain lists of dead and noisy single readout channels in order to flag those that are faulty (see Figure 10). At launch, 3661 readout strips were declared bad (0.31% of the total); almost 40% of these were located in Tower 0, the first tower to be assembled. Most bad strips appeared to be too quiet: 2047 were apparently strips disconnected from the readouts, collecting no signal and having an abnormally low noise level, while 413 channels had a dead preamplifier showing zero noise and no signal. There were also 998 partially disconnected strips where one or more of the wire bonds along the ladder are defective, leading to intermediate noise levels. Finally, 203 strips were electronically disabled for being too noisy (0.02% of total). After 10 yr of operation in orbit, the number of bad strips increased to 4087 (0.46%). Surprisingly, the number of disconnected and dead channels diminished to 2008 and 405, respectively, indicating that in some cases the true cause of the malfunction is not clear and in some cases is reversible. The number of partially disconnected strips increased to 1076 while the number of noisy channels increased to 598 (0.07%), mostly in Towers 0 and 3. In Tower 0, most of the noisy strips are in one ladder on silicon layer 14, and were either noisy before launch or became noisy in the first few weeks after launch. In Tower 3, most of the noisy strips are in one ladder on silicon layer 35 at the top of the tower, nearest the ACD shell. In Tower 3, the noisy ladder started becoming noisy in 2010, and more strips in that ladder became noisy over a period of several




Figure 11. Tracker bias currents for the tracker towers after the Fermi safe mode and LAT power-off period in 2018 March. The bias current for Tower 3 (the top blue curve) shows a slower recovery over a duration of many days compared to the other towers, as the noise in layer 35 of tower 3 slowly recovered to prefault levels over that extended time period.

years. After 10 yr, 58% of the 598 masked noisy strips were on two of the 2304 silicon strip ladders in the LAT. The recovery from the extended LAT power-off period in 2018 April provided more evidence of nonstandard behavior of the noisy ladder in Tower 3. Figure 11 shows the TKR bias currents for the 16 towers after the 2018 March LAT power-off period. The bias current for Tower 3 (the top blue curve) shows a slower recovery over a duration of many days compared to the other towers, as the noise in layer 35 of Tower 3 slowly recovered to prefault levels over that extended time period.

In Figure 12, left, the TKR readout noise is averaged over all channels (ENC is equivalent noise charge). The relative increase is within expectations (a couple percent). For comparison in Figure 12, right, the long-term bias current increase of each module and the total are shown: the noise increase can be correlated with the overall increase in leakage current, attributed to radiation damage. The rapid variations on top of the slowly increasing trend in the plot can be correlated to temperature variations as measured by the sensors in the TKR, in Figure 13.

In Figure 14, left, the noise level for each channel in the tracker is shown for two runs, one after the launch and one after almost 10 yr of the mission. Very low values correspond to the dead or disconnected channels, the uncertainty on the measurement of the noise level for a single channel is $\sim 7\%$. The distribution of the ratio of the values in the later run over the earlier run is shown in Figure 14, right: a weak trend toward values greater than 1 (i.e., increasing noise) is apparent. To put the noise in context, a MIP releases ~ 5 fC, as seen in Figure 8, i.e., $\sim 31,000$ electron-hole pairs.

The TKR hit efficiency can be estimated by selecting MIP tracks and searching for missing hits where the tracks cross the active volumes (see Figure 15, left). The efficiency uncertainty is better than 10^{-4} . An exponential fit shows that the efficiency decreases by 3×10^{-5} per year, with no practical consequences. Similarly, by looking at how many recorded tracks that should have produced a trigger in an adjacent tower failed to issue one, we can estimate the TKR trigger efficiency (see Figure 15, right). The uncertainty on this estimate is $\sim 10^{-4}$, and an exponential fit gives an increase of trigger efficiency of 2×10^{-5} per year, mostly due to a small increase in the first year. Note that the exact values in Figure 15 are also affected

by the details of the reconstruction process. Notably, with the migration to the Pass 8 reconstruction, a systematic shift of 0.1% is scaled away from the data, being irrelevant for performance monitoring.

The mechanical alignment of the TKR is also regularly calibrated. One rotation matrix aligns the reconstructed direction of each event from the TKR to the spacecraft (and thence the celestial) reference systems. Also, translation and rotation offsets of each Si plane in the TKR mechanical coordinate reference system are accounted for by additional calibration constants. The relative alignment of the LAT reference system to the spacecraft reference system, derived from the spacecraft star tracker observations, is determined by minimizing the residuals of the measured locations of known gamma-ray sources in the sky. The relative alignment is known with an accuracy better than $5''$ on the three rotation angles and shows no significant evolution in time. The relative alignment of each Si plane in a tower with respect to the LAT reference system is given as three offset translations and three rotation angles. The sensitivity of the measurement is better than $1 \mu\text{m}$ ($2 \mu\text{m}$) for the x , y (z) translations, and then 0.02 mrad (0.01 mrad) for rotations around the x , y (z) axes. For each of the 288 X - Y silicon planes in the TKR, we compute the time average of the three lateral offsets and of the three rotation angles, and in Figure 16 we show the deviation from the mean as a function of time of all planes; no evolution can be observed. The overall translation uncertainties are well below the position resolution of a charge cluster given the TKR strip pitch of $228 \mu\text{m}$, and we estimate that spurious rotations would need to be ~ 0.3 mrad at the tower level before affecting the high-energy PSF.

3.2. Tracker Configurations

The lists of *dead* and *hot* channels, defined in Section 3.1, are also used as configurations. In addition, the *trigger thresholds* define the minimum amount of charge to generate a trigger request, with one value for each front-end ASIC. In Figure 17, left, the distribution of threshold values for each tracker channel is shown for two runs, one after the launch and one after almost 10 yr of mission; no difference is evident. The distribution of the ratio, i.e., the values for the later run divided by the values for the earlier run, is shown in Figure 17, right. The standard deviation is comparable with the uncertainty on the determination of the threshold for a single channel ($\sim 1\%$).

3.3. Tracker Component Performance

In 2006, before launch, a single Tracker Readout Controller failed during a LAT powerup during ground testing. Schedule constraints dictated that this failed circuit in silicon Layer 0 of Tower 10 would not be repaired or replaced. The split point described in Section 3.1 was moved such that the entire plane was read by the left-end controller. On orbit, no other tracker component failures have occurred (dead strips are not considered).

4. Calorimeter

The calorimeter (CAL) detector subsystem provides the primary contribution to the energy estimate for all but the lowest energy events, $\lesssim 100$ MeV, where depositions in the TKR are important. The calorimeter is segmented on dimensions comparable to the characteristic longitudinal and




Figure 12. Left: tracker equivalent noise charge (ENC) trending (average over all strips; statistical error bars are smaller than markers dimension). Right: tracker bias current trending.

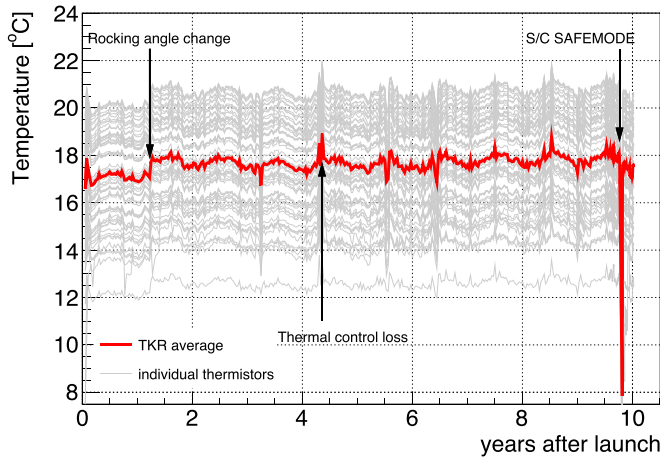


Figure 13. Tracker temperature, average (red) and each thermistor (gray) as a function of time (years).

transverse length scales of electromagnetic showers, the radiation length, and the Moliere radius, respectively. This segmentation allows three-dimensional imaging of shower development and enables good calorimetry ($\Delta E < 20\%$ FWHM) up to ~ 2 TeV (Bruehl 2012) in a relatively thin calorimeter ($8.3 X_0$ on axis). A moment analysis of energy depositions within the CAL locates the energy centroid and provides directional information used in the reconstruction of the full event. Information from the CAL is also used in rejecting cosmic-ray background events through a comparison of the location of the energy centroid and the projected TKR direction and through analyzing the CAL shower topology.

The CAL is a 4×4 array of modules matched to the array of 16 TKR towers in the LAT. A tower electronics module (TEM) provides the interface between the TKR tower and CAL module to the global-trigger and dataflow electronics system in the LAT. Each CAL module consists of eight layers of 12 CsI (TI) crystals in a hodoscopic arrangement, with successive layers rotated by 90° with respect to each other. Each of the 1536 CsI(Tl) crystals is $2.7 \text{ cm} \times 2.0 \text{ cm} \times 32.6 \text{ cm}$, read out at each $2.7 \text{ cm} \times 2.0 \text{ cm}$ end face with a dual PIN photodiode assembly. Each crystal provides three coordinates of the centroid of energy deposition within it. Two come from the location of the crystal within the CAL module. The third,

the longitudinal position, comes from measuring the scintillation light asymmetry at the two ends of the crystal. The longitudinal position has \sim millimeter–centimeter resolution, which is comparable to the transverse dimensions of an individual crystal. With this configuration, a three-dimensional image of the shower can be constructed from the ensemble of crystal energy measurements and position measurements.

Measuring event energies over the entire five orders of magnitude for which the LAT operates requires that each crystal end be read out in four gain ranges. The two photodiodes at each crystal end have different sizes: the large (small) one is for low (high) energy depositions, called LE (HE). Two readout paths provide readout channels with gains of $\times 1$ (LEX1, HEX1) and $\times 8$ (LEX8, HEX8). Thus the 1536 crystals have 3072 end faces read out with 6144 photodiodes into 12,288 spectroscopy channels. A single analog front-end ASIC handles signals from the dual photodiode assembly at each crystal end face.

In addition, each ASIC contains two fast analog trigger discriminators, the Fast Low Energy (FLE) and Fast High Energy (FHE) discriminators nominally set to 100 and 1000 MeV. Each CAL module forms two trigger-request primitives, called CAL-LO and CAL-HI respectively, from the logical OR of the corresponding low-energy and high-energy trigger discriminator outputs of all the front-end ASICs within that module. Each CAL module presents its fast trigger primitives to the central trigger logic of the LAT, the Global-Trigger Electronics Module (GEM), along with trigger primitives coming from the TKR and ACD. Based on a configurable logic table, the GEM adjudicates whether the presented trigger primitives warrant reading out the LAT detector, and if so, generates a Trigger Acknowledge message that is distributed back to each TEM and ACD electronics module. In response to the Trigger Acknowledge, analog signals from all 3072 CAL front-end ASICs are digitized; thus, the entire CAL is read out for every such trigger. In the nominal instrument configuration for science data acquisition, only one of the four gain channels at each crystal end is digitized: the highest gain range that is unsaturated, which preserves the greatest energy resolution in each crystal. CAL data are sparsified with zero-suppression logic in each front-end ASIC that eliminates energy depositions below a configurable Log-Accept (LAC) threshold, nominally set to 2 MeV.

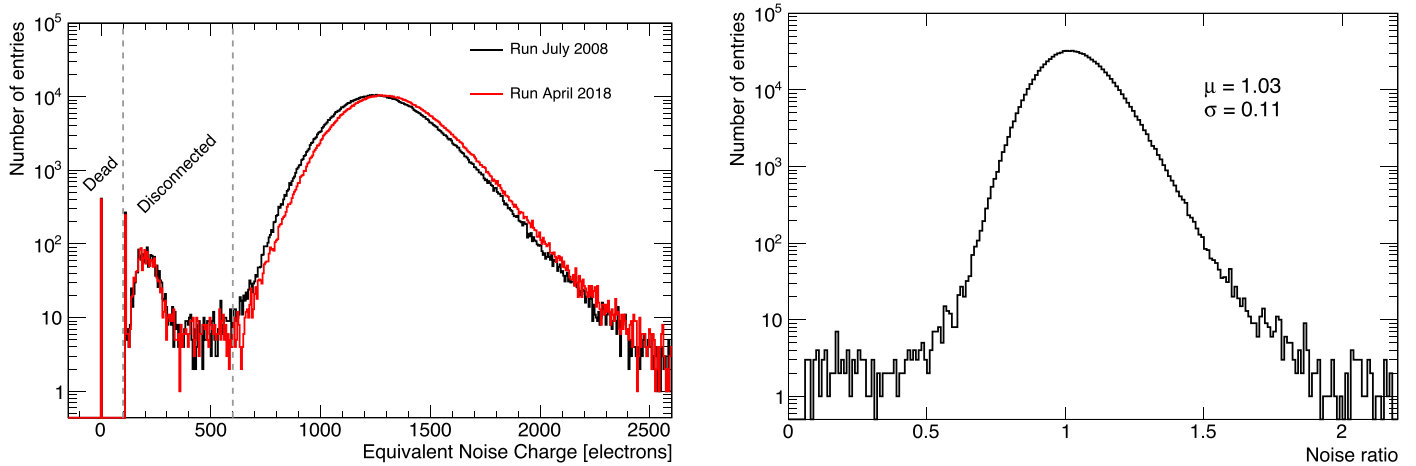


Figure 14. Left: tracker equivalent noise charge (ENC) distribution for two runs, one early (black data) and one late (red data) in the mission. Right: the distribution of the ratio of the ENC values for the two runs.




Figure 15. Left: tracker hit efficiency. Right: trigger efficiency. Pass 8 data are scaled for trending purposes.

See Grove & Johnson (2010) and Abdo et al. (2009) for more complete descriptions of the CAL and the LAT trigger system.

4.1. Calorimeter Calibration

Abdo et al. (2009) detail the calibrated quantities and calibration procedure for the CAL. Here we briefly review the calibration constants and discuss how some have evolved during the mission.

On-orbit calibration of the CAL detector subsystem is performed in 6 Ms intervals (five times per year). For each individual CAL crystal, the three quantities that need to be monitored and calibrated over time are the energy scale, the positional light yield dependence, and the pedestal values (i.e., the front-end output for zero input). There are two primary aspects to calibration of the crystal energy scales: monitoring nonlinearities in the front-end electronics, and accounting for degradation in the CsI light yield. Twice per year, the electronics response is monitored by injecting pulses of known amplitudes into the CAL input channels and fitting the resulting digital signals. Characterizations of the nonlinearities have not changed by more than 0.5% from values found during

on-ground testing. Therefore, we have not updated the electronics calibration constants for any CAL crystals.

Radiation exposure in the Fermi orbit, primarily from the intense particle fluxes within the SAA (see Section 6.3) and secondarily from Galactic cosmic rays, degrades the performance of the CsI(Tl) crystals over time. Radiation exposure induces visible and infrared absorption bands in the CsI(Tl) that reduce the scintillation light output and increase the attenuation of scintillation light as it propagates along the length of the crystal (Woody et al. 1992; Kazui et al. 1997; Chowdhury et al. 1999). Both of these changes are observed in the CAL. The light yield for a given energy deposition is calibrated on orbit using cosmic-ray events identified by reconstruction algorithms. Since launch, we have observed a decrease in the average energy scale of $\lesssim 1\%$ per year, with the degradation rate decreasing with time (see Figure 18). On-orbit calibration minimizes the effect on the reconstruction of event energy and shower parameters.

The light asymmetry, used to infer where on the crystal an interaction occurred, is the logarithm of the ratio of signals from diodes at opposite crystal ends. The conversion from asymmetry to longitudinal position uses parameters that were determined for each crystal before launch and are regularly updated at each 6 Ms calibration epoch. With two different




Figure 16. Alignment of the 288 X–Y silicon planes in the LAT reference system, measured 10 times per year: the offsets along the x , y , z directions (left top, middle, bottom) and the rotations around the same axes (right top, middle, bottom).

diodes at each end, there are four possible readout combinations; accordingly, four asymmetry maps are necessary. Each crystal is separated into twelve bins evenly spaced longitudinally, and the asymmetry is calibrated using cosmic rays. The maps of the light asymmetry in each crystal are generated from a fit to the inner ten bins, while the bins closest to the photodiodes are ignored because of the known change in light collection visible in Figure 7 of Atwood et al. (2009). As the attenuation of scintillation light within each crystal increases (i.e., as the attenuation length decreases) over time from radiation exposure, the light asymmetry increases. If this change in asymmetry were not accounted for, the reconstructed event position would be biased increasingly with time. Figure 19 shows an example of how the “bias” between true event position and measured event position would have evolved through the mission had the increasing attenuation not been accounted for, i.e., had the crystal light asymmetry maps not been updated with time. The data show the longitudinal position bias, averaged over all crystals, when reading the larger diodes at both ends: the increasing light attenuation would have introduced a bias of up to ~ 1 mm per year in reconstructed event positions. However, the regular

calibration updates to the light asymmetry maps account for the changing attenuation and remove this bias in the event reconstruction on the ground. Results are similar for the other three possible diode combinations.

The pedestals are offset voltages that define the zero-point for the energy scale of each readout channel for each crystal. The pedestal values will drift if the satellite pointing changes in such a way that temperatures change in the CAL. The pedestals are constantly monitored using periodic triggers issued by the LAT at 2 Hz, with new average pedestal values calculated every 6 Ms.

4.2. Calorimeter Configurations

More than 24,000 registers within or directly related to the CAL must be set in order to configure the CAL for science data acquisitions. Section 2.2 described the reinitialization process that occurs roughly every orbit. About half of these configuration registers are intended to be set at particular energies (e.g., trigger threshold registers), and therefore the register settings can be expected to need to be varied as the instrument response evolves with time in orbit. Other registers would not be




Figure 17. Left: tracker threshold distributions for two runs, one early and one late in the mission. Right: the distribution of the ratio of the threshold values for the two runs.




Figure 18. Relative energy scale, averaged over all CAL crystals, as a function of time, referenced to the values just after launch. The calibration at the start of the seventh year was neglected during the Pass 8 software upgrade.

expected to change, except in response to a component failure (e.g., a trigger-enabling mask for an individual readout ASIC would be changed if the trigger threshold discriminator were to become noisy). Still other registers, e.g., those defining the crystal readout order, should remain static throughout the Fermi mission because the definition of the nominal science data acquisition has not changed.

Registers defining the CAL threshold discriminators are listed in Table 3, followed by descriptions of the parameters. Each discriminator was calibrated—i.e., the correspondence between register setting and threshold value in energy units was measured—before launch and verified during on-orbit commissioning (Abdo et al. 2009). Prior to the start of nominal science operations in 2008 August, each register was set to an intended value in energy units. While the threshold settings have been constant for all nominal science data acquisitions since then, the register settings for each channel have changed as the crystal scintillation light yield has declined with cumulative radiation exposure in orbit (see the discussion in Section 4.1).

The Log-Accept, or LAC, zero-suppression threshold sparsifies CAL data by qualifying for inclusion into the CAL

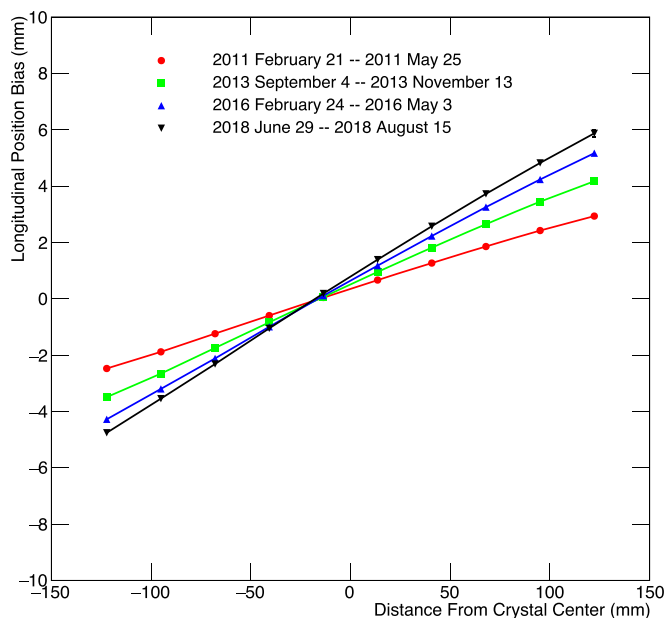


Figure 19. Average longitudinal position “bias” (see text) as a function of distance from crystal center during four different 6 Ms time periods corresponding to approximately 2.5 yr (red circles); 5 yr (green squares); 7.5 yr (blue upward-pointing triangles); and 10 yr (black downward-pointing triangles) since launch.

Table 3
Calorimeter Settings Applied to Each of the 3072 Crystal end Faces

Name	Purpose	Intent
LAC	Zero-suppression threshold	2 MeV
FLE	Low-energy trigger threshold	100 MeV
FHE	High-energy trigger threshold	1000 MeV
ULD	Range-selection upper level	−5%

data stream only those crystals for which at least one end face registers a pulse in excess of the nominal setting of 2 MeV, shown as the threshold “intent” in Table 3. The LAC settings for all CAL channels have been adjusted three times: twice in the early science mission to apply flight calibrations as the

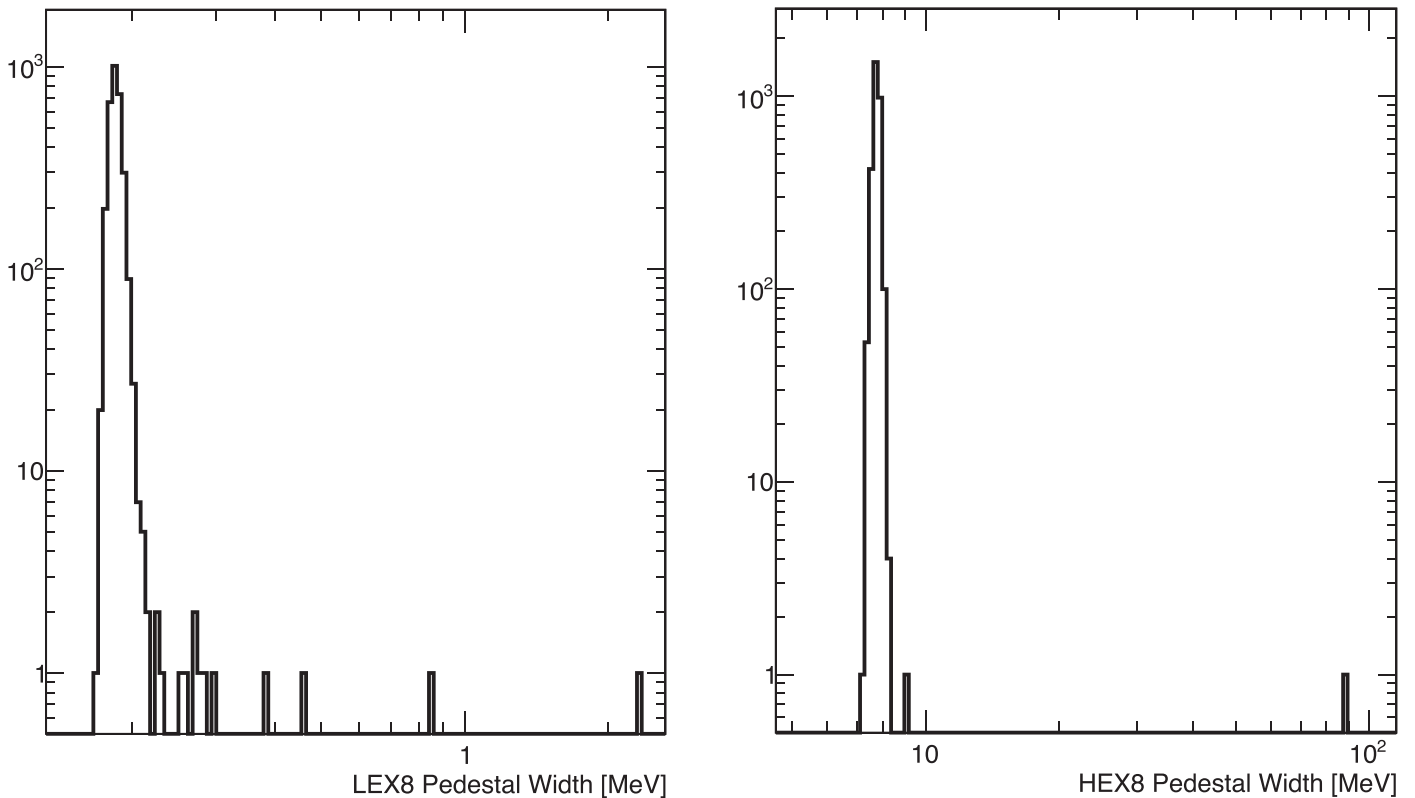


Figure 20. Distribution of electronic pedestal widths for the CAL low-energy photodiodes (left) and high-energy photodiodes (right) converted to MeV. The distributions combine readout channels from both faces of each crystal and are for the 6 Ms calibration interval at the end of year 10 of the mission. The outlier channels shown in each plot are discussed in Section 4.3.

temperature of the CAL slowly settled, and once in 2011, to account for the declining scintillation light yield from the CAL crystals and maintain the 2 MeV intent for the LAC settings, keeping the actual LAC threshold near the intended 2 MeV value despite the reduced light output.

The digital-to-analog converter that controls the LAC threshold setting has a step size of 0.44 MeV, which sets a natural cadence for adjusting the setting. The nominal 2 MeV LAC threshold is well above the typical pedestal noise (pedestal width) of ~ 0.3 MeV. Figure 20 shows the pedestal widths for all 3072 LEX8 and HEX8 channels. If a front end becomes noisy, the threshold will sometimes be exceeded, and the noisy channel will appear at a rate in the CAL data stream that increases with the noise, i.e., with the pedestal width. Crystal occupancy (the LAC rate) is reported continuously by Data Quality Monitoring (Section 2.4) as are the pedestal widths. LAT event reconstruction (Atwood et al. 2013; Bruel et al. 2018) is robust against a modest number of CAL noise hits, so there is no immediate need to respond to noisy channels when they occur. Over the first 10 yr in orbit, the typical pedestal width has remained constant, although several channels show increased noise, detailed in Section 4.3. To accommodate the increased noise in those few cases, we have raised the LAC threshold to its maximum possible value on three channels, effectively disabling them. Because a given crystal is qualified for inclusion in the CAL data stream by the LAC discriminator on either end face, and because the energy deposited in the crystal is calculated from the geometric mean of the energy calculated at each end face, the modest number of noisy channels has no effect on the science performance of the CAL.

The FLE trigger serves only a secondary role in LAT triggering. The CAL-LO trigger-request primitive formed from the logical OR of the FLE discriminators is not allowed to open a trigger window, i.e., it cannot by itself initiate a LAT event readout. Thus, for example, a side-entering gamma-ray that misses the TKR and deposits >100 MeV in one or more crystals (but not more than 1000 MeV in any crystal) will not cause a LAT event to be generated. The CAL-LO primitive is instead used by the GEM solely to help classify an event—for example, if asserted, as one of the trigger conditions required to qualify an event as a heavy cosmic ray to be used for calibration of the four CAL gain ranges. A minor error or drift in an FLE threshold has no effect on the science performance of the CAL.

In contrast, FHE does independently initiate a LAT event readout, regardless of the state of TKR or ACD trigger primitives. Importantly, this independent trigger condition dramatically reduces the self-vetoing of high-energy gamma-rays by secondary shower particles splashing back into the ACD and preserves the effective area of LAT at high energies. An error or drift in an FHE threshold, or a disabling of an FHE discriminator, can therefore affect LAT performance. However, the high redundancy provided by the segmentation of the CAL into 1536 crystals with 3072 FHE discriminators means that the electromagnetic shower of a high-energy gamma-ray typically has many opportunities to register >1000 MeV in a single discriminator.

The range-selection upper-level discriminator (ULD) settings have not been changed since launch. This has no effect on the CAL performance. As the light yield has declined with time, the equivalent energy at which a given CAL front end

presents a particular gain range for digitization has increased, but the gain scales remain calibrated, and the event energy is properly reconstructed.

The values of the above parameters were set before launch. They are checked every 6 Ms. The change in light yield is small enough that gamma-ray detection, filtering, and subsequent transmission to the ground are unaffected. Thus, no settings changes have been made since launch. The light yield changes are updated in the ground calibrations so that the final physical quantities (energy, longitudinal position, and so on) are stable.

4.3. Calorimeter Component Performance

During the 10 yr of on-orbit operations, only a single hardware component has failed in CAL, a single preamplifier among the 12,288 in the CAL front-end electronics. Signals from the high-energy photodiode from the corresponding crystal end face stopped being received in 2010 July, with no increase in pedestal noise visible in periodic trigger events. However, by 2017 August, the noise had increased enough to solicit an FHE trigger request with some frequency. This channel (Tower 4, Layer 2, Column 4, at the $+X$ end) appears above 800 MeV in the right-hand panel of Figure 20. At that time, we disabled automatic range selection and FHE triggers from the affected front-end electronics chip. No root cause for the failure has been identified. Reconstruction of any event that showers in the calorimeter and either misses this crystal or deposits less than 1 GeV in this crystal is unaffected by this failure. The overwhelming majority of LAT events is thus utterly unaffected by this failure. For a high-energy event that showers in the calorimeter and deposits more than 1 GeV of energy in this crystal, the reconstructed event energy and direction will be incorrect. The magnitude of the errors on reconstructed total event energy and direction has not been estimated but it is expected to be quite small.

Of the 3072 LEX8 pedestal widths in the left-hand panel of Figure 20, 12 are outside of the body of the distribution and 5 are $>2\times$ the average value, after 10 yr in orbit. Before launch, there were four out-of-family LE channels. Besides the failed channel discussed above, only one HE channel is out-of-family. These noisy channels have no discernible effect on the quality of the science data.

5. Anticoincidence Detector

The Anticoincidence Detector (ACD) is the first and foremost means of rejection of the predominant background of charged cosmic rays detected by the LAT. Cosmic ray detections outnumber the gamma-ray signal by 3–5 orders of magnitude across the Fermi-LAT’s energy range. As charged particles pass through individual ACD scintillators, they deposit energy via ionization which is converted to optical light and transmitted via wavelength shifting (WLS) fibers to ACD detector elements. Signals in the ACD are used in both the onboard event veto and ground-based event reconstruction for discriminating between the charged cosmic-ray background and the desired gamma-ray signal.

The ACD is composed of 89 plastic scintillator tiles and eight scintillating fiber ribbons, each component being read by two photomultiplier (PMT) tubes, for a total 194 channels. The dynamic range of each PMT is enhanced by a double readout with different gains for a low range and high range, for a total

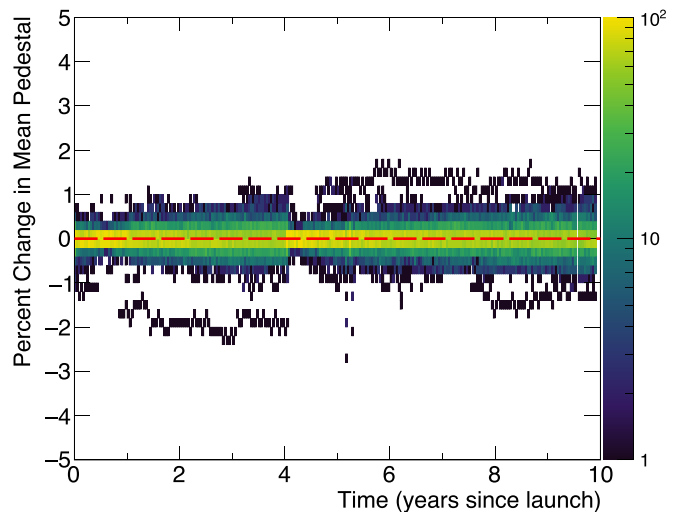


Figure 21. Trending of the ACD low-range pedestal. The x-axis is in units of years since launch. The y-axis is in units of percent change from the reference calibration. The z-axis (color bar) is the number of PMT channels. The red dashed line shows zero change, to guide the eye. The jump in pedestal values at the four-year mark is typical of ACD calibration updates.

of 388 ranged readouts. The ACD covers five sides of the Fermi-LAT extending far enough to cover the entire TKR subsystem. The segmentation of the ACD is used to minimize signal variation across the large area of the LAT and self-veto due to “backsplash” electrons and positrons induced by gamma-rays propagating upwards through the TKR and ACD after interacting with the CAL subsystem. The ACD particle detection efficiency, estimated through a combination of ground measurements and simulations, exceeds 0.9997.

The ACD can issue two trigger primitives. One signals the passage of an ionizing particle, with a low discriminator value typically used as a veto in defining the gamma trigger. The second indicates a heavily ionizing particle event, useful for calibrating the instrument. For more details see Moiseev et al. (2007).

5.1. ACD Calibration

The calibration constants for the ACD system are actively monitored and evaluated biweekly. The ACD has proven to be very stable over 10 yr of operation and there has not been a need, as of yet, to update the monitored calibration constants on board; the calibration constants used for ground processing have been updated once in 2012 September. For the ACD detector elements, the scintillator tiles and ribbons, four calibrations are actively monitored and are required to determine energy deposited: pedestals, coherent noise, low-range gain, and high-range gain. Detailed descriptions of each of these calibrations can be found in Abdo et al. (2009).

The electronic readout pedestals define the zero of the energy scale for each channel. The ACD low-range pedestal is measured using a 2 Hz periodic trigger for which a single orbit provides over 10,000 samples. The ACD high-range pedestal requires specialized charge injection and therefore cannot be monitored as actively. Figure 21 shows the low-range pedestal peak trending of all 194 channels over 10 yr. Some channels show a few-percent drift of either low-range or high-range pedestals, but nevertheless, the on-orbit calibrations are within an overall 5% allowable drift specified for the ACD. The effect of the updated calibrations for ground data processing can barely be seen 4 yr after launch.




Figure 22. Comparison of the ACD channel pedestal values for a reference run in the early mission with a run 10 yr into the mission (left) and the pedestal value ratios for the two runs (right). Channels associated with ACD tiles are shown in black while channels associated with ACD ribbons are shown in red.




Figure 23. Trending of low-range ACD gain peak (left) and the low-range ACD gain width (right). The x -axes are in units of years since launch. The y -axes are in units of percent change from the respective reference calibration. The red dashed lines show zero change to guide the eye. The jump in gains at the four-year mark is typical of ACD calibration updates.

Figure 22 compares the reference pedestal calibration constants and the pedestal calibration constants after 10 yr. Clearly, the ACD low-range pedestals are and have remained stable over the 10 yr of operation. The ACD low-range and high-range pedestals are highly correlated with the ACD temperature and during the safe mode event in 2018 March, several channels took days to return to pre-safe mode values.

Coherent noise is modeled and corrected for in an offline analysis to better estimate the energy deposition in the ACD elements. This is used specifically for the low range of the ACD readout and represents a small correction of up to 0.5% increase or decrease in the energy deposited in a single ACD detector element. The three variables used to measure the coherent noise, amplitude, decay, and phase, show little change over 10 yr. Amplitude and phase show almost no change over 10 yr, under 1% drift. Decay lifetime shows slightly more drift, about 2% over 10 yr.

The gain represents the conversion from the measured signal, in pulse height amplitude (PHA) units, to units of minimum ionizing particles (MIPs) and equivalently deposited

energy for each channel. The ACD low-range gain is measured using cosmic-ray protons with a confidently reconstructed track pointing toward the ACD tile or ribbon which recorded the signal. The ACD high-range gain is measured using cosmic-ray carbon events identified within the first three layers of the deposited energy in the CAL and having a confidently reconstructed track. The signal in either range is then pedestal subtracted and path-length corrected and fit with a Gaussian distribution to represent the signal and a line to represent a background of poorly reconstructed events. The peak of the Gaussian distribution is used as the on-orbit gain for that respective channel and range. The gain measurement is only accurate within 5% and is not meant to be a precise measurement of all factors in the energy deposited in an ACD scintillator element.

Figure 23 shows the low-range gain peak and width trending of all 194 channels over 10 yr. The effect on the gain spread due to the updated calibrations for ground data processing is particularly evident. Figure 24 compares the reference low-gain calibration constants from the early science mission, for all the

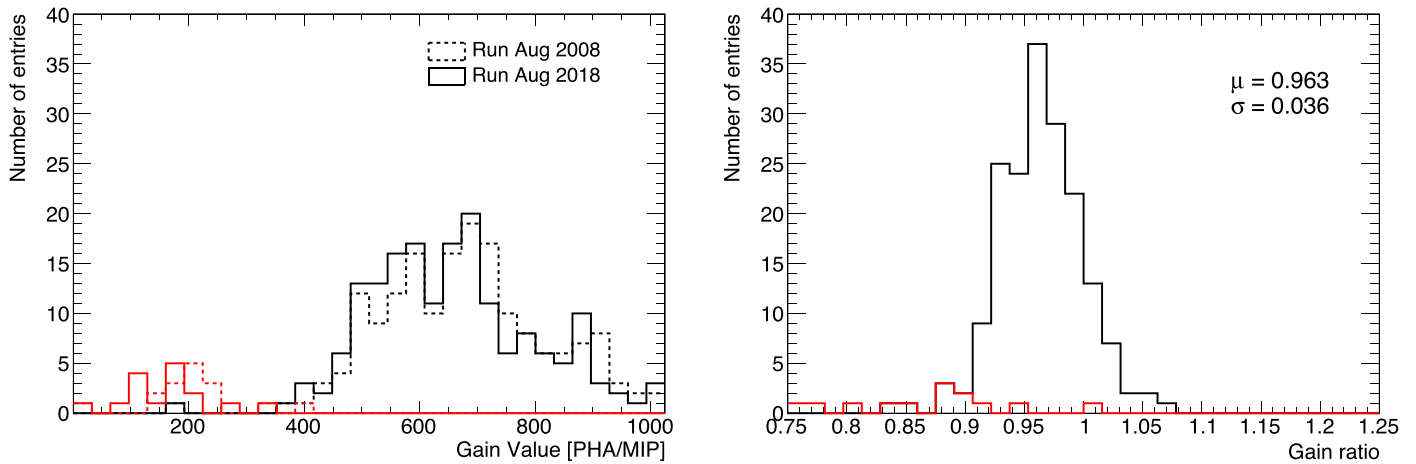


Figure 24. Comparison of the ACD channel gains for a reference run in the early mission with a run 10 yr into the mission (left) and the gain ratios for the two runs (right). Channels associated with ACD tiles are shown in black while channels associated with ACD ribbons are shown in red.

ACD channels, and the equivalent low-gain calibration constants from late in the 10 yr mission. Both Figures 23 (left) and 24 (right) clearly show an overall average reduction in channel gains, with an overall slow change of about $\sim 3\%$ over the 10 yr of the mission. Nevertheless, that average calibration change is still within the required overall 10% change limit, which represents less than 0.1% in the particle detection efficiency for single ACD detector elements.

5.2. ACD Configurations

The ACD pedestals described in Section 5.1 are part of the ACD configuration. The two remaining parts of the ACD configuration are the veto and high-energy discriminator, which define the minimum energy to issue a low-energy and high-energy trigger primitive. These calibrations are set on board and measured on the ground using events specifically collected and downlinked for diagnostics, described in Section 3.1.2 of Ackermann et al. (2012). For the low range, we measure the 50% efficiency in PHA units between events with and without a veto trigger.

For the high range, this configuration is complicated by the fact that the high-energy discriminator is OR'ed between 18 ACD channels. As a result, a given channel can have its associated high-energy discriminator bit set even if that channel did not trigger the high-energy discriminator. The process is similar to that of the veto discriminator: we measure the 95% efficiency in PHA units between events with and without a high-energy discriminator trigger. The 95% efficiency is used to compensate for the multiplexed signals. These values are set to 0.4 MIPs for the low range and 25 MIPs for the high range.

5.3. ACD Component Performance

Prior to launch two ACD channels were identified during testing as having larger than average noise and one appeared to be dead. One noisy channel is associated with a side bottom tile and the other is associated with a ribbon. These two channels are operational but have been running with reduced performance since launch. The dead one is associated with a ribbon; most likely this is a loose optical contact between the ribbon and the PMT, which occurred during one of the several thermal tests during integration. Since the ribbon can be read from the other end this does not affect the overall performance of the ACD.

Table 4
Configurations of the Trigger System

Configuration	Dimensions	Updates
Time coincidence window	1	None
Fast trigger delays	1 per subsystem	None
Latch delays	1 per subsystem	None

After 10 yr of operation, there have been no component failures within the ACD in either the detector or the electronic readout chain. The ACD is being monitored for degradation of both detector elements and the electronics readout chain through continuous observations of its low-range pedestals and gains via the active data monitoring system and through biweekly recalculation of calibration constants.

6. Trigger, Timestamp Accuracy, South Atlantic Anomaly

In addition to the detectors, the LAT's overall performance depends on some other subsystems. These include the trigger electronics and computers, the accuracy and stability of the system clocks, and the response and resilience of all the electronics to the hostile charged-particle environment around Fermi, especially during transits through the SAA region. Below we discuss these elements.

6.1. Trigger and Readout

The trigger system relies on several configurations, listed in Table 4, to ensure optimal behavior of the LAT readout. All are time delays: because the system clock runs at 20 MHz, the granularity is 50 ns (1 tick). Optimal delays were set on the ground, and updated on orbit in the initial checkout phase of the mission, but not updated since then (Abdo et al. 2009).

Time coincidence window specifies the amount of time the system waits while collecting trigger requests after the first one received starts the procedure. The value was set to 700 ns during the initial calibration phase after launch and has not been changed since. Trigger signals from TKR, CAL and ACD to the GEM follow different time paths. One *Fast Trigger Delay* per subsystem synchronizes them. On-orbit delays were set to 750 ns for ACD, 200 ns for TKR, and none for CAL, and never changed. Once a trigger is issued, the optimal time delays

for readout, corresponding to maximum amplitudes in the detectors, are set by the *Latch Delays*. On-orbit delays were set to 200 ns for ACD, 2450 ns for TKR, and 2500 ns for CAL, and never changed. The dependence of the overall efficiency on the latch delays varies little around the optimal point, e.g., a shift of ± 10 ticks, or 500 ns, around the optimal latch delay for the CAL causes a decrease in the MIP peak amplitude of only 1%.

As illustrated in Figure 5, the long-term average trigger rate is steady overall, although subject to environmental changes (see Section 2.1). At the other end of the data chain, the steady flux of bright pulsars such as Vela reported in the 4FGL catalog further demonstrates that the trigger efficiency has not deteriorated since launch (Abdollahi et al. 2020).

6.2. Timing

Fermi has a second scientific instrument in addition to the LAT: the Gamma-ray Burst Monitor (GBM), designed to have an all-sky field of view, except as occulted by Earth, to optimally detect gamma-ray burst (GRB) transient events. The GBM detected the gamma-ray burst GRB 170817A 1.74 ± 0.05 s after the LIGO gravitational wave experiment detected gravity waves from GW170817 (Abbott et al. 2017). The time lag between the GW event and the GRB is most likely due to the time required for particle acceleration and gamma-ray emission to occur, but were the lag due to a difference between the speed of gravity and the speed of light, c , that difference would be between $-3 \times 10^{-15}c$ and $+0.7 \times 10^{-15}c$. The combined LIGO/VIRGO GW detectors system should soon start detecting dozens of GW events per year, some of which will coincide with GRBs. Should one of these cataclysms yield near-simultaneous GW and GRB arrivals at Earth, confidence in Fermi's clocks could become an important element in how an observed lag is interpreted.

Section 9 of Abdo et al. (2009) describes prelaunch measurements of the accuracy of the timestamps assigned to LAT events. Atmospheric muons traversing both the LAT and a standalone muon detector allowed us to compare times recorded using a standalone GPS receiver with those recorded by the LAT. The timestamps were shown to be accurate to ~ 300 ns rms. The least significant digit of timestamp data in the data files used for astronomical analyses is $1 \mu\text{s}$, contributing an rms of 300 ns to the pulsar results described below.

To compare an arrival time at the satellite with an event time at a different observatory, Fermi's position must also be known accurately. The satellite's Guidance & Navigation Control system maintains an orbit model, updated periodically with positions provided by one of its two (redundant) Viceroy GPS receivers.⁸⁴ GPS positions close to the predicted location receive higher weights, applied with a Kalman filter. The solution filters anomalous GPS positions, and propagates the position during GPS outages, yielding results more accurate and robust than individual GPS measurements. Telemetry monitoring shows position rms residuals < 20 m, meaning that the additional uncertainty introduced when propagating a Fermi time to another location is negligible. Here, we review telemetry data and pulsar data for the 10 yr in orbit and confirm the timestamp performance. GBM timing derives from

the same spacecraft signals as the LAT timing. Prelaunch muon measurements also showed that LAT and GBM dates agree to better than $2 \mu\text{s}$ (Meegan et al. 2009). LAT's timing accuracy and stability imply that GBM's on-orbit event times are also reliable.

6.2.1. Telemetry Monitoring of the Clock Stability

The Fermi GPS receivers generate pulse per second (PPS) signals to mark the instant of validity of their time tone message (TTM), an integer number of seconds since a reference time. The PPS signals transit through the spacecraft's UDL (Up-Down Link) module before being sent to the GBM and to the LAT GEM (Global Electronics Module). The UDL PPS is generated by a 10 MHz oven-controlled crystal oscillator; therefore its rms accuracy is of order $100 \text{ ns}/\sqrt{12} = 30 \text{ ns}$. Scalers in the GEM count 50 ns "ticks" from a nominally 20 MHz oscillator. The number of ticks N_{PPS} between successive PPS signals is stored in the data streams, to provide the oscillator's true average frequency during each second. A LAT event trigger latches a GEM scaler counting the same 20 MHz ticks. A LAT timestamp is thus the TTM plus the fraction of a second given by the number of ticks since the most recent PPS. The overall timestamp accuracy is that of the PPS convolved with the variations in GEM trigger time that depend on event geometry: which TKR towers were crossed, whether the CAL participated in the trigger, and so on. The quantities described here—the measured oscillator rate, GPS lock, and others—are part of the instrument telemetry stream and are included in the DQM described in Section 2.4.

Several times per day, the Viceroy's lose "lock," sometimes for several minutes, generally because of where the GPS satellites appear in the sky relative to the receiver antennae. The UDL PPS signal specification is to drift by no more than $\pm 0.6 \mu\text{s}$ (± 12 GEM oscillator ticks) per minute during GPS lock loss. DQM of N_{PPS} is sensitive at the level of several ticks. N_{PPS} varies little, on timescales of days and weeks. Were a putative UDL PPS drift to exceed specifications, it would likely be noticed. We note that the half-day $13 \mu\text{s}$ mishap⁸⁵ in the GPS system of 2016 January 26 was reported in the LAT DQM system, and the data were flagged accordingly.

6.2.2. Timestamp End-to-end Tests Using Pulsars

With pulsars, we can validate the complex timing chain by comparing measured gamma-ray arrival times with predictions derived from radio telescope observations of the same pulsars. This amounts to comparing LAT clocks with ground-based observatory clocks. We perform two tests, for clock stability and for absolute timing. The high rotational frequency F of millisecond pulsars (MSPs) tests clock stability. Gamma-ray photon statistics for MSPs limit us to $1 \mu\text{s}$ accuracy for 1 yr data integrations, but confirm nevertheless that the timestamps cannot often be much worse than our published 300 ns rms value. The high braking rate $\dot{F} = \frac{dF}{dt}$ of the young Crab pulsar constrains the LAT absolute time.

Figure 25 summarizes our MSP results. Of the pulsars described by Abdo et al. (2013), hereafter 2PC, we chose those with short spin periods $P < 6$ ms and large H-test pulsed test statistic values (de Jager & Büsching 2010), meaning that they are bright in gamma-rays and have narrow gamma-ray peaks.

⁸⁴ General Dynamics Viceroy-4 Global Positioning System (GPS) Spaceborne Receiver.

⁸⁵ <https://www.navcen.uscg.gov/pdf/gps/AirForceOfficialPressRelease.pdf>

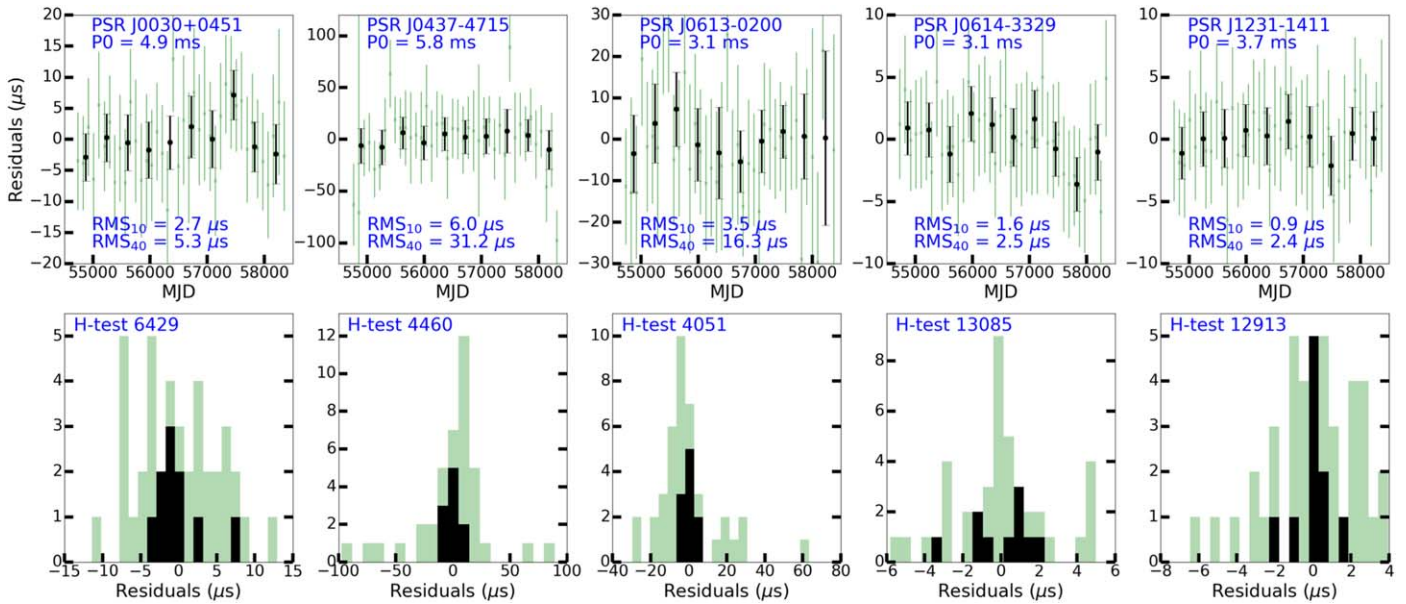


Figure 25. Stability of gamma-ray millisecond pulse times compared to radio references. Top: black dots show ten one-year average values, and green crosses show 40 three-month average values. The rms variations are dominated by the statistical uncertainties in both the gamma-ray and radio data. Bottom: the same points, displayed as histograms. The H-test indicates the statistical significance of the pulsed signal (see text).

The Nançay radio telescope (e.g., Guillemot et al. 2016) provided accurate rotation ephemerides, except for PSR J0437–4715, observed by the Parkes radio telescope (Reardon et al. 2016). We used the ephemerides to calculate a rotational phase for each gamma-ray event with the `fermi` plug-in (Ray et al. 2011) to Tempo2 (Hobbs et al. 2006). We then determined gamma-ray times-of-arrival (ToAs) using the methods described by Kerr et al. (2015). Tempo2 calculated the residuals shown in Figure 25, the differences between the ToAs predicted by the ephemerides and those measured with the LAT.

We repeated this twice for each MSP: once for gamma-ray pulse profiles integrated over a full year of LAT data, giving 10 ToAs, and then again for 40 ToAs. We expect $\text{rms}_{40} \approx 2\text{rms}_{10}$, roughly confirmed in Figure 25. The pulsars with the largest H-test values yield the smallest RMSs. With the observed jitter being dominated by statistics, we conclude that the brightest two pulsars best constrain LAT timing. These are PSRs J0614–3329 and J1231–1411. For both, there are about 43,000 photons with energy >100 MeV within 2° of the pulsar position for 10 yr of data, and the unweighted H-test value is about 14,000. The annual timing variance is of the order of a microsecond, meaning that any contribution from LAT clock jitter is likely to be less than $1/\sqrt{2}$ of that.

A less constraining test is provided by PSR J1959+2048, interesting nevertheless because the analysis is simpler. One of the fastest known MSPs ($P = 1.61$ ms), it also has the narrowest known gamma-ray peak. With only 3 yr of LAT data, Guillemot et al. (2012) found a full width at half maximum of $23 \pm 11 \mu\text{s}$. We repeated the study using 10 yr of data and a Nançay rotation ephemeris and obtained $26_{-6}^{+8} \mu\text{s}$. The uncertainty has decreased in proportion to the larger data sample, and the width agrees with the earlier measurement. The narrow width excludes timing “accidents” of tens of microseconds for any substantial fraction of the LAT’s observation time. Unfortunately, the background level around PSR J1959+2048 is too high to allow more than a few accurate LAT ToAs.

The above observations constrain putative variations in the LAT timestamps but not their absolute accuracy. Postulating

that the radio and gamma-ray signals for PSRs J1939+2134 and J1959+2048 are emitted from the same regions near the neutron stars, at the same time, allows a test. $P = 1.6$ ms for both MSPs, so the 2PC gamma-ray peak phases of -0.003 ± 0.005 for both pulsars correspond to time offsets of $\delta = -5 \pm 8 \mu\text{s}$, which could then be interpreted as a worst-case estimate of the LAT’s on-orbit absolute accuracy. However, it requires an assumption about the magnetospheric emission.

To understand how the Crab’s high braking rate constrains LAT clock accuracy without such an assumption, consider an idealized pulsar whose spindown is described only by F and \dot{F} . In a time T , it will turn $N = FT + \frac{1}{2}\dot{F}T^2$ times. If that time is wrong by ε , giving recorded timestamp values $T' = T + \varepsilon$, we would count $N' = FT' + \frac{1}{2}\dot{F}T'^2$ turns. The observed phase of a gamma-ray peak relative to a radio reference would then drift as $(N' - N) = \dot{F}T\varepsilon$, and the drift in time would be $\delta(T) = P\dot{F}T\varepsilon$. Of pulsars known to be bright in gamma-rays, the Crab (PSR J0534+2200) has $P = 33.6$ ms and $\dot{F} = -3.711 \times 10^{-10} \text{ s}^{-2}$ and thus the largest value of $P\dot{F}$. A 10 ms clock error would cause a $\approx 40 \mu\text{s}$ peak drift in 10 yr. Modeling the Crab’s spindown evolution requires more parameters than just F and \dot{F} , but Tempo2 with an accurate rotation ephemeris confirmed the adequacy of the simple expression for $\delta(T)$ for the present purposes.

Figure 26 shows that $\varepsilon < 3$ ms, which we now explain. At left, we find the time lag δ between the main gamma-ray peak and the radio reference peak, for 110 one-month data integrations (the rotation ephemeris covers a bit less than the full 10 yr). The mean, irrelevant for this test, is set to $\delta = -148 \mu\text{s}$ as per the erratum⁸⁶ to Abdo et al. (2010). At right, we show the chi-squared value of these points relative to a series of values $-8 < \varepsilon < 8$ ms. The offset of the minimum value of chi-squared from zero varies by ± 2 ms depending on

⁸⁶ <https://fermi.gsfc.nasa.gov/ssc/data/access/lat/ephems/0534+2200/NewEphem/README>

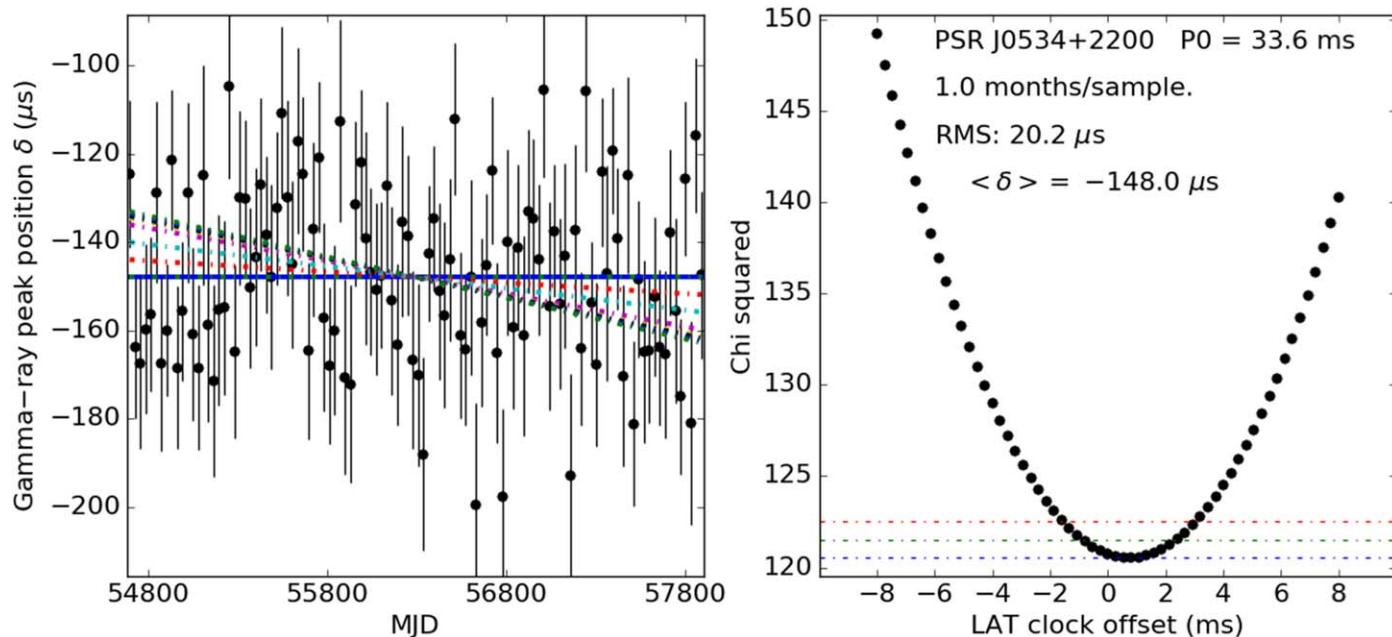


Figure 26. Accuracy of Crab gamma-ray pulse arrival times compared to a radio reference. Left: black dots show monthly average arrival times. The solid horizontal line shows the mean offset $\langle \delta \rangle = -148 \mu\text{s}$ between the gamma-ray peak and the main radio pulse. Dashed diagonals show how the mean would vary in the presence of a putative LAT clock offset relative to UTC. The steepest slope shown corresponds to δ drifting by $30 \mu\text{s}$ over 10 yr. For clarity, only a sample of negative slopes is shown. Right: each black dot is the chi-squared of the monthly pulse arrival times compared to one of the diagonals. Slopes have been converted to putative LAT clock offsets (see text). The horizontal dashed lines show the minimum chi-squared value plus 0, 1, and 2, showing $\varepsilon < 3 \text{ ms}$ (1σ limit).




Figure 27. LAT computer memory errors inside and outside the SAA region defined for the LAT (shown by the green polygon).

details of the radio rotation ephemeris used to calculate the gamma-ray phases.

The spacecraft GPS cannot be so far off: for $\varepsilon = 1 \text{ ms}$, $\varepsilon c = 300 \text{ km}$, incompatible with the $< 20 \text{ m}$ position accuracy cited at the beginning of Section 6.2.2. The integer part of the LAT time (the TTM seconds) is thus correct, since a 1 s error would induce a huge drift in the Crab peak compared to the radio peak. The fractional part derives from the well-monitored scaler counts of the 20 MHz oscillator between successive PPS signals, and is thus also reliable. The prelaunch ground tests showed mean offsets of $\pm 300 \text{ ns}$ between the LAT and simple

GPS times, varying from run to run. We attribute this to inaccuracies in the simple GPS but have not demonstrated it unequivocally. In summary, the onboard timestamp measurements at the few-microsecond level indicate that the clocks are performing as designed, confirming the published 300 ns accuracy.

6.3. South Atlantic Anomaly

For mission planning and operations purposes, the boundary of the South Atlantic Anomaly (SAA) is defined by 12-vertex




Figure 28. Counts of LAT computer memory errors in successive 10 million second time intervals during the Fermi mission. The first data points in 2018 are lower than expected because that time bin includes the 17 days power-off period for the LAT in 2018 March and April.

polygons defined separately for the LAT and GBM, specified in terms of geocentric latitude–longitude pairs. These polygons denote the geospatial location of the SAA, a region of space where geomagnetically trapped charged-particle flux densities increase dramatically. The LAT is not able to collect useful science event data inside the SAA region because of excessive dead time, and also because the anticoincidence rejection of charged particles is effectively lost, as the high voltages are reduced on the PMTs that read out the ACD to avoid excessive current draw off of the PMT photocathodes and therefore preserve the long-term performance of the PMTs. Therefore, the SAA boundary location is critical to the mission planning process for defining the LAT science data acquisition start and stop commanding. The FSSC and MOC use the LAT SAA boundary polygon to supply predicted SAA entry and exit times, with the corresponding LAT stop and start commands planned to be outside the predicted SAA by 30 s, to allow for uncertainty between the predicted times and the actual times. The polygon definition is also uploaded to Fermi and used there to supply the “in SAA” flag from the spacecraft flight software to the LAT flight software.

Another notable effect of the SAA on the LAT is an associated increase in memory errors in the LAT onboard computers. Figure 27 shows the positions of Fermi when memory cell upsets in the memory electronics of the onboard LAT computers occurred, during the first seven months of 2018. It shows that over 90% of the LAT memory errors occurred while Fermi was passing through the SAA region, which occupies about 15% of the total Fermi orbit time. This clearly demonstrates the high density of trapped charged particles within the SAA region.

Figure 28 shows the number of LAT onboard computer memory cell upsets in successive 10 million second time intervals over the 10 yr mission. The slow variation seen is another likely indication of the influence of the 11 yr solar cycle on the density of trapped charged particles in the SAA,

with the density being reduced during the period of maximum solar activity, which peaked in 2013–2014.

7. Conclusion

The LAT remains in excellent operating condition after more than 10 years in space. The Fermi spacecraft is also in good operating condition with no performance concerns. There are no consumables that will limit the lifetime of the LAT or Fermi. In Fermi’s 13th year in orbit, continued monitoring confirms the LAT’s ongoing smooth operation.

The LAT and the Fermi mission both continue to produce excellent science. The most recent NASA Senior Review in 2019 highlighted their crucial roles in the extremely important new frontiers of multimessenger astrophysics that were opened by Fermi and the LAT. A gamma-ray burst that was detected by GBM and followed up by the LAT was coincident with the first neutron star–neutron star merger observed by its gravitational radiation with the LIGO and Virgo detectors (Abbott et al. 2017). The LAT detected a flaring of gamma-ray emission from an active galaxy in association with the detection of a high-energy neutrino from the same direction by the IceCube neutrino detector, marking the first association between neutrino and photon emission from an active galaxy (IceCube Collaboration et al. 2018). These ground-breaking multimessenger science results herald the important science still to come from Fermi and the LAT, as multimessenger capabilities and sensitivities continue to improve.

In addition to multimessenger science, the LAT and Fermi are expected to centrally participate in coming major advances in multiwavelength studies. New observatories such as the Rubin Observatory⁸⁷ and the Cherenkov Telescope Array⁸⁸ are either nearing completion or in active development and will or may overlap with the Fermi mission. Data from these new observatories will leverage further use of both LAT survey data

⁸⁷ <https://project.lsst.org/>

⁸⁸ <https://www.cta-observatory.org/>

and results, and synergize with the detection of transient sources by both the LAT and by the other telescopes.

The continued good performance of the LAT, and the new era of multimessenger science that includes important contributions from Fermi and the LAT, plus further enhancements by new multiwavelength telescopes and sky surveys coming in the near future, should produce outstanding science for the foreseeable future.

The Fermi-LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the United States, the Commissariat à l'Énergie Atomique and the Centre National de la Recherche Scientifique/Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden.

Additional support for science analysis during the operations phase is gratefully acknowledged from the Istituto Nazionale di Astrofisica in Italy and the Centre National d'Études Spatiales in France. This work was performed in part under DOE Contract DE-AC02-76SF00515.

ORCID iDs

M. Ajello <https://orcid.org/0000-0002-6584-1703>
 M. Axelsson <https://orcid.org/0000-0003-4378-8785>
 L. Baldini <https://orcid.org/0000-0002-9785-7726>
 D. Bastieri <https://orcid.org/0000-0002-6954-8862>
 R. Bellazzini <https://orcid.org/0000-0002-2469-7063>
 E. Bissaldi <https://orcid.org/0000-0001-9935-8106>
 R. Bonino <https://orcid.org/0000-0002-4264-1215>
 J. Bregeon <https://orcid.org/0000-0002-6790-5328>
 R. A. Cameron <https://orcid.org/0000-0003-0942-2747>
 P. A. Caraveo <https://orcid.org/0000-0003-2478-8018>
 S. Chen <https://orcid.org/0000-0003-1586-3653>
 C. C. Cheung <https://orcid.org/0000-0002-4377-0174>
 S. Ciprini <https://orcid.org/0000-0002-0712-2479>
 I. Cognard <https://orcid.org/0000-0002-1775-9692>
 S. Cutini <https://orcid.org/0000-0002-1271-2924>
 F. D'Ammando <https://orcid.org/0000-0001-7618-7527>
 P. de la Torre Luque <https://orcid.org/0000-0002-4150-2539>
 N. Di Lalla <https://orcid.org/0000-0002-7574-1298>
 L. Di Venere <https://orcid.org/0000-0003-0703-824X>
 A. Domínguez <https://orcid.org/0000-0002-3433-4610>
 E. C. Ferrara <https://orcid.org/0000-0001-7828-7708>
 A. Fiori <https://orcid.org/0000-0003-3174-0688>
 Y. Fukazawa <https://orcid.org/0000-0002-0921-8837>
 P. Fusco <https://orcid.org/0000-0002-9383-2425>
 F. Gargano <https://orcid.org/0000-0002-5055-6395>
 D. Gasparri <https://orcid.org/0000-0002-5064-9495>
 M. Giroletti <https://orcid.org/0000-0002-8657-8852>
 T. Glanzman <https://orcid.org/0000-0001-9649-3871>
 M.-H. Grondin <https://orcid.org/0000-0002-8383-251X>

L. Guillemot <https://orcid.org/0000-0002-9049-8716>
 S. Guiriec <https://orcid.org/0000-0001-5780-8770>
 E. Hays <https://orcid.org/0000-0002-8172-593X>
 G. Jóhannesson <https://orcid.org/0000-0003-1458-7036>
 M. Kerr <https://orcid.org/0000-0002-0893-4073>
 M. Kuss <https://orcid.org/0000-0003-1212-9998>
 S. Larsson <https://orcid.org/0000-0003-0716-107X>
 L. Latronico <https://orcid.org/0000-0002-0984-1856>
 M. Lemoine-Goumard <https://orcid.org/0000-0002-4462-3686>
 J. Li <https://orcid.org/0000-0003-1720-9727>
 F. Longo <https://orcid.org/0000-0003-2501-2270>
 F. Loparco <https://orcid.org/0000-0002-1173-5673>
 P. Lubrano <https://orcid.org/0000-0003-0221-4806>
 S. Maldera <https://orcid.org/0000-0002-0698-4421>
 A. Manfreda <https://orcid.org/0000-0002-0998-4953>
 M. N. Mazziotta <https://orcid.org/0000-0001-9325-4672>
 I. Mereu <https://orcid.org/0000-0003-0219-4534>
 M. Meyer <https://orcid.org/0000-0002-0738-7581>
 W. Mitthumsiri <https://orcid.org/0000-0002-3776-072X>
 T. Mizuno <https://orcid.org/0000-0001-7263-0296>
 M. E. Monzani <https://orcid.org/0000-0002-8254-5308>
 I. V. Moskalenko <https://orcid.org/0000-0001-6141-458X>
 M. Negro <https://orcid.org/0000-0002-6548-5622>
 E. Nuss <https://orcid.org/0000-0002-3528-5777>
 V. S. Paliya <https://orcid.org/0000-0001-7774-5308>
 J. S. Perkins <https://orcid.org/0000-0001-9608-4023>
 M. Pesce-Rollins <https://orcid.org/0000-0003-1790-8018>
 H. Poon <https://orcid.org/0000-0002-6004-4270>
 T. A. Porter <https://orcid.org/0000-0002-2621-4440>
 J. L. Racusin <https://orcid.org/0000-0002-4744-9898>
 S. Rainò <https://orcid.org/0000-0002-9181-0345>
 R. Rando <https://orcid.org/0000-0001-6992-818X>
 B. Rani <https://orcid.org/0000-0001-5711-084X>
 M. Razzano <https://orcid.org/0000-0003-4825-1629>
 S. Razzaque <https://orcid.org/0000-0002-0130-2460>
 A. Reimer <https://orcid.org/0000-0001-8604-7077>
 O. Reimer <https://orcid.org/0000-0001-6953-1385>
 P. M. Saz Parkinson <https://orcid.org/0000-0001-6566-1246>
 D. Serini <https://orcid.org/0000-0002-9754-6530>
 C. Sgrò <https://orcid.org/0000-0001-5676-6214>
 D. A. Smith <https://orcid.org/0000-0002-7833-0275>
 D. J. Suson <https://orcid.org/0000-0003-2911-2025>
 D. J. Thompson <https://orcid.org/0000-0001-5217-9135>
 L. Tibaldo <https://orcid.org/0000-0001-7523-570X>
 D. F. Torres <https://orcid.org/0000-0002-1522-9065>
 G. Tosti <https://orcid.org/0000-0002-0839-4126>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *ApJL*, 848, L13
 Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, *Aph*, 32, 193
 Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, 708, 1254
 Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *ApJS*, 208, 17
 Abdollahi, S., Acero, F., Ackermann, M., et al. 2020, *ApJS*, 247, 33
 Ackermann, M., Ajello, M., Albert, A., et al. 2012, *ApJS*, 203, 4
 Atwood, W., Albert, A., Baldini, L., et al. 2013, arXiv:1303.3514
 Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, *ApJ*, 697, 1071
 Atwood, W. B., Bagagli, R., Baldini, L., et al. 2007, *Aph*, 28, 422
 Baldini, L., Brez, A., Himel, T., et al. 2006, *ITNS*, 53, 466
 Briggs, M. S., Xiong, S., Connaughton, V., et al. 2013, *JGRA*, 118, 3805
 Bruehl, P. 2012, *JPhCS*, 404, 012033
 Bruehl, P., Burnett, T. H., Digel, S. W., et al. 2018, arXiv:1810.11394

- Chowdhury, M., Watts, S., Imrie, D., McKemey, A., & Holmes-Siedle, A. 1999, *NIMPA*, 432, 147
- de Jager, O. C., & Büsching, I. 2010, *A&A*, 517, L9
- Grove, J. E., & Johnson, W. N. 2010, *Proc. SPIE*, 7732, 77320J
- Guillemot, L., Johnson, T. J., Venter, C., et al. 2012, *ApJ*, 744, 33
- Guillemot, L., Smith, D. A., Laffon, H., et al. 2016, *A&A*, 587, A109
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, 369, 655
- IceCube Collaboration, Aartsen, M. G., Ackermann, M., et al. 2018, *Sci*, 361, eaat1378
- Kazui, K., Watanabe, A., Osone, S., et al. 1997, *NIMPA*, 394, 46
- Kerr, M., Ray, P. S., Johnston, S., Shannon, R. M., & Camilo, F. 2015, *ApJ*, 814, 128
- Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, 702, 791
- Moiseev, A. A., Hartman, R. C., Ormes, J. F., et al. 2007, *Aph*, 27, 339
- Ray, P. S., Kerr, M., Parent, D., et al. 2011, *ApJS*, 194, 17
- Reardon, D. J., Hobbs, G., Coles, W., et al. 2016, *MNRAS*, 455, 1751
- Schartmann, M., Burkert, A., Alig, C., et al. 2012, *ApJ*, 755, 155
- Woody, C., Kierstead, J., Levy, P., & Stoll, S. 1992, *ITNS*, 39, 524