

Muon ($g - 2$): experiment and theory

To cite this article: James P Miller *et al* 2007 *Rep. Prog. Phys.* **70** 795

View the [article online](#) for updates and enhancements.

Related content

- [Anomalous magnetic moment of the muon](#)
I B Logashenko and S I Eidelman
- [Spin-polarized charged particle beams](#)
S R Mane, Yu M Shatunov and K Yokoya
- [Review: Standard model prediction of muon anomalous magnetic moment](#)
M Passera

Recent citations

- [Excitation probability of the hyperfine ground state of Doppler-broadened muonic hydrogen through dipole-forbidden transition](#)
Rakesh Mohan Das *et al*
- [Search for the Electric Dipole Moment and anomalous magnetic moment of the tau lepton at tau factories](#)
Xin Chen and Yongcheng Wu
- [Observation of e+e and search for +](#)
Y. Jin *et al*



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

Muon ($g - 2$): experiment and theory

James P Miller¹, Eduardo de Rafael² and B Lee Roberts^{1,3}

¹ Department of Physics, Boston University, Boston, MA 02215, USA

² Centre de Physique Théorique, CNRS-Luminy, Case 907, F-13288 Marseille Cedex 9, France

E-mail: roberts@bu.edu

Received 28 February 2007

Published 23 April 2007

Online at stacks.iop.org/RoPP/70/795

Abstract

A review of the experimental and theoretical determinations of the anomalous magnetic moment of the muon is given. The anomaly is defined by $a = (g - 2)/2$, where the Landé g -factor is the proportionality constant that relates the spin to the magnetic moment. For the muon, as well as for the electron and tauon, the anomaly a differs slightly from zero (of the order 10^{-3}) because of radiative corrections. In the Standard Model, contributions to the anomaly come from virtual ‘loops’ containing photons and the known massive particles. The relative contribution from heavy particles scales as the square of the lepton mass over the heavy mass, leading to small differences in the anomaly for e , μ and τ . If there are heavy new particles outside the Standard Model which couple to photons and/or leptons, the relative effect on the muon anomaly will be $\sim (m_\mu/m_e)^2 \approx 43 \times 10^3$ larger compared with the electron anomaly. Because both the theoretical and experimental values of the muon anomaly are determined to high precision, it is an excellent place to search for the effects of new physics or to constrain speculative extensions to the Standard Model. Details of the current theoretical evaluation and of the series of experiments that culminates with E821 at the Brookhaven National Laboratory, are given. At present the theoretical and the experimental values are known with a similar relative precision of 0.5 ppm. There is, however, a 3.4 standard-deviation difference between the two, strongly suggesting the need for continued experimental and theoretical study.

(Some figures in this article are in colour only in the electronic version)

This article was invited by Professor P Söding

³ Author to whom any correspondence should be addressed.

Contents

	Page
1. Introduction and history of g -factors	798
1.1. Muon decay	801
1.2. History of the muon ($g - 2$) experiments	803
2. The Brookhaven experiment E821: experimental technique	807
2.1. The proton and muon beamlines	810
2.2. The inflector magnet	811
2.3. The fast muon kicker	812
2.4. The electrostatic quadrupoles	815
2.5. The superconducting storage ring	816
2.6. Measurement of the precision magnetic field	818
2.7. Detectors and electronics	820
2.7.1. Electromagnetic calorimeters: design considerations.	820
2.7.2. Electromagnetic calorimeters: construction.	822
2.7.3. Special detectors: FSD, PSD, fibre harps, traceback chambers.	823
2.7.4. Electronics.	824
3. Beam dynamics	826
3.1. Monitoring the beam profile	829
3.2. Corrections to ω_a : pitch and radial electric field	830
4. Analysis of the data for ω_p and ω_a	833
4.1. The average magnetic field: the ω_p analysis	833
4.1.1. Calibration of the trolley probes.	834
4.1.2. Mapping the magnetic field.	835
4.1.3. Tracking the magnetic field in time.	836
4.1.4. Determination of the average magnetic field: ω_p .	836
4.1.5. Summary of the magnetic field analysis.	837
4.2. The muon spin precession frequency: the ω_a analysis	837
4.2.1. Distribution of decay electrons.	838
4.2.2. Modification of the five-parameter fitting function.	840
5. The permanent electric dipole moment of the muon	844
6. Results and summary of E821	847
7. The theory of the muon $g - 2$	847
7.1. The QED contributions	848
7.1.1. Diagrams with virtual photons and muon loops.	848
7.1.2. Vacuum polarization diagrams from electron loops.	850
7.1.3. Vacuum polarization diagrams from tau loops.	855
7.1.4. Light-by-light scattering diagrams from lepton loops.	857
7.2. Contributions from hadronic vacuum polarization	861
7.2.1. Higher-order hadronic vacuum polarization contributions.	865
7.3. Contributions from hadronic light-by-light scattering	868
7.4. Electroweak contributions	871
7.5. Summary of the Standard Model contributions to the muon anomaly	874

8. Future experimental possibilities	875
9. Summary and conclusions	876
Acknowledgments	877
References	877

1. Introduction and history of g -factors

A charged elementary particle with half-integer intrinsic spin has a magnetic dipole moment $\vec{\mu}$ aligned with its spin \vec{s} :

$$\vec{\mu} = g_s \left(\frac{q}{2m} \right) \vec{s}, \quad (1)$$

where $q = \pm e$ is the charge of the particle in terms of the magnitude of the electron charge e and the proportionality constant g_s is the Landé g -factor. For the charged leptons, e , μ or τ , g_s is slightly greater than 2. From a quantum-mechanical view, the magnetic moment must be directed along the spin, since in the absence of an external electric or magnetic field, the spin provides the only preferred direction in space.

While the behaviour of the spin of an individual particle must be treated quantum mechanically, the polarization, the average behaviour of the spins of a large ensemble of particles, can be treated to a large extent as a classical collection of spinning bar magnets⁴. An externally applied magnetic field, \vec{B} , will exert a torque on the magnetic moment which tends to align the polarization with the direction of the field. However, because the magnetic moments are associated with the intrinsic angular momentum, the polarization precesses in this case like a classical collection of gyroscopes in the plane perpendicular to the field. In both the classical and quantum-mechanical cases, the torque on the particle is given by $\vec{N} = \vec{\mu} \times \vec{B}$, the energy depends on the orientation of the magnetic dipole, $\mathcal{H} = -\vec{\mu} \cdot \vec{B}$, and there is a net force on the particle if the field is non-uniform.

The study of atomic and sub-atomic magnetic moments, which continues to this day, began in 1921 with the famous ‘Stern–Gerlach’ experiments [1,2]. A beam of silver atoms was passed through a gradient magnetic field to separate the individual quantum states. Two spatially separated bands of atoms were observed, signifying two quantum states. From this separation, the magnetic moment of the silver atom was determined to be one Bohr magneton [3], $e\hbar/2m_e$, to within 10%. Phipps and Taylor [4] repeated the experiment with a hydrogen beam in 1927, and they also observed two bands. From the splitting of the bands they concluded that, like silver, the magnetic moment of the hydrogen atom was one Bohr magneton. In terms of spin, which was proposed independently by Compton [5], and by Uhlenbeck and Goudsmit [6], the two-band structure indicated that the z -component of the angular momentum had two values, $s_z = \pm\hbar/2$. Later, the magnetic moments of both atoms would be traced to the un-paired spin of an atomic electron, implying (from equation (1)) that the g -factor of the electron is 2.

In 1928 Dirac published his relativistic wave equation for the electron. It employed four-vectors and the famous 4×4 matrix formulation, with the spin degree of freedom emerging in a natural way. Dirac pointed out that his wave equation for an electron in external electric and magnetic fields has ‘the two extra terms’⁵

$$\frac{eh}{c}(\boldsymbol{\sigma}, \mathbf{H}) + i \frac{eh}{c} \rho_1(\boldsymbol{\sigma}, \mathbf{E}), \quad (2)$$

which when divided by the factor $2m$ can be regarded as the additional potential energy of the electron due to its new degree of freedom’ [7]. These terms represent the magnetic dipole (Dirac) moment and electric dipole moment interactions with the external magnetic and electric fields. Dirac theory predicts that the electron magnetic moment is one Bohr-magneton ($g_s = 2$), consistent with the value measured by the Stern–Gerlach and Phipps–Taylor experiments.

⁴ Following the custom in the literature, in this article we will often refer to the motion of the ‘spin’ when more correctly we are speaking of the motion of the polarization of an ensemble of a large number of particles which have spin.

⁵ This expression uses Dirac’s original notation.




Figure 1. The Feynman graphs for: (a) $g = 2$; (b) the lowest-order radiative correction first calculated by Schwinger; and (c) the vacuum polarization contribution, which is an example of the next-order term. The * emphasizes that in the loop the muon is off-shell.

Dirac later commented: ‘It gave just the properties that one needed for an electron. That was an unexpected bonus for me, completely unexpected’ [8]. The non-relativistic reduction of the Dirac equation for the electron in a weak magnetic field [9],

$$i\hbar \frac{\partial \phi}{\partial t} = \left[\frac{(\vec{p})^2}{2m} - \frac{e}{2m} (\vec{L} + 2\vec{S}) \cdot \vec{B} \right] \phi \tag{3}$$

shows clearly that $g_s = 2$ and the g -factor for orbital angular momentum, $g_\ell = 1$.

With the success of the Dirac equation, it was believed that the proton should also have a g -factor of 2. However, in 1933 Stern and his collaborators [10] showed that the g -factor of the proton was ~ 5.5 . In 1940 Alvarez and Bloch [11] found that the neutron likewise had a large magnetic moment, which was a complete surprise, since $q = 0$ for the neutron. These two results remained quite mysterious for many years, but we now understand that the baryon magnetic moments are related to their internal quark-gluon structure.

Theoretically it is useful to break the magnetic moment into two pieces:

$$\mu = (1 + a) \frac{q\hbar}{2m} \quad \text{where} \quad a = \frac{g - 2}{2}. \tag{4}$$

The first piece, predicted by the Dirac equation and called the Dirac moment, is 1 in units of the appropriate magneton, $q\hbar/2m$. The second piece is the anomalous (Pauli) moment [12], where the dimensionless quantity a is referred to as *the anomaly*.

In 1947, motivated by measurements of the hyperfine structure in hydrogen that obtained splittings larger than expected from the Dirac theory [13–15], Schwinger [16]⁶ showed that from a theoretical viewpoint these ‘discrepancies can be accounted for by a small additional electron spin magnetic moment’ that arises from the lowest-order radiative correction to the Dirac moment.⁷ This radiative correction, which we now call the one-loop correction to $g = 2$, is shown diagrammatically in figure 1(b). Schwinger obtained the value $a_e = \alpha/(2\pi) \simeq 0.00116 \dots$ (which is also true for a_μ and a_τ). In the same year, Kusch and Foley [18] measured a_e with 4% precision and found that the measured electron anomaly agreed well with Schwinger’s prediction. In the intervening time since the Kusch and Foley paper, many improvements have been made in the precision of the electron anomaly [19]. Most recently a_e has been measured to a relative precision of 0.65 ppb (parts per billion) [20], a factor of 6 improvement over the famous experiments at the University of Washington [21].

The ability to calculate loop diagrams such as those shown in figures 1 (b) and (c) is intimately tied to the renormalizability of the theory, which provides a prescription to deal

⁶ The former paper contains a misprint in the expression for a_e that is corrected in the longer paper.

⁷ In response to Nafe, *et al* [13], Breit [17] conjectured that this discrepancy could be explained by the presence of a small Pauli moment. It is not clear whether this paper influenced Schwinger’s work, but in a footnote Schwinger states: ‘However, Breit has not correctly drawn the consequences of his empirical hypothesis.’

with the infinities encountered in calculating radiative corrections and was important to the development of quantum electrodynamics. In his original paper Schwinger [16]⁶ described a new procedure that transformed the Dirac Hamiltonian to include the electron self-energy which arises from the emission and absorption of virtual photons. By doing so, he eliminated the divergences encountered in calculating the lowest-order radiative correction. He pointed to three important features of his new Hamiltonian: ‘it involves the experimental electron mass’ (known today as the ‘dressed or physical mass’); ‘an electron now interacts with the radiation field only in the presence of an external field;’ (i.e. the virtual photons from the self-energy are absent) and ‘the interaction energy of an electron with an external field is now subject to a *finite* radiative correction’ [16]. This concept of renormalization also played an important role in the development of the Standard Model, and the lowest-order contribution from virtual W and Z gauge bosons to a_μ was calculated very soon after the electroweak theory was shown to be renormalizable [22].

The diagram in figure 1(a) corresponds to $g = 2$, and the first-order (Schwinger) correction which dominates the anomaly is given in figure 1(b). More generally, the Standard-Model value of the electron, muon or tauon anomaly, $a(\text{SM})$, arises from loops (radiative corrections) containing virtual leptons, hadrons and gauge bosons. By convention, these contributions are divided into three classes: the dominant QED terms, like Schwinger’s correction, which contain only leptons and photons; terms which involve hadrons, particularly the hadronic vacuum polarization correction to the Schwinger term; and electroweak terms, which contain the Higgs, W and Z . Some of the terms are identical for all three leptons, but as noted below, there are mass-dependent terms which are significant for the muon and tauon but not for the electron. As a result, the muon anomaly is slightly larger than that of the electron.

Both the QED and electroweak contributions can be calculated to high precision. For the muon, the former has been calculated to 8th order (4th order in α/π) and the 10th order terms are currently being evaluated (see section 7.1). The electroweak contributions have been calculated to two loops, as discussed in section 7.4.

In contrast, the hadronic contribution to a_μ cannot be accurately calculated from low-energy quantum chromodynamics (QCD) and contributes the dominant theoretical uncertainty on the Standard-Model prediction. The lowest-order hadronic contribution, determined using experimental data via a dispersion relation as discussed in section 7.2, accounts for over half of this uncertainty. The Hadronic light by light scattering contribution (see section 7.3) contributes the rest. This latter contribution cannot be related to existing data, but rather must be calculated using models that incorporate the features of QCD.

Since the value of a_μ arises from all particles that couple, directly or indirectly, to the muon, a precision measurement of the muon anomaly serves as an excellent probe for new physics. Depending on their mass and coupling strength, as yet unknown particles could make a significant contribution to a_μ . Lepton or W -boson substructure might also have a measurable effect. Conversely, the comparison between the experimental and Standard-Model values of the anomaly can be used to constrain the parameters of speculative theories [23–25].

While the current experimental uncertainty of ± 0.5 ppm (parts per million) on the muon anomaly is 770 times larger than that on the electron anomaly, the former is far more sensitive to the effects of high mass scales. In the lowest-order diagram where mass effects appear, the contribution of heavy virtual particles scales as $(m_{\text{lepton}}/m_{\text{HV}})^2$, giving the muon a factor of $(m_\mu/m_e)^2 \simeq 43\,000$ increase in sensitivity over the electron. Thus, at a precision of 0.5 ppm, the muon anomaly is sensitive to physics at a few hundred GeV scale, while the reach of the electron anomaly at the current experimental uncertainty of 0.65 ppb is limited to the few hundred MeV range.

To observe effects of new physics in a_e , it is necessary to have a sufficiently precise Standard-Model value to compare with the experimental value. Since the QED contribution depends directly on a power series expansion in the fine-structure constant α , a meaningful comparison requires that α be known *from an independent experiment* to the same relative precision as a_e . At present, independent measurements of α have a precision of ~ 6.7 ppb [27], about ten times larger than the error on a_e . On the other hand, if one assumes that there are no new physics contributions to a_e , i.e. the Standard-Model radiative corrections are the *only* source of a_e , one can use the measured value of a_e and the Standard-Model calculation to obtain the most precise determination of α [28].

1.1. Muon decay

Since the kinematics of muon decay are central to the measurements of a_μ , we discuss the general features in this section. Specific issues that relate to the design of E821 at Brookhaven are discussed in the detector section, section 2.7.1.

After the discovery of parity violation in β -decay [29] and in muon decay [30,31], it was realized that one could make beams of polarized muons in the pion decay reactions

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \quad \text{or} \quad \pi^+ \rightarrow \mu^+ + \nu_\mu.$$

The pion has spin zero, the neutrino (antineutrino) has helicity of -1 ($+1$) and the forces in the decay process are very short-ranged, so the orbital angular momentum in the final state is zero. Thus conservation of angular momentum requires that the μ^- (μ^+) helicity be $+1$ (-1) in the pion rest frame. The muons from pion decay at rest are always polarized.

From a beam of pions traversing a straight beam-channel consisting of focusing and defocusing elements (FODO), a beam of polarized, high energy muons can be produced by selecting the ‘forward’ or ‘backward’ decays. The forward muons are those produced, in the pion rest frame, nearly parallel to the pion laboratory momentum and are the decay muons with the highest laboratory momenta. The backward muons are those produced nearly anti-parallel to the pion momentum and have the lowest laboratory momenta. The forward μ^- (μ^+) are polarized along (opposite) their lab momenta, respectively; the polarization reverses for backward muons. The E821 experiment uses forward muons, produced by a pion beam with an average momentum of $p_\pi \approx 3.15$ GeV/ c . Under a Lorentz transformation from the pion rest frame to the laboratory frame, the decay muons have momenta in the range $0 < p_\mu < 3.15$ GeV/ c . After momentum selection, the average momentum of muons stored in the ring is 3.094 GeV/ c and the average polarization is in excess of 95%.

The pure ($V - A$) three-body weak decay of the muon, $\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e$ or $\mu^+ \rightarrow e^+ + \bar{\nu}_\mu + \nu_e$, is ‘self-analysing’, that is the parity-violating correlation between the directions in the muon rest frame (MRF) of the decay electron and the muon spin can provide information on the muon spin orientation at the time of the decay. (In the following text, we use ‘electron’ generically for either e^- and e^+ from the decay of the μ^\mp .) Consider the case when the decay electron has the maximum allowed energy in the MRF, $E'_{\max} \approx (m_\mu c^2)/2 = 53$ MeV. The neutrino and anti-neutrino are directed parallel to each other and at 180° relative to the electron direction. The $\nu\bar{\nu}$ pair carries zero total angular momentum, since the neutrino is left-handed and the anti-neutrino is right-handed; the electron carries the muon’s angular momentum of $1/2$. The electron, being a lepton, is preferentially emitted left-handed in a weak decay and thus has a larger probability to be emitted with its momentum *anti-parallel* rather than parallel to the μ^- spin. By the same line of reasoning, in μ^+ decay, the highest-energy positrons are emitted *parallel* to the muon spin in the MRF.

At other extreme, when the electron kinetic energy is zero in the MRF, the neutrino and anti-neutrino are emitted back-to-back and carry a total angular momentum of one. In this




Figure 2. Number of decay electrons per unit energy, N (arbitrary units), value of the asymmetry, A , and relative figure of merit, NA^2 (arbitrary units), as a function of electron energy. Detector acceptance has not been incorporated, and the polarization is unity. For the third CERN experiment and E821, $E_{\max} \approx 3.1$ GeV ($p_{\mu} = 3.094$ GeV/ c) in the laboratory frame. (a) Muon rest frame and (b) laboratory frame.

case, the electron spin is directed opposite to the muon spin in order to conserve angular momentum. Again, the electron is preferentially emitted with helicity -1 , however in this case its momentum will be preferentially directed *parallel* to the μ^- spin. The positron, in μ^+ decay, is preferentially emitted with helicity $+1$ and therefore its momentum will be preferentially directed *anti-parallel* to the μ^+ spin.

With the approximation that the energy of the decay electron $E' \gg m_e c^2$, the differential decay distribution in the muon rest frame is given by [32],

$$dP(y', \theta') \propto n'(y')[1 \pm A(y') \cos \theta'] dy' d\Omega', \quad (5)$$

where y' is the momentum fraction of the electron, $y' = p'_e / p'_{e \max}$, $d\Omega'$ is the solid angle, $\theta' = \cos^{-1}(\hat{p}'_e \cdot \hat{s})$ is the angle between the muon spin and \vec{p}'_e , $p'_{e \max} c \approx E'_{\max}$ and the $(-)$ sign is for negative muon decay. The number distribution $n(y')$ and the decay asymmetry $A(y')$ are given by

$$n(y') = 2y'^2(3 - 2y') \quad \text{and} \quad A(y') = \frac{2y' - 1}{3 - 2y'}, \quad (6)$$

Note that both the number and asymmetry reach their maxima at $y' = 1$, and the asymmetry changes sign at $y' = \frac{1}{2}$, as shown in figure 2(a).

The CERN and Brookhaven based muon ($g - 2$) experiments stored relativistic muons in a uniform magnetic field, which resulted in the muon spin precessing with constant frequency $\vec{\omega}_a$, while the muons travelled in circular orbits. If *all* decay electrons were counted, the number detected as a function of time would be a pure exponential; therefore we seek cuts on the laboratory observables to select subsets of decay electrons whose numbers oscillate at the precession frequency. Recalling that the number of decay electrons in the MRF varies with the angle between the electron and spin directions, the electrons in the subset should have a preferred direction in the MRF when weighted according to their asymmetry as given in equation (5). At $p_{\mu} \approx 3.094$ GeV/ c the directions of the electrons resulting from muon decay in the laboratory frame are very nearly parallel to the muon momentum regardless of their energy or direction in the MRF. Therefore the only practical remaining cut is on the electron's laboratory energy. Typically, selecting an energy subset will have the desired effect: there will be a net component of electron MRF momentum either parallel or antiparallel to the laboratory muon direction. For example, suppose that we only count electrons with the

highest laboratory energy, around 3.1 GeV. Let \hat{z} indicate the direction of the muon laboratory momentum. The highest-energy electrons in the laboratory are those near the maximum MRF energy of 53 MeV and with MRF directions nearly parallel to \hat{z} . There are more of these high-energy electrons when the μ^- spins are in the direction opposite to \hat{z} than when the spins are parallel to \hat{z} . Thus the number of decay electrons reaches a maximum when the muon spin direction is opposite to \hat{z} and a minimum when they are parallel. As the spin precesses, the number of high-energy electrons will oscillate with frequency ω_a . More generally, at laboratory energies above ~ 1.2 GeV, the electrons have a preferred average MRF direction parallel to \hat{z} (see figure 2). In this discussion, it is assumed that the spin precession vector, $\vec{\omega}_a$, is independent of time, and therefore the angle between the spin component in the orbit plane and the muon momentum direction is given by $\omega_a t + \phi$, where ϕ is a constant.

Equations (5) and (6) can be transformed to the laboratory frame to give the electron number oscillation with time as a function of electron energy,

$$N_d(t, E) = N_{d0}(E)e^{-t/\gamma\tau}[1 + A_d(E) \cos(\omega_a t + \phi_d(E))] \quad (7)$$

or, taking all electrons above threshold energy E_{th} ,

$$N(t, E_{\text{th}}) = N_0(E_{\text{th}})e^{-t/\gamma\tau}[1 + A(E_{\text{th}}) \cos(\omega_a t + \phi(E_{\text{th}}))]. \quad (8)$$

In equation (7) the differential quantities are

$$A_d(E) = \mathcal{P} \frac{-8y^2 + y + 1}{4y^2 - 5y - 5}, \quad N_{d0}(E) \propto (y - 1)(4y^2 - 5y - 5), \quad (9)$$

and in equation (8),

$$N(E_{\text{th}}) \propto (y_{\text{th}} - 1)^2(-y_{\text{th}}^2 + y_{\text{th}} + 3), \quad A(E_{\text{th}}) = \mathcal{P} \frac{y_{\text{th}}(2y_{\text{th}} + 1)}{-y_{\text{th}}^2 + y_{\text{th}} + 3}. \quad (10)$$

In the above equations, $y = E/E_{\text{max}}$, $y_{\text{th}} = E_{\text{th}}/E_{\text{max}}$, \mathcal{P} is the polarization of the muon beam and E , E_{th} and $E_{\text{max}} = 3.1$ GeV are the electron laboratory energy, threshold energy and maximum energy, respectively.

The fractional statistical error on the precession frequency, when fitting data collected over many muon lifetimes to the five-parameter function (equation (8)), is given by

$$\delta\epsilon = \frac{\delta\omega_a}{\omega_a} = \frac{\sqrt{2}}{2\pi f_a \tau_\mu N^{\frac{1}{2}} A}, \quad (11)$$

where N is the total number of electrons and A is the asymmetry in the given data sample. For a fixed magnetic field and muon momentum, the statistical figure of merit is NA^2 , the quantity to be maximized in order to minimize the statistical uncertainty.

The energy dependences of the numbers and asymmetries used in equations (7) and (8), along with the figures of merit NA^2 , are plotted in figures 2 and 3 for the case of E821. The statistical power is greatest for electrons at 2.6 GeV (figure 2). When a fit is made to all electrons above some energy threshold, the optimal threshold energy is about 1.7–1.8 GeV (figure 3).

1.2. History of the muon ($g - 2$) experiments

In 1957, at the Columbia–Nevis cyclotron, the spin rotation of a muon in a magnetic field was observed for the first time. The torque exerted by the magnetic field on the muon's magnetic moment produces a spin precession frequency

$$\vec{\omega}_S = -\frac{qg\vec{B}}{2m} - \frac{q\vec{B}}{\gamma m}(1 - \gamma), \quad (12)$$




Figure 3. The integral N , A and NA^2 (arbitrary units) for a single energy-threshold as a function of the threshold energy; (a) in the laboratory frame, not including and (b) including the effects of detector acceptance and energy resolution for the E821 calorimeters discussed below. For the third CERN experiment and E821, $E_{\max} \approx 3.1$ GeV ($p_{\mu} = 3.094$ GeV/c) in the laboratory frame.

where $\gamma = (1 - \beta^2)^{-\frac{1}{2}}$, with $\beta = v/c$. Garwin *et al* [30] found that the observed rate of spin rotation was consistent with $g = 2$.

In a subsequent paper, Garwin *et al* [33] reported the results of a second experiment, a measurement of the muon anomaly to a relative precision of 6.6%, $g_{\mu} \geq 2(1.00122 \pm 0.00008)$, with the inequality coming from the poor knowledge of the muon mass. In a note added in proof, the authors reported that a new measurement of the muon mass permitted them to conclude that $g_{\mu} = 2(1.00113^{+0.00016}_{-0.00012})$. Within the experimental uncertainty, the muon's anomaly was equal to that of the electron. More generally, the muon was shown to behave like a heavy electron, a spin 1/2 fermion obeying QED.

In 1961 the first of three experiments to be carried out at CERN reported a more precise result obtained at the CERN synchrocyclotron [34]. In this experiment, highly polarized muons of momentum 90 MeV/c were injected into a 6 m long magnet with a graded magnetic field. As the muons moved in almost circular orbits which drifted transverse to the gradient, their spin vectors precessed with respect to their momenta. The rate of spin precession is readily calculated. Assuming that $\vec{\beta} \cdot \vec{B} = 0$, the momentum vector of a muon undergoing cyclotron motion rotates with frequency

$$\vec{\omega}_C = -\frac{q\vec{B}}{m\gamma}. \quad (13)$$

The spin precession *relative to the momentum* occurs at the difference frequency, ω_a , between the spin frequency in equation (12) and the cyclotron frequency,

$$\vec{\omega}_a = \vec{\omega}_S - \vec{\omega}_C = -\left(\frac{g-2}{2}\right)\frac{q\vec{B}}{m} = -a_{\mu}\frac{q\vec{B}}{m}. \quad (14)$$

The precession frequency ω_a has the important property that it is *independent* of the muon momentum. When the muons reached the end of the magnet, they were extracted and their polarizations measured. The polarization measurement exploited the self-analysing property of the muon: more electrons are emitted opposite than along the muon spin. For an ensemble of muons, ω_a is the average observed frequency and B is the average magnetic field obtained by folding the muon distribution with the magnetic field map.

The result from the first CERN experiment was [34] $a_{\mu^+} = 0.001145(22)$ (1.9%), which can be compared with $\alpha/2\pi = 0.001161410\dots$. With additional data this technique resulted in the first observation of the effects of the $(\alpha/\pi)^2$ term in the QED expansion [35].

The second CERN experiment used a muon storage ring operating at 1.28 GeV/c. Vertical focusing was achieved with magnetic gradients in the storage-ring field. While the use of magnetic gradients to focus a charged particle beam is quite common, it makes a precision determination of the (average) magnetic field which enters into equation (14) rather difficult for two reasons. Since the field is not uniform, information on where the muons are in the storage ring is needed to correct the average field for the gradients encountered. Also, the presence of gradient magnetic fields broadens the NMR line-shape, which reduces the precision on the NMR measurement of the magnetic field.

A temporally narrow bunch of 10^{12} protons at 10.5 GeV/c from the CERN proton synchrotron (PS) struck a target inside the storage ring, producing pions, a few of which decay in such a way that their daughter muons are stored in the ring. A huge flux of other hadrons was also produced, which presented a challenge to the decay electron detection system. The electron detectors could only be placed in positions around the ring well removed from the production target, which limited their geometric coverage. Of the pions which circulated in the ring for several turns and then decayed, only one in a thousand produced a stored muon, resulting in about 100 stored muons per injected proton bunch. The polarization of the stored muons was 26% [36].

In all the experiments discussed in this review, the magnetic field was measured by observing the Larmor frequency of stationary protons, ω_p , in nuclear magnetic resonance (NMR) probes. Ignoring the signs of the muon and proton charges, we have the two equations:

$$\omega_a = \frac{e}{m} a_\mu B \quad \text{and} \quad \omega_p = g_p \left(\frac{eB}{2m_p} \right), \quad (15)$$

where ω_p is the Larmor frequency for a *free* proton. Dividing and solving for a_μ we find

$$a_\mu = \frac{\tilde{\omega}_a/\omega_p}{\lambda - \tilde{\omega}_a/\omega_p} = \frac{\mathcal{R}}{\lambda - \mathcal{R}}, \quad (16)$$

where the fundamental constant $\lambda = \mu_\mu/\mu_p$ and $\mathcal{R} = \tilde{\omega}_a/\omega_p$. The tilde on $\tilde{\omega}_a$ indicates that the measured frequency has been adjusted for any necessary (small) corrections, such as the pitch and radial electric field corrections discussed in section 3. Equation (16) was used in the second CERN experiment, and in all the subsequent muon ($g - 2$) experiments, to determine the value of the anomaly. The most accurate value of the ratio λ is derived from measurements of the muonium (the μ^+e^- atom) hyperfine structure [39, 40], i.e. from measurements made with the μ^+ , so that it is more properly written as $\lambda_+ = \mu_{\mu^+}/\mu_p$. Application of equation (16) to determine a_{μ^-} requires the assumption of CPT invariance.

The second CERN experiment obtained $a_\mu = (11\,661.6 \pm 3.1) \times 10^{-7}$ (0.27%) [36], testing quantum electrodynamics for the muon to the three-loop level. Initially this result was 1.8 standard deviations above the QED value, which stimulated a new calculation of the QED light-by-light scattering contribution [108] that brought theory and experiment into agreement.

In the third experiment at CERN, several of the undesirable features of the second experiment were eliminated. A secondary pion beam produced by the PS was brought into the storage ring, which was placed a suitable distance away from the production target. The $\pi \rightarrow \mu$ decay was used to kick muons onto a stable orbit, producing a stored muon beam with polarization of 95%. The ‘decay kick’ had an injection efficiency of 125 ppm, with many of the remaining pions striking objects in the storage ring and causing background in the detectors after injection. While still significant, the background in the electron calorimeters associated with beam injection was much reduced compared with the second CERN experiment.

Another very important improvement in the third CERN experiment, adopted by the Brookhaven based experiment as well, was the use of electrostatic focusing [37], which represented a major conceptual breakthrough. The magnetic gradients used by the second

CERN experiment made the precise determination of the average magnetic field difficult, as discussed above. The more uniform field greatly improves the performance of NMR techniques, to provide a measure of the field at the 10^{-7} – 10^{-8} level. With electrostatic focusing equations (12) and (13) must be modified. The cyclotron frequency becomes

$$\vec{\omega}_C = -\frac{q}{m} \left[\frac{\vec{B}}{\gamma} - \frac{\gamma}{\gamma^2 - 1} \left(\frac{\vec{\beta} \times \vec{E}}{c} \right) \right] \quad (17)$$

and the spin precession frequency becomes [42]

$$\vec{\omega}_S = -\frac{q}{m} \left[\left(\frac{g}{2} - 1 + \frac{1}{\gamma} \right) \vec{B} - \left(\frac{g}{2} - 1 \right) \frac{\gamma}{\gamma + 1} (\vec{\beta} \cdot \vec{B}) \vec{\beta} - \left(\frac{g}{2} - \frac{\gamma}{\gamma + 1} \right) \left(\frac{\vec{\beta} \times \vec{E}}{c} \right) \right]. \quad (18)$$

Substituting for $a_\mu = (g_\mu - 2)/2$, we find that the spin difference frequency is

$$\vec{\omega}_a = -\frac{q}{m} \left[a_\mu \vec{B} - a_\mu \left(\frac{\gamma}{\gamma + 1} \right) (\vec{\beta} \cdot \vec{B}) \vec{\beta} - \left(a_\mu - \frac{1}{\gamma^2 - 1} \right) \frac{\vec{\beta} \times \vec{E}}{c} \right]. \quad (19)$$

If $\vec{\beta} \cdot \vec{B} = 0$, this reduces to

$$\vec{\omega}_a = -\frac{q}{m} \left[a_\mu \vec{B} - \left(a_\mu - \frac{1}{\gamma^2 - 1} \right) \frac{\vec{\beta} \times \vec{E}}{c} \right]. \quad (20)$$

For $\gamma_{\text{magic}} = 29.3$ ($p_\mu = 3.09 \text{ GeV}/c$), the second term vanishes; one is left with the simpler result of equation (14), and the electric field does not contribute to the spin precession relative to the momentum. The great experimental advantage of using the magic γ is clear. By using electrostatic focusing, a gradient B -field is not needed for focusing and B can be made as uniform as possible. The spin precession is determined almost completely by equation (14), which is independent of muon momentum; *all* muons precess at the same rate. Because of the high uniformity of the B -field, a precision knowledge of the stored beam trajectories in the storage region is not required.

Since the spin precession period of $4.4 \mu\text{s}$ is much longer than the cyclotron period of 149 ns, during a single precession period, a muon samples the magnetic field over the entire azimuth 29 times. Thus the important quantity to be made uniform is the magnetic field averaged over azimuth. The CERN magnet was shimmed to an average azimuthal uniformity of ± 10 ppm, with the total systematic error from all issues related to the magnetic field of ± 1.5 ppm [38].

Two small corrections to ω_a are required to form $\vec{\omega}_a$, which is used in turn to determine a_μ in equation (16). (1) The vertical pitching motion (betatron oscillations) about the mid-plane of the storage region means that $\vec{\beta} \cdot \vec{B}$ differs slightly from zero (see equation (19)). (2) Only muons with the central radius of 7112 mm are at the magic value of γ , so that a radial electric field will effect the spin precession of muons with $\gamma \neq \gamma_{\text{magic}}$. The small corrections are described in section 3.2.

The design and shimming techniques of the CERN III storage ring were chosen to produce a very stable magnet mechanically. After cycling the magnet power, the field was reproducible to a few ppm [56]. The 14 m diameter storage ring ($B_0 = 1.5 \text{ T}$) was constructed using 40 identical C-magnets, bolted together in a regular polygon to make a ring, and the field was excited by four concentric coils connected in series. The coils that excited the field in the CERN storage ring were conventional warm coils, and it took about four days after powering for thermal and mechanical equilibrium to be reached [56, 57]. The shimming procedure involved grinding steel from the pole pieces, which introduced a periodic shape to the contour

map averaged over azimuth (see figure 5 of [38]). The magnetic field was mapped once or twice per running period, by removing the vacuum chamber and stepping NMR probes through the storage region.

The electrostatic quadrupoles had a twofold symmetry, covered approximately 80% of the azimuth and defined a rectangular storage region 120 mm (radial) by 80 mm (vertical). The ring behaved as a weak focusing storage ring (betatron) [43–45], with a field index of ~ 0.135 (see section 3).

The pion beam was brought into the storage ring through a pulsed co-axial ‘inflector’, which produced a magnetic field that cancelled the 1.5 T storage-ring field and permitted the beam to arrive undeflected, tangent to the storage circle but displaced 76 mm radially outwards from the central orbit. The inflector current pulse rose to a peak current of 300 kA in 12 μs . The transient effects of this pulsed device on the magnetic field seen by the stored muons during the precession measurement were estimated to be small [38].

The decay electrons were detected by 24 lead-scintillator sandwiches each viewed through an air light-guide by a single 5 in. photomultiplier tube. The photomultiplier signal from each detector was discriminated at several analog thresholds, thus providing a coarse pulse height measurement as well as the time of the pulse. Each of the discriminated signals was assigned to a time bin, in a manner independent of both the ambient rate and the electron arrival time relative to the clock time-boundary. The final precision of 7.3 ppm (statistics dominated) convincingly confirmed the effect of hadronic vacuum polarization, which was predicted to contribute to a_μ at the level of 60 ppm. The factor of 35 increase in sensitivity between the second and third experiments can be traced to the improved statistical power provided by the injected pion beam, both in the number and asymmetry of the muon sample and in reduced systematic errors resulting from the use of electrostatic focusing with the central momentum set to γ_{magic} .

Around 1984 a new collaboration (E821) was formed with the aim of improving on the CERN result to a relative precision of 0.35 ppm, a goal chosen because it was 20% of the predicted first-order electroweak contribution to a_μ . During the subsequent twenty-year period, over 100 scientists and engineers contributed to E821, which was built and performed at the Brookhaven National Laboratory (BNL) Alternating Gradient Synchrotron (AGS). A measurement at this level of precision would test the electroweak renormalization procedure for the Standard Model and serve as a very sensitive constraint on new physics [23–25].

The E821 experiment at Brookhaven [26, 47–51] has advanced the relative experimental precision of a_μ to 0.54 ppm, with a several standard-deviation difference from the prediction of the Standard Model. Assuming CPT symmetry, namely, $a_{\mu^+} = a_{\mu^-}$, the new world average is

$$a_\mu(\text{Expt}) = 11\,659\,208.0(6.3) \times 10^{-10} \quad (0.54 \text{ ppm}). \quad (21)$$

The total uncertainty includes a 0.46 ppm statistical uncertainty and a 0.28 ppm systematic uncertainty, combined in quadrature.

A summary showing the physics reach of each of the successive muon ($g - 2$) experiments is given in table 1. In the subsequent sections we review the most recent experiment, E821 at Brookhaven, and the Standard-Model theory with which it is compared.

2. The Brookhaven experiment E821: experimental technique

In general terms, such as the use of electrostatic focusing and the magic γ , E821 was modelled closely on the third CERN experiment. However, the statistical error goal and the ability of the Brookhaven AGS to provide the (necessary) tremendous increase in beam flux posed enormous challenges to the injection system as well as to the detectors, electronics and data

Table 1. Measurements of the muon anomalous magnetic moment. When the uncertainty on the measurement is the size of the next term in the QED expansion, or the hadronic or weak contributions, the term is listed under ‘sensitivity’. The ‘?’ indicates a result that differs by greater than two standard deviations from the Standard Model. For completeness, we include the experiment of Henry *et al* [46], which is not discussed in the text.

\pm	Measurement	σ_{a_μ}/a_μ	Sensitivity	References
μ^+	$g = 2.00 \pm 0.10$		$g = 2$	Garwin <i>et al</i> [30], Nevis (1957)
μ^+	$0.001\,13^{+0.000\,16}_{-0.000\,12}$	12.4%	$\frac{\alpha}{\pi}$	Garwin <i>et al</i> [33], Nevis (1959)
μ^+	0.001 145(22)	1.9%	$\frac{\alpha}{\pi}$	Charpak <i>et al</i> [34] CERN 1 (SC) (1961)
μ^+	0.001 162(5)	0.43%	$\left(\frac{\alpha}{\pi}\right)^2$	Charpak <i>et al</i> [35] CERN 1 (SC) (1962)
μ^\pm	0.001 166 16(31)	265 ppm	$\left(\frac{\alpha}{\pi}\right)^3$	Bailey <i>et al</i> [36] CERN 2 (PS) (1968)
μ^+	0.001 060(67)	5.8%	$\frac{\alpha}{\pi}$	Henry <i>et al</i> [46] solenoid (1969)
μ^\pm	0.001 165 895(27)	23 ppm	$\left(\frac{\alpha}{\pi}\right)^3 + \text{Hadronic}$	Bailey <i>et al</i> [37] CERN 3 (PS) (1975)
μ^\pm	0.001 165 911(11)	7.3 ppm	$\left(\frac{\alpha}{\pi}\right)^3 + \text{Hadronic}$	Bailey <i>et al</i> [38] CERN 3 (PS) (1979)
μ^+	0.001 165 919 1(59)	5 ppm	$\left(\frac{\alpha}{\pi}\right)^3 + \text{Hadronic}$	Brown <i>et al</i> [48] BNL (2000)
μ^+	0.001 165 920 2(16)	1.3 ppm	$\left(\frac{\alpha}{\pi}\right)^4 + \text{Weak}$	Brown <i>et al</i> [49] BNL (2001)
μ^+	0.001 165 920 3(8)	0.7 ppm	$\left(\frac{\alpha}{\pi}\right)^4 + \text{Weak} + ?$	Bennett <i>et al</i> [50] BNL (2002)
μ^-	0.001 165 921 4(8)(3)	0.7 ppm	$\left(\frac{\alpha}{\pi}\right)^4 + \text{Weak} + ?$	Bennett <i>et al</i> [51] BNL (2004)
μ^\pm	0.001 165 920 80(63)	0.54 ppm	$\left(\frac{\alpha}{\pi}\right)^4 + \text{Weak} + ?$	Bennett <i>et al</i> [51, 26] BNL WA (2004)

acquisition. At the same time, systematic errors, in the measurements of both the field and anomalous precession frequencies, had to be reduced significantly compared with the third CERN experiment. The challenges included the following.

- How to allow multiple beam bunches from the AGS to pass undeflected through the back leg of the storage ring magnet. The use of multiple bunches reduces the instantaneous rates in the detectors, but it was judged technically impractical to rapidly pulse an inflector like that used at CERN.
- How to provide a homogeneous storage ring field, uniform to 1 ppm when averaged over azimuth.
- How to store an injected *muon* beam, which would greatly reduce the hadronic flash seen by the detectors.
- How to maximize the statistical power of the detected electrons.
- How to make good measurements on the decay electrons—good energy resolution, and, in the face of high data rates with a wide dynamic range, how to maintain the stability of signal gain and timing pickoff.

Some of the solutions were the natural product of technical advances over the previous twenty years. Others, such as the superconducting inflector, were truly novel, and required considerable development. A comparison of E821 and the third CERN experiment is given in table 2.

Table 2. A comparison of the features of the E821 and the third CERN muon ($g-2$) experiment [38]. Both experiments operated at the ‘magic’ $\gamma = 29.3$, and used electrostatic quadrupoles for vertical focusing. Bailey, *et al* [38], do not quote a systematic error on the muon frequency ω_a .

Quantity	E821	CERN
Magnet	Superconducting	Room temperature
Yoke construction	Monolithic Yoke	40 Separate magnets
Magnetic field	1.45 T	1.47 T
Magnet gap	180 mm	140 mm
Stored energy	6 MJ	
Field mapped <i>in situ</i>	Yes	No
Central orbit radius	7112 mm	7000 mm
Averaged field uniformity	± 1 ppm	± 10 ppm
Muon storage region	90 mm diameter circle	120×80 mm ² rectangle
Injected beam	Muon	Pion
Inflector	Static superconducting	Pulsed coaxial Line
Kicker	Pulsed magnetic	$\pi \rightarrow \mu \nu_\mu$ decay
Kicker efficiency	$\sim 4\%$	125 ppm
Muons stored/fill	10^4	350
Ring symmetry	Four fold	Two fold
$\sqrt{\beta_{\max}/\beta_{\min}}$	1.03	1.15
Detectors	Pb-scintillating fibre	Pb-scintillator ‘sandwich’
Electronics	Waveform digitizers	Discriminators
Systematic error on B-field	0.17 ppm	1.5 ppm
Systematic error on ω_a	0.21 ppm	Not given
Total systematic error	0.28 ppm	1.5 ppm
Statistical error on ω_a	0.46 ppm	7.0 ppm
Final total error on a_μ	0.54 ppm	7.3 ppm

An engineering run with pion injection occurred in 1997, a brief μ -injection run where the new muon kicker was commissioned happened in 1998 and major data-collection periods of 3–4 months duration took place in 1999, 2000 and 2001. In each period, protons were accelerated by a linear accelerator, accelerated further by a booster synchrotron and then injected into the BNL alternating gradient synchrotron (AGS). Radio-frequency cavities in the AGS ring provide acceleration to a momentum of 24 GeV/ c and maintain the protons in a number of discrete, equally spaced bunches. The number of bunches (harmonic number) in the AGS during a 2.7 s acceleration cycle was different for each of these periods: eight in 1999, six in 2000 and twelve in 2001. The AGS has the ability to deliver up to 70×10^{12} protons (70 Tp) in one AGS cycle, providing a proton intensity per hour 180 times greater than that available at CERN in the 1970s.

The proton beam is extracted from the AGS one bunch at a time at 33 ms intervals. Each proton bunch results in a narrow time bunch of muons which is injected into the storage ring, and then the electrons from muon decays are measured for about 10 muon lifetimes, or about 640 μ s. A plan view of the Brookhaven Alternating Gradient Synchrotron, injection line and storage ring are shown in figure 4. Because the maximum total intensity available from the AGS is ≤ 70 Tp, the bunch intensity and the resulting pile-up (accidental coincidences between two electrons) in the detectors is minimized by maximizing the number of proton bunches. Pulse pile-up in the detectors following injection into the storage ring is one of the systematic issues requiring careful study in the data analysis.

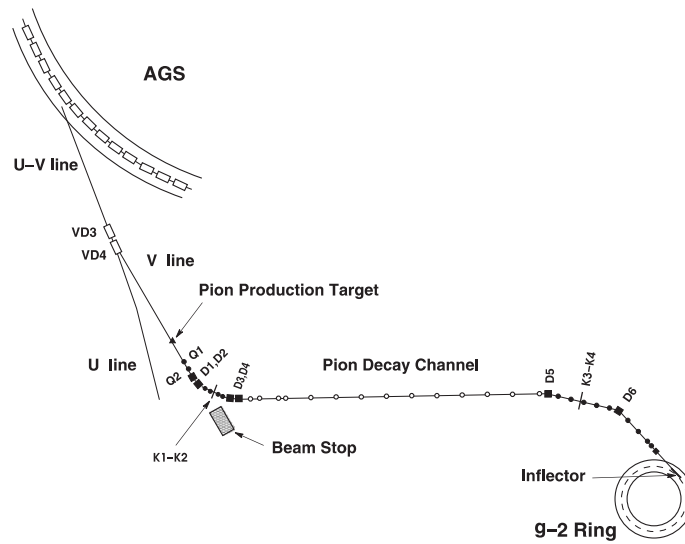


Figure 4. The E821 beamline and storage ring. Pions produced at 0° are collected by the quadrupoles Q1–Q2 and the momentum is selected by the collimators K1–K2. The pion decay channel is 72 m in length. Forward muons at the magic momentum are selected by the collimators K3–K4.

2.1. The proton and muon beamlines

The primary proton beam from the AGS is brought to a water-cooled nickel production target. Because of mechanical shock considerations, the intensity of a bunch is limited to less than 7 Tp. The beamline, shown in figure 4, accepts pions produced at 0° at the production target. They are collected by the first two quadrupoles, momentum analysed and brought into the decay channel by four dipoles. A pion momentum of $3.115 \text{ GeV}/c$, 1.7% higher than the magic momentum, is selected. The beam then enters a straight 80 m long focusing–defocusing quadrupole channel, where those muons from pion decays that are emitted approximately parallel to the pion momentum, so-called forward decays, are collected and transported downstream. Muons at the magic momentum of $3.094 \text{ MeV}/c$, and having an average polarization of 95%, are separated from the slightly higher momentum pions at the second momentum slit. However, after this momentum selection a rather large pion component remains in the beam, whose composition was measured to be $1 : 1 : 1, e^+ : \mu^+ : \pi^+$. The proton content was calculated to be approximately one-third of the pion flux [48]. The secondary muon beam intensity incident on the storage ring was about 2×10^6 per fill of the ring, which can be compared with 10^8 particles per fill with ‘pion injection [47]’ which was used in the 1997 engineering run.

The detectors are exposed to a large flux of background pions, protons, neutrons, electrons and other particles associated with the beam injection process, called the injection ‘flash’. The intensity varies around the ring, being most intense in the detectors adjacent to the injection point (referred to below as ‘upstream’ detectors). The flash consists of prompt and delayed components, with the prompt component caused by injected particles passing directly into the detectors. The delayed component is mainly due to γ rays following neutron capture. The neutrons are produced primarily by the protons and pions in the beam striking the magnet, calorimeters, etc. Many neutrons thermalize in the calorimeter and surrounding materials over tens of microseconds, where they can undergo nuclear capture. These γ rays cause a dc baseline offset in the PMT signal which steadily decays with a lifetime of $\sim 50\text{--}100 \mu\text{s}$. To reduce the decay time of the neutron background, the epoxy in the upstream subset of




Figure 5. (a) A plan view of the inflector-storage-ring geometry. The dot-dash line shows the central muon orbit at 7112 mm. The beam enters through a hole in the back of the magnet yoke, then passes into the inflector. The inflector cryostat has a separate vacuum from the beam chamber, as can be seen in the cross-sectional view. The cryogenic services for the inflector are provided through a radial penetration through the yoke at the upstream end of the inflector. (b) A cross-sectional view of the pole pieces, the outer-radius coil-cryostat arrangement and the downstream end of the superconducting inflector. The muon beam direction at the inflector exit is into the page. The centre of the storage ring is to the right. The outer-radius coils which excite the storage-ring magnetic field are shown, but the inner-radius coils are omitted.

detectors is doped with a natural boron carbide powder, thus taking advantage of the large neutron capture cross-section on ^{10}B .

This injection flash is most severe with the pion injection scheme. The upstream photomultiplier tubes had to be gated off for $120\ \mu\text{s}$ (1.8 muon lifetimes) following injection, to allow the signals to return to the nominal baseline. To reduce this ‘flash’ and to increase significantly the number of muons stored per fill of the storage ring, a fast muon kicker was developed which permitted direct muon injection into the storage ring.

2.2. The inflector magnet

The geometry of the incoming beam is shown in figure 5. A unique superconducting inflector magnet [58,59] was built to cancel the magnetic field permitting the beam to arrive undeflected at the edge of the storage region (see figure 6(b)). The inflector is a truncated, double-cosine theta magnet, shown in figure 5(b) at its downstream end, with the muon velocity going *into* the page. In the inflector, the current flows into the page down the central ‘C’-shaped layer of superconductor, then out of the page through the ‘backward-D’-shaped outer conductor layer. At the inflector exit, the centre of the injected beam is 77 mm from the central orbit. For μ^+ stored in the ring, the main field points up and $q\vec{v} \times \vec{B}$ points to the right in figure 5(b), towards the ring centre.

In this design shown in figures 5(b) and 6(b), the beam channel aperture is rather small compared with the flux return area. The field is $\sim 3\ \text{T}$ in the return area (inflector plus central storage-ring field), and the flux density is sufficiently high to lower the critical current in the superconductor placed in that region. If the beam channel aperture were to be increased by pushing the coil further into the flux return area, the design would have to be changed, either by employing a superconductor with larger critical current or by using more conductors in a revised geometry, further complicating the fabrication of this magnet. The result of the small inflector aperture is a rather poor phase-space match between the inflector and the storage ring and, as a consequence, a loss of stored muons.

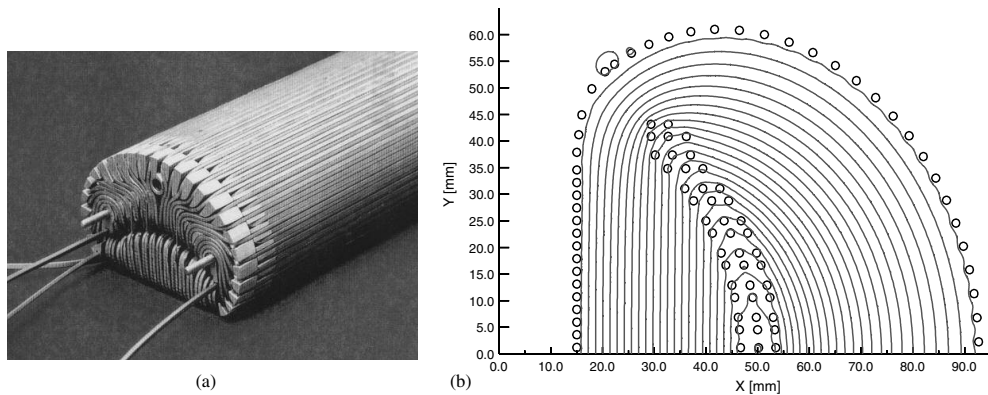


Figure 6. (a) A photo of the prototype inflector showing the crossover between the two coils. The beam channel is covered by the lower crossover. (b) The magnetic design of the inflector. Note that the magnetic flux is largely contained inside of the inflector volume.

As can be seen from figure 6(a), the entrance (and exit) to the beam channel is covered with superconductor, as well as by aluminium windows that are not visible in the photograph. This design was chosen to maximize the mechanical stability of the superconductor in the magnetic field, thus reducing the risk of motion which would quench the magnet. However, multiple scattering in the material at both the entrance and exit windows causes about half the incident muon beam to be lost.

The distribution of conductor on the outer surface of the inflector magnet (the ‘D-shaped’ arrangement) prevents most of the magnetic flux from leaking outside of the inflector volume, as seen from figure 6(b). To prevent flux leakage from entering the beam storage region, the inflector is wrapped with a passive superconducting shield that extends beyond both inflector ends, with a 2 m seam running longitudinally along the inflector side away from the storage region. With the inflector at zero current and the shield warm, the main storage ring magnet is energized. Next the inflector is cooled down, so that the shield goes superconducting and pins the precision field inside the inflector region. When the inflector magnet is powered, the supercurrents in the shield prevent the leakage field from penetrating into the storage region behind the shield.

A 0.5 m long prototype of the inflector was fabricated and then tested, first by itself, then in an external magnetic field. Afterwards the full-size 1.7 m inflector was fabricated. Unfortunately, the full-size inflector was damaged during initial tests, making it necessary to cut through the shield to repair the interconnect between the two inflector coils. A patch was placed over the cut, and this repaired inflector was used in the 1997 π -injection run as well as in the 1998 and 1999 run periods. Magnetic flux leakage around the patch reduced the main storage-ring field locally by ~ 600 ppm. This inhomogeneity complicated the field measurement due to the degradation in NMR performance and contributed an additional uncertainty of 0.2 ppm to the average magnetic field seen by the muons. The leakage field from a new inflector, installed before the 2000 running period, was immeasurably small.

2.3. The fast muon kicker

Left undisturbed, the injected muon beam would pass once around the storage ring, strike the inflector and be lost. As shown in figure 7, the role of the fast muon kicker is to briefly reduce a section of the main storage field and move the centre of the muon orbit to the geometric




Figure 7. A sketch of the beam geometry. $R = 7112$ mm is the storage ring radius and $x_c = 77$ mm is the distance between the inflector centre and the centre of the storage region. This is also the distance between the centres of the circular trajectory that a particle entering at the inflector centre (at $\theta = 0$ with $x' = 0$) will follow and the circular trajectory a particle at the centre of the storage volume (at $\theta = 0$ with $x' = 0$) will follow.

centre of the storage ring. The 77 mm offset at the injection point, between the centre of the entering beam and the central orbit, requires that the beam be kicked outwards by approximately 10 mrad. The kick should be made at about 90° around the ring, plus a few degrees correction due to the defocusing effect of the electric quadrupoles between the injection point and the kicker, as shown in figures 5(b) and 7.

The requirements on the fast muon kicker are rather stringent. While electric, magnetic and combination electromagnetic kickers were considered, the collaboration settled on a magnetic kicker design [72] because it was thought to be technically easier and more robust than the other options. Because of the very stringent requirements on the storage ring magnetic field uniformity, no magnetic materials could be used. Thus the kicker field had to be generated and shaped solely with currents, rather than using ferrite cores. Even with the kicker field generated by currents, there existed the potential problem of inducing eddy currents which might affect the magnetic field seen by the stored muons.

The length of the kicker is limited to the ~ 5 m azimuthal space between the electrostatic quadrupoles (see figure 10), so each of the three sections is 1.76 m long. The cross-section of the kicker is shown in figure 8. The two parallel conductors are connected with cross-overs at each end, forming a single current loop. The kicker plates also have to serve as ‘rails’ for the NMR field-mapping trolley (discussed below), and the trolley is shown riding on the kicker rails in figure 8.

The kicker current pulse is formed by an under-damped LCR circuit. A capacitor is charged to 95 kV through a resonant charging circuit. Just before the beam enters the storage ring, the capacitor is shorted to ground by firing a deuterium thyratron. The peak current in an LCR circuit is given by $I_0 = V_0/(\omega_d L)$ making it necessary to keep the system inductance, L , low to maximize the magnetic field for a given voltage V_0 . For this reason, the kicker was divided into three sections, each powered by a separate pulse-forming network.

The electrode design and the time-varying fields were modelled using the package OPERA 2d⁸. Calculations for electrodes made of Al, Ti and Cu were carried out with the geometry shown in figure 8. The aluminium electrodes were best for minimizing eddy currents 20 μ s after injection, which was fortunate, since mechanical tests on prototype electrodes made from the three metals showed that only aluminium was satisfactory. This was also true from the detector’s point of view, since the low atomic number of aluminium minimizes the multiple scattering and showering of the decay electrons on their way to the calorimeters. The magnetic

⁸ OPERA, Electromagnetic Fields Analysis Program, Vector Fields Ltd, 24 Bankside, Oxford OX5 1JE, England.




Figure 8. An elevation view of the kicker plates, showing the ceramic cage supporting the kicker plates, and the NMR trolley riding on the kicker plates. (The trolley is removed during data collection.)




Figure 9. (a) The main magnetic field of the kicker measured with the Faraday effect. (b) The residual magnetic field measured using the Faraday effect. The solid points are calculations from OPERA, and the small \times are the experimental points measured with the Faraday effect. The solid horizontal lines show the ± 0.1 ppm band for affecting $\int \vec{B} \cdot d\vec{\ell}$.

field produced by the kicker, along with the residual field, was measured using the Faraday effect [72], and these are shown in figure 9.

The cyclotron period of the ring is 149 ns, substantially less than the kicker pulse base-width of ~ 400 ns, so that the injected beam is kicked on the first few turns. Nevertheless, approximately 10^4 muons are stored per fill of the ring, corresponding to an injection efficiency of about 3–5% (ratio of stored to incident muons). The storage efficiency with muon injection is much greater than that obtained with pion injection for a given proton flux, with only 1% of the hadronic flash.




Figure 10. The layout of the storage ring, as seen from above, showing the location of the inflector, the kicker sections (labelled K1–K3) and the quadrupoles (labelled Q1–Q4). The beam circulates in a clockwise direction. Also shown are the collimators, which are labelled ‘C’ or ‘ $\frac{1}{2}C$ ’, indicating whether the Cu collimator covers the full aperture or half the aperture. The collimators are rings with inner radius: 45 mm, outer radius: 55 mm, thickness: 3 mm. The scalloped vacuum chamber consists of 12 sections joined by bellows. The chambers containing the inflector, the NMR trolley garage and the trolley drive mechanism are special chambers. The other chambers are standard, with either quadrupole or kicker assemblies installed inside. An electron calorimeter is placed behind each of the radial windows, at the position indicated by the calorimeter number.

2.4. The electrostatic quadrupoles

The electric quadrupoles, which are arranged around the ring with fourfold symmetry, provide vertical focusing for the stored muon beam. The quadrupoles cover 43% of the ring in azimuth, as shown in figure 10. While the ideal vertical profile for a quadrupole electrode would be hyperbolic, beam dynamics calculations determined that the higher multipoles present with flat electrodes, which are much easier to fabricate, would not cause an unacceptable level of beam losses. The flat electrodes are shown in figure 11. Only certain multipoles are permitted by the fourfold symmetry, and a judicious choice of the electrode width relative to the separation between opposite plates minimizes the lowest of these. With this configuration, the 20-pole is the largest, being 2% of the quadrupole component and an order of magnitude greater than the other allowed multipoles [73].

In the quadrupole regions, the combined electric and magnetic fields can lead to electron trapping. The electron orbits run longitudinally along the inside of the electrode, and then return on the outside. Excessive trapping in the relatively modest vacuum of the storage ring can cause sparking. To minimize trapping, the leads were arranged to introduce a dipole field at the end of the quadrupole, thus sweeping away trapped electrons. In addition, the quadrupoles are pulsed, so that after each fill of the ring all trapped particles could escape. Since some protons (antiprotons) were stored in each fill of μ^+ (μ^-), they were also released at the end of each storage time. This lead arrangement worked so well in removing trapped electrons that for the μ^+ polarity it would have been possible to operate the quadrupoles in a dc mode. For




Figure 11. A photograph of an electrostatic quadrupole assembly inside a vacuum chamber. The cage assembly doubles as a rail system for the NMR trolley which is resting on the rails. The locations of the NMR probes inside the trolley are shown as black circles (see section 2.6). The probes are located just behind the front face. The inner (outer) circle of probes has a diameter of 3.5 cm (7 cm) at the probe centres. The storage region has a diameter of 9 cm. The vertical locations of three upper fixed probes is also shown. The fixed probes are located symmetrically above and below the vacuum chamber.

the storage of μ^- , this was not true; the trapped electrons necessitated an order of magnitude better vacuum and limited the storage time to less than $700 \mu\text{s}$.

Beam losses during the measurement period, which could distort the expected time spectrum of decay electrons, had to be minimized. Beam scraping is used to remove, just after injection, those muons which would likely be lost later on. To this end, the quadrupoles are initially powered asymmetrically and then brought to their final symmetric voltage configuration. The asymmetric voltages lower the beam and move it sideways in the storage ring. Particles whose trajectories reach too near the boundaries of the storage volume (defined by collimators placed at the ends of the quadrupole sectors) are lost. The scraping time was $17 \mu\text{s}$ during all data-collection runs except 2001, where $7 \mu\text{s}$ was used. The muon loss rates without scraping were on the order of 0.6% per lifetime at late times in a fill, which dropped to $\sim 0.2\%$ with scraping.

2.5. The superconducting storage ring

The storage ring magnet combined with the electrostatic quadrupoles forms a Penning trap that, while very different in scale, has common features with the electron g -value measurements [21, 54]. In the electron experiments, a single electron is stored for long periods of time, and the spin and cyclotron frequencies are measured. Since the muon decays in $2.2 \mu\text{s}$, the study of a single muon is impossible. Rather, an ensemble of muons is stored at relativistic energies, and their spin precession in the magnetic field is measured using the parity-violating decay to analyse the spin motion (see equations (8), (17) and (18)).

The design goal of ± 1 ppm was placed on the field uniformity when averaged over azimuth in the storage ring. A ‘superferric’ design, where the field configuration is largely determined by the shape and magnetic properties of the iron, rather than by the current distribution in the superconducting coils, was chosen. To reach the ppm level of uniformity it was important to minimize discontinuities such as holes in the yoke, spaces between adjacent pole pieces and, especially, the spacing between pole pieces across the magnet gap containing the beam vacuum chamber. Every effort was made to minimize penetrations in the yoke, and where they




Figure 12. The storage-ring magnet. The cryostats for the inner-radius coils are clearly visible. The kickers have not yet been installed. The racks in the centre are the quadrupole pulsers, and a few of the detector stations are installed, especially the quadrant of the ring closest to the person. The magnet power supply is in the upper left, above the plane of the ring (courtesy of Brookhaven National Laboratory).




Figure 13. (a) A cross-sectional view of the storage ring magnet. The beam centre is at a radius of 7112 mm. The pole pieces are separated from the yoke by an air gap. (b) An expanded view of the pole pieces, which shows the edge shims and the wedge shim.

are necessary, such as for the beam entrance channel, additional iron is placed around the hole to minimize the effect of the hole on the magnetic flux circuit.

The storage ring, shown in figure 12, is designed as a continuous C-magnet [60] with the yoke made up of twelve sectors with minimum gaps where the yoke pieces come together. A cross-section of the magnet is shown in figure 13. The largest gap between adjacent yoke pieces after assembly is 0.5 mm. The pole pieces are built in 36 pieces, with keystone rather than radial boundaries to ensure a close fit. They are electrically isolated from each other with $80\ \mu\text{m}$ kapton to prevent eddy currents from running around the ring, especially during a quench or energy extraction from the magnet. The vertical mismatch from one pole piece to the next when going around the ring in azimuth is held to $\pm 10\ \mu\text{m}$, since the field strength depends critically on the pole-piece spacing across the magnet gap.

The field is excited by 14 m diameter superconducting coils, which in 1996 were the largest-diameter coils ever fabricated. The coil at the outer radius consists of two identical coils on a common mandrel, above and below the plane of the beam, each with 24 turns. Each of the inner-radius coils, which are housed in separate cryostats, also consists of 24 turns (see

figures 5(b) and 13(a)). The nominal operating current is 5200 A, which is driven by a power supply. The choice of using an extremely stable power supply, further stabilized with feedback from the NMR system, was chosen over operating in a ‘persistent mode’, for two reasons. The switch required to change from the powering mode to persistent mode was technically very complicated, and unlike the usual superconducting magnet operated in persistent mode, we anticipated the need to cycle the magnet power a number of times during a three-month running period.

The pole pieces are fabricated from continuous vacuum-cast low-carbon magnet steel (0.0004% carbon) and the yoke from standard AISI 1006 (0.07% carbon) magnet steel [60]. At the design stage, calculations suggested that the field could be made quite uniform and that when averaged over azimuth, a uniformity of ± 1 ppm could be achieved. It was anticipated that, at the initial turn-on of the magnet, the field would have a uniformity of about 1 part in 10^4 and that an extensive program of shimming would be necessary to reach a uniformity of one ppm.

A number of tools for shimming the magnet were therefore built into the design. The air gap between the yoke and pole pieces dominates the reluctance of the magnetic circuit outside of the gap that includes the storage region and decouples the field in the storage region from possible voids or other defects in the yoke steel. Iron wedges placed in the air gap were ground to the wedge angle needed to cancel the quadrupole field component inherent in a C-magnet. The dipole can be tuned locally by moving the wedge radially. The edge shims screwed to the pole pieces were made oversized. Once the initial field was mapped and the sextupole content was measured, the edge shims were reduced in size individually in two steps of grinding, followed by field measurements, to cancel the sextupole component of the field.

After mechanical shimming, these higher multipoles were found to be quite constant in azimuth. They are shimmed out on average by adjusting currents in conductors placed on printed circuit boards going around the ring in concentric circles spaced by 0.25 cm. These boards are glued to the top and bottom pole faces between the edge shims and connected at the pole ends to form a total of 240 concentric circles of conductor, connected in groups of four, to sixty ± 1 A power supplies. These correction coils are quite effective in shimming multipoles up through the octupole. Multipoles higher than octupole are less than 1 ppm at the edge of the storage aperture and, with our use of a circular storage aperture are unimportant in determining the average magnetic field seen by the muon beam (see section 4.1.4).

2.6. Measurement of the precision magnetic field

The magnetic field is measured and monitored by pulsed nuclear magnetic resonance techniques [74]. The free-induction decay is picked up by the coil L_S in figure 14 after a pulsed excitation rotates the proton spin in the sample by 90° to the magnetic field and is mixed with the reference frequency to form the frequency f_{FID} . The reference frequency, $f_{\text{ref}} = 61.74$ MHz, is obtained from a frequency synthesizer phase locked to the LORAN C standard [76], that is chosen with $f_{\text{ref}} < f_{\text{NMR}}$ such that for all probes $f_{\text{FID}} \simeq 50$ kHz. The relationship between the actual field B_{real} and the field corresponding to the reference frequency is given by

$$B_{\text{real}} = B_{\text{ref}} \left(1 + \frac{f_{\text{FID}}}{f_{\text{ref}}} \right). \quad (22)$$

The field measurement process has three aspects: calibration, monitoring the field during data collection and mapping the field. The probes used for these purposes are shown in figure 14. To map the field, an NMR trolley was built with an array of 17 NMR probes arranged in concentric circles, as shown in figure 11. While it would be preferable to have

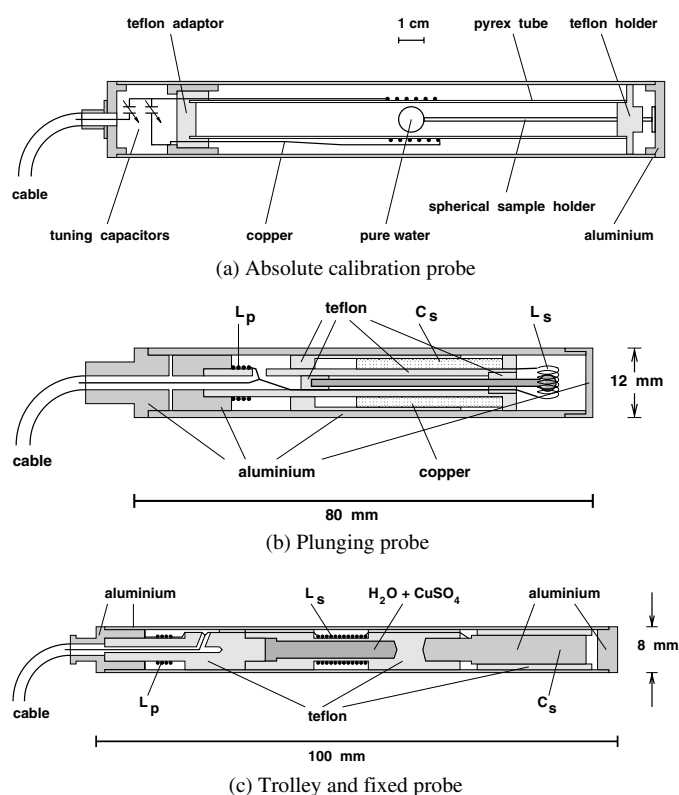


Figure 14. Schematics of different NMR probes. (a) Absolute probe featuring a spherical sample of water. This probe was the same one used in reference [39] to determine λ , the muon-to-proton magnetic-moment ratio. (b) Plunging probe, which can be inserted into the vacuum at a specially shimmed region of the storage ring to transfer the calibration to the trolley probes. (c) The standard probes used in the trolley and as fixed probes. The resonant circuit is formed by the two coils with inductances L_s and L_p and a capacitance C_s made by the Al-housing and a metal electrode.

information over the full 90 mm aperture, space limitations inside the vacuum chamber, which can be understood by examining figures 8 and 11, prevent a larger diameter trolley. The trolley is pulled around the storage ring by two cables, one in each direction circling the ring. One of these cables is a standard thin (Lemo-connector compatible) co-axial signal cable and the other is non-conducting because of its proximity to the kicker high voltage. Power, data in and out and a reference frequency are all carried on the single Lemo cable.

During data-collection periods, the trolley is parked in a garage (see figure 10) in a special vacuum chamber. Every few days, at random times, the field is mapped using the trolley. To map the field, the trolley is moved into the storage region and pulled through the vacuum chamber, measuring the field at over 6000 locations. During the data-collection periods, the storage-ring magnet remains powered continuously for periods lasting from five to twenty days; thus the conditions during mapping are identical to those during the data collection.

To cross calibrate the trolley probes, bellows are placed at one location in the ring to permit a special NMR plunging probe or an absolute calibration probe with a spherical water sample [75], to plunge into the vacuum chamber. Measurements of the field at the same spatial point with the plunging, calibration and trolley probes provide both relative and absolute calibration of the trolley probes. During the calibration measurements before and after each

running period, the spherical water probe is used to calibrate the plunging probe, the centre probe of the trolley and several other trolley probes. The absolute calibration probe provides the calibration to the Larmor frequency of the free proton [65], which is called ω_p below.

To monitor the field on a continuous basis during data collection, a total of 378 NMR probes are placed in fixed locations above and below the vacuum chamber around the ring. Of these, about half provide useful data in monitoring the field with time. Some of the others are noisy, or have bad cables or other problems, but a significant number of fixed probes are located in regions near the pole-piece boundaries where the magnetic gradients are sufficiently large to reduce the free-induction decay time in the probe, limiting the precision on the frequency measurement. The number of probes at each azimuthal position around the ring alternates between two and three, at radial positions arranged symmetrically about the magic radius of 7112 mm. Because of this geometry, the fixed probes provide a good monitor of changes in the dipole and quadrupole components of the field around the storage ring.

Initially the trolley and fixed probes contained a water sample. Over the course of the experiment, the water samples in many of the probes were replaced with petroleum jelly. The jelly has several advantages over water: low evaporation, favourable relaxation times at room temperature, a proton NMR signal almost comparable to that from water and a chemical shift (and the accompanying NMR frequency shift) with a temperature coefficient much smaller than that of water, and thus negligible for our experiment.

2.7. Detectors and electronics

The detector system consists of a variety of particle detectors: calorimeters, position-sensitive hodoscope detectors and a set of tracking chambers. There are also horizontal and vertical arrays of scintillating-fibre hodoscopes which could be temporarily inserted into the storage region. A number of custom electronics modules were developed, including event simulators, multi-hit time-to-digital converters (MTDC) and the waveform digitizers (WFD) which are at the heart of the measurement. We refer to the data collected from one muon injection pulse in one detector as a ‘spill,’ and we will speak of ‘early-to-late’ effects, namely, the gain or time stability requirements in a given detector at early compared with late decay times in a spill.

2.7.1. Electromagnetic calorimeters: design considerations. The electromagnetic calorimeters, together with the custom WFD readout system, are the primary source of data for determining the precession frequency. They provide the energies and arrival times of the electrons, and they also provide signal information immediately before and after the electron pulses, allowing studies of baseline changes and pulse pile-up.

There are 24 calorimeters placed evenly around the 45 m circumference of the storage region, adjacent to the inside radius of the storage vacuum region as shown in figures 10 and 12. Nearly all decay electrons have momenta ($0 < p(\text{lab}) < 3.1 \text{ GeV}/c$, see figure 2) below the stored muon momentum ($3.1 \text{ GeV}/c \pm 0.2\%$), and they are swept by the B -field to the inside of the ring where they can be intercepted by the calorimeters. The storage-region vacuum chamber is scalloped so that electrons pass nearly perpendicular to the vacuum wall before entering the calorimeters, minimizing electron pre-showering (see figure 10). The calorimeters are positioned and sized in order to maximize the acceptance of the highest-energy electrons, which have the largest statistical figure of merit NA^2 . The variations of N and A as a function of electron energy are shown in figures 2 and 3.

The electrons with the lowest laboratory energies, while more numerous than high energy electrons, generally have a lower figure of merit and therefore carry relatively little information on the precession frequency. These electrons have relatively small radii of curvature and exit the

ring vacuum chamber closer to the radial direction than electrons at higher energies, with most of them missing the detectors entirely. Detection of these electrons would require detectors that cover a much larger portion of the circumference than is needed for high-energy electrons, and is not cost effective.

Consequently the detector system is designed to maximize the acceptance of the high energy decay electrons above approximately 1.8 GeV, with the acceptance falling rapidly below this energy. The detector acceptance reaches a maximum of 87% at 2.3 GeV, decreases to 70% at 1.8 GeV and continues to decrease roughly linearly to zero as the energy decreases. With increasing energy above 2.3 GeV, the acceptance also decreases because the highest-energy electrons tend to enter the calorimeters at the outer radial edge, increasing the loss of registered energy due to shower leakage and reducing the acceptance to 80% at 3 GeV.

In a typical analysis, the full data sample consists of all electrons above a threshold energy of about 1.8 GeV, where NA^2 is approximately a maximum, with about 65% of the electrons above that energy detected (figure 3). The average asymmetry is about 0.35. The loss of efficiency is from the low-energy tail in the detector response characteristic of electromagnetic showers in calorimeters and from lower energy electrons missing the detectors altogether. The statistical error improves by only 5% if the data sample contains all electrons above 1.8 GeV compared with all above 2.0 GeV. For threshold energies below 1.7 GeV, the decline in the average asymmetry more than cancels the additional number of electrons in NA^2 , and the statistical error actually increases. Some of the independent analyses fit time spectra of data formed from electrons in narrow energy bands (about 200 MeV wide). When the results of the separate fits are combined, there is a 10% reduction in the statistical error on ω_a . However, there is also a slight increase in the systematic error contribution from gain shifts, because the relative number of events moved by a gain shift from one energy band to another is increased. One analysis used data weighted by the asymmetry as a function of energy. It can be shown [71] that this produces the same statistical improvement as dividing the data into energy bands.

Gain and timing shift limitations are much more stringent within a single spill than from spill to spill. Shifts at late decay times compared with early times in a given spill, so-called ‘early-to-late’ shifts, can lead directly to serious systematic errors on ω_a . Shifts of gain or the $t = 0$ point from one spill to the next are generally much less serious; they will usually only change the asymmetry, average energy, phase, etc, *but to first order, ω_a will be unaffected.*

The calorimeters should have pulses with narrow time widths to minimize the probability of two pulses overlapping (pile-up) during the very high electron decay data rates encountered at early decay times, which can reach a MHz in a single detector. The scintillator is chosen to have minimal long-lived components to reduce the afterglow from the intense detector flash associated with beam injection. Laser calibration studies show that the timing stability for a typical detector over any 200 μ s time interval is better than 15 ps, easily meeting the demands of the measurement of a_μ . For example, a 20 ps timing shift would lead to an uncertainty in a_μ of about 0.1 ppm, which is small compared with the final error. Modest detector energy resolution (≈ 10 – 15% at 2 GeV) is required in order to select the desired high energy electrons for analysis. Better energy resolution also reduces the amount of calibration data needed to monitor the stability of the detector gains.

The stability requirement for the electron energy measurement (‘gain’) versus time in the spill is largely determined by the energy dependence of the phase of the ($g - 2$) oscillation. In a fit of the data to the 5-parameter function, the oscillation phase is highly correlated to ω_a . Therefore a shift in the gain from early to late decay times, combined with an energy dependence in the ($g - 2$) phase, can lead to a systematic error in the determination of ω_a . There are two main contributing factors to the energy dependence, which appear with opposite signs: (1) the phase ϕ in the 5-parameter function (equation (8)) depends on the electron drift time. High-energy




Figure 15. Schematic of a detector station. The calorimeter consists of scintillating fibres embedded in lead, with the fibres being directed radially towards the centre of the storage ring. The electron entrance face of the calorimeter is covered with an array of five horizontally-oriented FSD scintillators. Six of the detector stations also have a PSD scintillating xy hodoscope covering the entrance face.

electrons must travel further, on average, from the point of muon decay to the detector and therefore have longer drift times. The change in drift time with energy implies a corresponding energy dependence in the $(g - 2)$ phase. (2) For decay electrons at a given energy, those with positive (radially outwards) components of momentum at the muon decay point travel further to reach the detectors than electrons with negative (radially inwards) components. They spread out more in the vertical direction and may miss the detectors entirely. Consequently, electrons with positive (outward) radial momentum components will have slightly lower acceptance than those with negative components, causing the average spin direction to rotate slightly, leading to a shift in ϕ . Recalling the correlation between electron direction and muon spin, the overall effect is to shift the time when the number oscillation reaches its maximum, causing a shift in the precession phase in the 5-parameter function. The size of the shift depends on the electron energy. From studies of the data sample and simulations, it is established that the detector gains need to be stable to better than 0.2% over any 200 μs time interval in a spill, in order to keep the systematic error contribution to ω_a less than 0.1 ppm from gain shifts. This requirement is met by all the calorimeters. The gain from one spill to the next is not coupled to the precession frequency and therefore the requirement on the spill to spill stability is far less stringent than the stability requirement within an individual spill.

2.7.2. Electromagnetic calorimeters: construction. The calorimeters (see figure 15) consist of 1 mm diameter scintillating fibres embedded in a lead-epoxy matrix [61]. The final dimensions are 22.5 cm (radial) \times 14 cm high \times 15 cm deep, where the vertical dimension is limited by the vertical magnet gap. The radiation length is $X_0 = 1.14$ cm with a Moliere radius of about 2.5 cm. This gives a calorimeter depth of 13 radiation lengths, which means between 96% and 93% of the energies from 1 to 3 GeV are deposited in the calorimeter, with the remaining energy escaping. The fibres are oriented along the radial direction in the ring, so that incoming electrons are incident approximately perpendicular to the fibres. The large-radius ends of the fibres are reflective in order to improve the uniformity of the pulse height response as a function of the radial direction. There is an average 6% change in the photomultiplier (PMT) pulse height depending on the radial electron entrance point, and there is a change of a

few per cent depending on the vertical position. The measured versus actual electron energy deviates from a linear relationship by about 1% between 0 and 3 GeV, primarily due to shower leakage.

Of the electron shower energy deposited in the calorimeter, about 5% is deposited in the scintillator, producing on the order of 500 photo-electrons per GeV of incident electron energy, with the remainder being deposited in the lead or epoxy. The average overall energy resolution is $\sigma(E)/E \approx 10\%/\sqrt{E(\text{GeV})}$, dominated by the statistics of shower sampling (i.e. variation in the energy that gets deposited in the scintillator versus the other calorimeter material), with some contribution from photo-electron statistics and shower leakage out the back or sides of the calorimeter.

Scintillation light is gathered by four 52.5 cm^2 , nearly square, acrylic light-guides which cover the inner radius side of the calorimeter. Each light-guide necks down to a circle which mates to a 12-dynode 5 cm diameter Hamamatsu R1828 photo multiplier tube. The photomultiplier gains in the four segments of each calorimeter are balanced initially using monochromatic 2 GeV electrons from a BNL test beam. Some detector calibrations shifted when the detectors were moved to their final storage ring locations. In these cases the gains in the quadrants were balanced using the energy spectra from the muon decay electrons themselves. The electron signals, which are the sum of the four signals from the photomultipliers on the calorimeter quadrants, are about 5 ns wide (FWHM) and are sent to waveform digitizers for readout.

In one spill during the time interval from a few tens of microseconds to $640\ \mu\text{s}$ after injection, approximately 20 decay-electrons above 1.8 GeV are recorded on average in each detector. The instantaneous rate of decay electrons above 1 GeV changes from about 300 kHz to almost zero over this period. The gain and timing of photomultipliers can depend on the data rate. The necessary gain and timing stability is achieved with custom, actively stabilized photomultiplier bases [62].

To prevent paralysis of the photomultipliers due to the injection flash, the amplifications in the photomultipliers are temporarily reduced by a factor of about 1 million during the beam injection. Depending on the intensity of the flash and the duration of the background levels encountered at a particular detector station, the amplifications are restored at times between 2 and $50\ \mu\text{s}$ after injection. The switching of the gain is accomplished in the Hamamatsu R1828 photomultipliers by swapping the bias voltages on dynodes 4 and 7. With the proper selection of the delay time after injection to let backgrounds die down, the gains typically return to 99.8% of their steady-state value within several microseconds after the tube is turned back on. Other gating schemes, such as switching the photocathode voltage, were found to have either required a much longer time for the gain to recover or failed to give the necessary reduction in the gain when the tube is gated off.

2.7.3. Special detectors: FSD, PSD, fibre harps, traceback chambers. Several specialized detector systems, the front scintillation detectors (FSD), position-sensitive detectors (PSD), fibre harp scintillators and traceback detectors are all designed to provide information on the phase-space parameters of the stored muon beam and their decay electrons. Such measurements are compared with simulation results and are important, for example, in the study of coherent betatron motion of the stored beam and detector acceptances and in placing a limit on the electric dipole moment of the muon (see sections 3 and 5). A modest knowledge of the beam phase space is necessary in order to calculate the average magnetic and electric fields seen by the stored muons.

The FSDs cover the front (electron entrance) faces of the calorimeters at about half the detector stations. They provide information on the vertical positions of the electrons when

they enter the calorimeters. Each consists of an array of five horizontally-oriented strips of scintillator 23.5 cm long (radial dimension) \times 2.8 cm high \times 1.0 cm thick. Each strip is coupled to a 2.8 cm diameter Hamamatsu R6427 PMT. The signal is discriminated and sent to multi-hit time-to-digital converters (MTDC) for time registration.

The FSDs are an integral part of the muon loss monitoring system. A lost muon is identified as a particle which passes through three successive detector stations, firing the three successive FSD detectors in coincidence. The candidate lost muon is also required to leave minimal energy in each of the calorimeters, since a muon loses energy by ionization, rather than creating an electromagnetic shower. This technique only identifies muons lost in the inward radial direction. Those lost in other directions, for example, above or below the storage region, are not monitored. The scaling between registered and actual muon losses is estimated based on models of the muon loss mechanisms.

Five detector stations are equipped with PSDs. They provide better vertical position resolution than the FSDs and also provide radial position information. Each PSD consists of an xy array of scintillator sticks which, like the FSDs, cover the electron entrance face of the calorimeters. There are 20 horizontally- and 32 vertically-oriented 7 mm wide, 8 mm thick elements read out with multi-anode photomultipliers followed by discriminators and MTDCs.

The traceback-chamber system consists of four sets of drift chambers which can provide up to 12 vertical and 12 horizontal position measurements along a track, with about 300 μm resolution. It is positioned 270° around the ring from the injection point, where a thin window is installed for the electrons to pass through the vacuum wall with minimal scattering. A regular calorimeter is placed behind the array to serve as a trigger. Using the momentum derived from the track information, plus knowledge of the B -field, it is possible to extrapolate the track back to the point where the trajectory is tangent to the storage ring. Since the decay electrons are nearly parallel to the muon momentum, and the muons are travelling with small angles relative to the tangent, the extrapolation point is a very good approximation ($\sigma_V \simeq 9$ mm, $\sigma_R \simeq 15$ mm) to the position of muon decay. Thus, the traceback provides information on the muon distribution for that portion of the ring just upstream of the traceback position.

The storage ring is equipped with two sets of scintillating-fibre beam monitors inside the vacuum chamber. Their purpose is to measure the beam profile as a function of time, and they can be inserted at will into the storage region. Each set consists of an x and y plane of seven 0.5 mm diameter scintillating fibres with 13 mm spacing that covers the central part of the beam region as shown in figure 16. These fibre ‘harps’ are located at 180 and 270° around the ring from the injection point as shown in figure 10. Each fibre is mated to a clear fibre, which connects to a vacuum optical feedthrough and then to a PMT which is read out with a WFD that digitizes continuously during the measurement time. The scintillating fibres are sufficiently thin so that the beam can be monitored for many tens of microseconds after injection, as shown in figure 21, before the stored beam is degraded (effective lifetime about 30 μs), and therefore they are very helpful in providing information on the early-time position and width distributions of the stored muon beam.

2.7.4. Electronics. For each calorimeter, the arrival times of the signals from the four photomultipliers are matched to within a few hundred picoseconds, and the analog sum is formed. The resulting signal is fed to a custom waveform digitizer (WFD) with a 400 MHz equivalent sampling rate, which provides several pulse height samples from each candidate electron.

WFD data are added to the data stream only if a trigger is formed, i.e. when the energy associated with a pulse exceeds a pre-assigned threshold, usually taken to be 900 MeV. When a trigger occurs, WFD samples from about 15 ns before the pulse to about 65 ns after the pulse




Figure 16. A sketch of the fibre beam monitors. The fibres can be rotated into the beam for calibration such that each fibre sees the same portion of the beam.

are recorded. There is the possibility of two or more electron pulses being over-threshold in the same 80 ns time window. In that case, the length of the readout period is extended to include both pulses. At the earliest decay times, the detector signals have a large pedestal due to the ‘flash,’ and some of the upstream detectors are continuously over-threshold at early times.

The energy and time of an electron is obtained by fitting a standard pulse shape to the WFD pulse using a conventional χ^2 minimization. The standard pulse shape is established for each calorimeter. It is based on an average of the shapes of a large number of late-time pulses where the problems associated with overlapping pulses and backgrounds are greatly reduced. There are three fitting parameters, time and height of the pulse and the constant pedestal, with the fits typically spanning 15 samples centred on the pulse. The typical time resolution of an individual electron approached 60 ps.

The period after each pulse is searched for any additional pulses from other electrons. These accidental pulses have the advantage that they do not need to be over the hardware threshold (~ 900 mV), but rather over the much lower (~ 250 MeV) minimum pulse height that can be discriminated from background by the pulse-fitting algorithm. A pile-up spectrum is constructed by combining the triggering pulses and the following accidental pulses as described later in this document.

Zero time for a given fill is defined by the trigger pulse to the AGS kicker magnet that extracts the proton bunch and sends it to the pion production target. The resolution of the zero time needs only to be much better than the ($g - 2$) precession period in order to minimize loss of the asymmetry amplitude.

A pulsed UV laser signal is fanned out simultaneously to all elements of the calorimeter stations to monitor the gain and time stabilities. For each calorimeter quadrant, an optical fibre carries the laser signal to a small bunch of the detector’s scintillating fibres at the outer-radius edge. The laser pulses produce an excitation in the scintillators which is a good approximation to that produced by passing charged particles. The intensity and timing of the laser pulses are monitored with a separate solid-state photo-diode and a PMT, which are shielded from the beam in order to avoid beam-related rate changes and background which might cause shifts in the registered time or pulse height. The times of the laser pulses are chosen to appropriately map the gain and time stability during the 6 to 12 injection bunches per AGS cycle and over the 10 muon lifetimes per injection. Dedicated laser runs are made 6 times per day, for about 20 min each. The average timing stability is typically found to be better than 10 ps in any 200 μ s-interval when averaged over a number of events, with many stable to 5 ps. This level of timing instability contributes less than a 0.05 ppm systematic error on ω_a .

3. Beam dynamics

The behaviour of the beam in the $(g - 2)$ storage ring directly affects the measurement of a_μ . Since the detector acceptance for decay electrons depends on the radial coordinate of the muon at the point where it decays, coherent radial motion of the stored beam can produce an amplitude modulation in the observed electron time spectrum. Resonances in the storage ring can cause particle losses, thus distorting the observed time spectrum, and must be avoided when choosing the operating parameters of the ring. Care must be taken in setting the frequency of coherent radial beam motion, the ‘coherent betatron oscillation’ (CBO) frequency, which lies close to the second harmonic of $f_a = \omega_a/(2\pi)$. If f_{CBO} is too close to $2f_a$ the difference frequency $f_- = f_{\text{CBO}} - f_a$ complicates the extraction of f_a from the data and can introduce a significant systematic error.

A pure quadrupole electric field provides a linear restoring force in the vertical direction, and the combination of the (defocusing) electric field and the central magnetic field provides a linear restoring force in the radial direction. The $(g - 2)$ ring is a weak focusing ring [43–45] with the field index

$$n = \frac{\kappa R_0}{\beta B_0}, \quad (23)$$

where κ is the electric quadrupole gradient. For a ring with a uniform vertical dipole magnetic field and a uniform quadrupole field that provides vertical focusing covering the full azimuth, the stored particles undergo simple harmonic motion called betatron oscillations, in both the radial and vertical dimensions.

The horizontal and vertical motion are given by

$$x = x_e + A_x \cos\left(\nu_x \frac{s}{R_0} + \delta_x\right) \quad \text{and} \quad y = A_y \cos\left(\nu_y \frac{s}{R_0} + \delta_y\right), \quad (24)$$

where s is the arc length along the trajectory and $R_0 = 7112$ mm is the radius of the central orbit in the storage ring. The horizontal and vertical tunes are given by $\nu_x = \sqrt{1 - n}$ and $\nu_y = \sqrt{n}$. Several n -values were used in E821 for data acquisition: $n = 0.137$, 0.142 and 0.122. The horizontal and vertical betatron frequencies are given by

$$f_x = f_C \sqrt{1 - n} \simeq 0.929 f_C \quad \text{and} \quad f_y = f_C \sqrt{n} \simeq 0.37 f_C, \quad (25)$$

where f_C is the cyclotron frequency and the numerical values assume that $n = 0.137$. The corresponding betatron wavelengths are $\lambda_{\beta_x} = 1.08(2\pi R_0)$ and $\lambda_{\beta_y} = 2.7(2\pi R_0)$. It is important that the betatron wavelengths are not simple multiples of the circumference, as this minimizes the ability of ring imperfections and higher multipoles to drive resonances that would result in particle losses from the ring.

The field index, n , also determines the acceptance of the ring. The maximum horizontal and vertical angles of the muon momentum are given by

$$\theta_{\text{max}}^x = \frac{x_{\text{max}} \sqrt{1 - n}}{R_0} \quad \text{and} \quad \theta_{\text{max}}^y = \frac{y_{\text{max}} \sqrt{n}}{R_0}, \quad (26)$$

where $x_{\text{max}}, y_{\text{max}} = 45$ mm is the radius of the storage aperture. For a betatron amplitude A_x or A_y less than 45 mm, the maximum angle is reduced, as can be seen from the above equations.

Resonances in the storage ring will occur if $L\nu_x + M\nu_y = N$, where L , M and N are integers, which must be avoided in choosing the operating value of the field index. These resonances form straight lines on the tune plane shown in figure 17, which shows resonance lines up to fifth order. The operating point lies on the circle $\nu_x^2 + \nu_y^2 = 1$.




Figure 17. The tune plane showing the three operating points used during our three years of running.

For a ring with discrete quadrupoles, the focusing strength changes as a function of azimuth and the equation of motion looks like an oscillator whose spring constant changes as a function of azimuth s . The motion is described by

$$x(s) = x_e + A\sqrt{\beta(s)} \cos(\psi(s) + \delta), \quad (27)$$

where $\beta(s)$ is one of the three Courant–Snyder parameters [44]. The layout of the storage ring is shown in figure 10. The fourfold symmetry of the quadrupoles was chosen because it provided quadrupole-free regions for the kicker, traceback chambers, fibre monitors and trolley garage; but the most important benefit of fourfold symmetry over the twofold used at CERN [38] is that $\sqrt{\beta_{\max}/\beta_{\min}} = 1.03$. The twofold symmetry used at CERN [38] gives $\sqrt{\beta_{\max}/\beta_{\min}} = 1.15$. The CERN magnetic field had significant non-uniformities on the outer portion of the storage region, which when combined with the 15% beam ‘breathing’ from the quadrupole lattice made it much more difficult to determine the average magnetic field weighted by the muon distribution (equation (20)).

The detector acceptance depends on the radial position of the muon when it decays, so that any *coherent* radial beam motion will amplitude modulate the decay e^\pm distribution. The principal frequency will be the ‘coherent betatron frequency,’

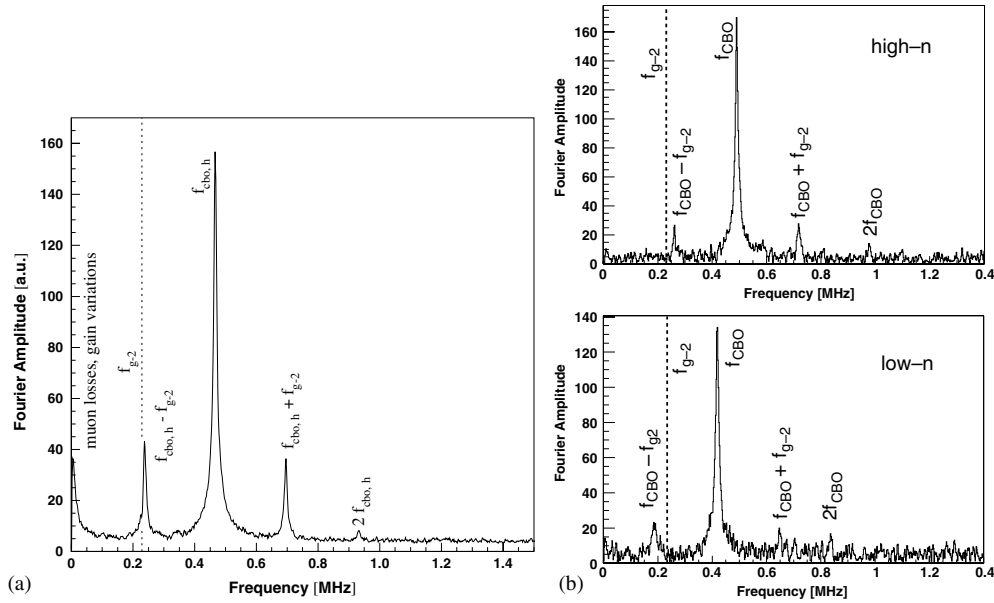
$$f_{\text{CBO}} = f_c - f_x = (1 - \sqrt{1 - n})f_c \simeq 470 \text{ kHz}, \quad (28)$$

which is the frequency at which a single fixed detector sees the beam coherently moving back and forth radially. This CBO frequency is close to the second harmonic of the ($g - 2$) frequency, $f_a = \omega_a/2\pi \simeq 228 \text{ Hz}$.

An alternative way of thinking about the CBO motion is to view the ring as a spectrometer where the inflector exit is imaged at each successive betatron wavelength, λ_{β_x} . In principle, an inverted image appears at half a betatron wavelength; but the radial image is spoiled by the $\pm 0.3\%$ momentum dispersion of the ring. A given detector will see the beam move radially with the CBO frequency, which is also the frequency at which the horizontal waist precesses

Table 3. Frequencies in the $(g - 2)$ storage ring, assuming that the quadrupole field is uniform in azimuth and that $n = 0.137$.

Quantity	Expression	Frequency (MHz)	Period
f_a	$\frac{e}{2\pi mc} a_\mu B$	0.228	$4.37 \mu\text{s}$
f_C	$\frac{1}{2\pi R_0}$	6.7	149 ns
f_x	$\sqrt{1-n} f_C$	6.23	160 ns
f_y	$\sqrt{n} f_C$	2.48	402 ns
f_{CBO}	$f_C - f_x$	0.477	$2.10 \mu\text{s}$
f_{VW}	$f_C - 2f_y$	1.74	$0.574 \mu\text{s}$

**Figure 18.** The Fourier transform to the residuals from a fit to the five-parameter function, showing clearly the coherent beam frequencies. (a) is from 2000, when the CBO frequency was close to $2\omega_a$, and (b) shows the Fourier transform for the two n -values used in the 2001 run period.

around the ring. Since there is no dispersion in the vertical dimension, the vertical waist (VW) is reformed every half wavelength $\lambda_{\beta_y}/2$. A number of frequencies in the ring are tabulated in table 3.

The CBO frequency and its sidebands are clearly visible in the Fourier transform to the residuals from a fit to the five-parameter fitting function, equation (8), and are shown in figure 18. The vertical waist frequency is barely visible. In 2000, the quadrupole voltage was set such that the CBO frequency was uncomfortably close to the second harmonic of f_a , thus placing the difference frequency $f_- = f_{\text{CBO}} - f_a$ next to f_a . This nearby sideband forced us to work very hard to understand the CBO and how its related phenomena affect the value of ω_a obtained from fits to the data. In 2001, we carefully set f_{CBO} at two different values, one well above, the other well below $2f_a$, which greatly reduced this problem.

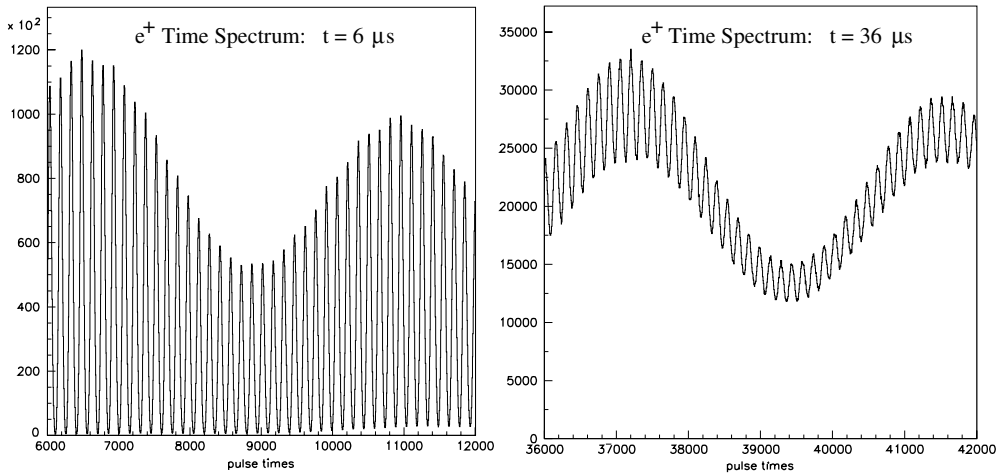


Figure 19. The time spectrum of a single calorimeter soon after injection. The spikes are separated by the cyclotron period of 149 ns.

3.1. Monitoring the beam profile

Three tools are available to us to monitor the muon distribution. Study of the beam de-bunching after injection yields information on the distribution of equilibrium radii in the storage ring. The FSDs provide information on the vertical centroid of the beam. The wire chamber system and the fibre beam monitors, described above, also provide valuable information on the properties of the stored beam.

The beam bunch that enters the storage ring has a time spread with $\sigma \simeq 23$ ns, while the cyclotron period is 149 ns. The momentum distribution of stored muons produces a corresponding distribution in radii of curvature. The distributions depend on the phase-space acceptance of the ring, the phase space of the beam at the injection point and the kick given to the beam at injection. The narrow horizontal dimension of the beam at the injection point, about 18 mm, restricts the stored momentum distribution to about $\pm 0.3\%$. As the muons circle the ring, the muons at smaller radius (lower momentum) eventually pass those at larger radius repeatedly after multiple transits around the ring, and the bunch structure largely disappears after $60 \mu\text{s}$. This de-bunching can be seen in figure 19 where the signal from a single detector is shown at two different times following injection. The bunched beam is seen very clearly in the left figure, with the 149 ns cyclotron period being obvious. The slow amplitude modulation comes from the ($g - 2$) precession. By $36 \mu\text{s}$ the beam has largely de-bunched.

Only muons with orbits centred at the central radius have the ‘magic’ momentum, so knowledge of the momentum distribution, or equivalently the distribution of equilibrium radii, is important in determining the correction to ω_a caused by the radial electric field used for vertical focusing. Two methods of obtaining the distribution of equilibrium radii from the beam debunching are employed in E821. One method uses a model of the time evolution of the bunch structure. A second, alternative procedure uses modified Fourier techniques [52]. The results from these analyses are shown in figure 20. The discrete points were obtained using the model, and the dotted curve was obtained with the modified Fourier analysis. The two analyses agree. The measured distribution is used both in determining the average magnetic field seen by the muons and the radial electric field correction discussed below.

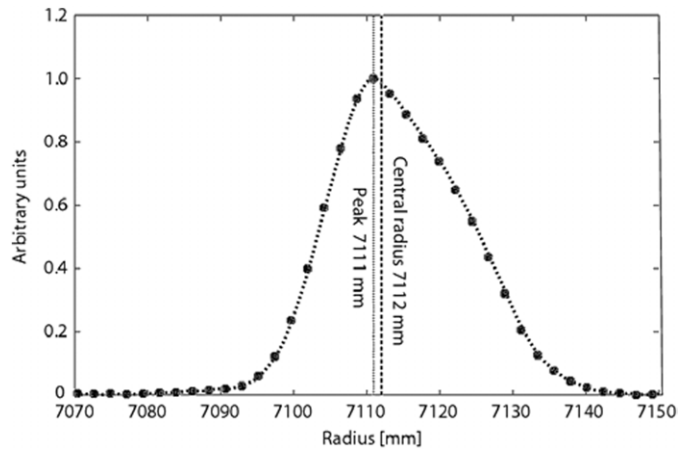


Figure 20. The distribution of equilibrium radii obtained from the beam de-bunching. The solid circles are from a de-bunching model fit to the data, and the dotted curve is obtained from a modified Fourier analysis.




Figure 21. (a) The horizontal beam centroid motion with beam scraping and without, using data from the scintillating fibre hodoscopes; note the tune change between the two. (b) A Fourier transform of the pulse from a single horizontal fibre, which shows clearly the vertical waist motion, as well as the vertical tune. The presence of stored protons is clearly seen in this frequency spectrum.

The scintillating-fibre monitors show clearly the vertical and horizontal tunes as expected. In figure 21, the horizontal beam centroid motion is shown, with the quadrupoles powered asymmetrically during scraping and then symmetrically after scraping. A Fourier transform of the latter signal shows the expected frequencies, including the cyclotron frequency of protons stored in the ring. The traceback system also sees the CBO motion.

3.2. Corrections to ω_a : pitch and radial electric field

If the velocity is not transverse to the magnetic field, or if a muon is not at γ_{magic} , the difference frequency is modified as indicated in equation (19). Thus the measured frequency ω_a must

be corrected for the effect of a radial electric field (because of the $\vec{\beta} \times \vec{E}$ term) and for the vertical pitching motion of the muons (which enters through the $\vec{\beta} \cdot \vec{B}$ term). These are the only corrections made to the ω_a data. We sketch the derivation for E821, below [53]. For a general derivation the reader is referred to [54, 55].

First we calculate the effect of the electric field, for the moment neglecting the $\vec{\beta} \cdot \vec{B}$ term. If the muon momentum is different from the magic momentum, the precession frequency is given by

$$\omega'_a = \omega_a \left[1 - \beta \frac{E_r}{B_y} \left(1 - \frac{1}{a_\mu \beta^2 \gamma^2} \right) \right]. \quad (29)$$

Using $p = \beta \gamma m = (p_m + \Delta p)$, after some algebra one finds

$$\frac{\omega'_a - \omega_a}{\omega_a} = \frac{\Delta \omega_a}{\omega_a} = -2 \frac{\beta E_r}{B_y} \left(\frac{\Delta p}{p_m} \right). \quad (30)$$

Thus the effect of the radial electric field reduces the observed frequency from the simple frequency ω_a given in equation (14). Now

$$\frac{\Delta p}{p_m} = (1 - n) \frac{\Delta R}{R_0} = (1 - n) \frac{x_e}{R_0}, \quad (31)$$

where x_e is the muon's equilibrium radius of curvature relative to the central orbit. The electric quadrupole field is

$$E = \kappa x = \frac{n \beta B_y}{R_0} x. \quad (32)$$

We obtain

$$\frac{\Delta \omega}{\omega} = -2n(1 - n) \beta^2 \frac{x x_e}{R_0^2 B_y}, \quad (33)$$

so clearly the effect of muons not at the magic momentum is to lower the observed frequency. For a quadrupole focusing field plus a uniform magnetic field, the time average of x is just x_e , so the electric field correction is given by

$$C_E = \frac{\Delta \omega}{\omega} = -2n(1 - n) \beta^2 \frac{\langle x_e^2 \rangle}{R_0^2 B_y}, \quad (34)$$

where $\langle x_e^2 \rangle$ is determined from the fast-rotation analysis (see figure 19). The uncertainty on $\langle x_e^2 \rangle$ is added in quadrature with the uncertainty in the placement of the quadrupoles of $\delta R = \pm 0.5$ mm (± 0.01 ppm) and with the uncertainty in the mean vertical position of the beam, ± 1 mm (± 0.02 ppm). For the low- n 2001 sub-period, $C_E = 0.47 \pm 0.054$ ppm.

The vertical betatron oscillations of the stored muons lead to $\vec{\beta} \cdot \vec{B} \neq 0$. Since the $\vec{\beta} \cdot \vec{B}$ term in equation (18) is quadratic in the components of $\vec{\beta}$, its contribution to ω_a will not generally average to zero. Thus the spin precession frequency has a small dependence on the betatron motion of the beam. It turns out that the only significant correction comes from the vertical betatron oscillation; therefore it is called the pitch correction (see equation (19)). As the muons undergo vertical betatron oscillations, the 'pitch' angle between the momentum and the horizontal (see figure 22) varies harmonically as $\psi = \psi_0 \cos \omega_y t$, where ω_y is the vertical betatron frequency $\omega_y = 2\pi f_y$, given in equation (25). In the approximation that all muons are at the magic γ , we set $a_\mu - 1/(\gamma^2 - 1) = 0$ in equation (19) and obtain

$$\vec{\omega}'_a = -\frac{q}{m} \left[a_\mu \vec{B} - a_\mu \left(\frac{\gamma}{\gamma + 1} \right) (\vec{\beta} \cdot \vec{B}) \vec{\beta} \right], \quad (35)$$

where the prime indicates the modified frequency as it did in the discussion of the radial electric field given above, and $\vec{\omega}_a = -(q/m) a_\mu \vec{B}$. We adopt the (rotating) coordinate system shown




Figure 22. The coordinate system of the pitching muon. The angle ψ varies harmonically. The vertical direction is \hat{y} and \hat{z} is the azimuthal (beam) direction.

in figure 22, where $\vec{\beta}$ lies in the zy -plane, z being the direction of propagation and y being vertical in the storage ring. Assuming $\vec{B} = \hat{y}B_y$, $\vec{\beta} = \hat{z}\beta_z + \hat{y}\beta_y = \hat{z}\beta \cos \psi + \hat{y}\beta \sin \psi$, we find

$$\vec{\omega}'_a = -\frac{q}{m} [a_\mu \hat{y} B_y - a_\mu \left(\frac{\gamma}{\gamma + 1} \right) \beta_y B_y (\hat{z}\beta_z + \hat{y}\beta_y)]. \quad (36)$$

The small-angle approximation $\cos \psi \simeq 1$ and $\sin \psi \simeq \psi$ gives the component equations

$$\omega'_{ay} = \omega_a \left[1 - \left(\frac{\gamma - 1}{\gamma} \right) \psi^2 \right] \quad (37)$$

and

$$\omega'_{az} = -\omega_a \left(\frac{\gamma - 1}{\gamma} \right) \psi. \quad (38)$$

Rather than use the components given above, we can resolve ω'_a into components along the coordinate system defined by $\vec{\beta}$ (see figure 22) using the standard rotation formula. The transverse component of ω' is given by

$$\omega_\perp = \omega'_{ay} \cos \psi - \omega'_{az} \sin \psi. \quad (39)$$

Using the small-angle expansion for $\cos \psi \simeq 1 - \psi^2/2$, we find

$$\omega_\perp \simeq \omega_a \left[1 - \frac{\psi^2}{2} \right]. \quad (40)$$

As can be seen from table 3, the pitching frequency ω_y is an order of magnitude larger than the frequency ω_a , so that in one $(g - 2)$ period ω_\perp oscillates more than ten times, thus averaging out its effect on ω'_a so $\omega'_a \simeq \omega_\perp$. Thus

$$\omega_a \simeq -\frac{q}{m} a_\mu B_y \left(1 - \frac{\psi^2}{2} \right) = -\frac{q}{m} a_\mu B_y \left(1 - \frac{\psi_0^2 \cos^2 \omega_y t}{2} \right). \quad (41)$$

Taking the time average yields a pitch correction

$$C_p = -\frac{\langle \psi^2 \rangle}{2} = -\frac{\langle \psi_0^2 \rangle}{4} = -\frac{n \langle y^2 \rangle}{4 R_0^2}, \quad (42)$$

where we have used equation (26) $\langle \psi_0^2 \rangle = n \langle y^2 \rangle / R_0^2$. The quantity $\langle y_0^2 \rangle$ was both determined experimentally and from simulations. For the 2001 period, $C_p = 0.27 \pm 0.036$ ppm, the amount the precession frequency is lowered from that given in equation (20) because $\vec{\beta} \cdot \vec{B} \neq 0$.

We see that both the radial electric field and the vertical pitching motion *lower* the observed frequency from the simple difference frequency $\omega_a = (e/m)a_\mu B$, which enters into our determination of a_μ using equation (16). Therefore our observed frequency must be *increased* by these corrections to obtain the measured value of the anomaly. Note that if $\omega_y \simeq \omega_a$ the situation is more complicated, with a resonance behaviour that is discussed in [54, 55].

4. Analysis of the data for ω_p and ω_a

In the data analysis for E821, great care was taken to insure that the results were not biased by previous measurements or the theoretical value expected from the Standard Model. This was achieved by a blind analysis which guaranteed that no single member of the collaboration could calculate the value of a_μ before the analysis was complete. Two frequencies, ω_p , the Larmor frequency of a free proton which is proportional to the B field, and ω_a , the frequency that muon spin precesses relative to its momentum, are measured. The analysis was divided into two separate efforts, ω_a , ω_p , with no collaboration member permitted to work on the determination of both frequencies.

In the first stage of each year's analysis, each independent ω_a (or ω_p) analyzer presented intermediate results with his own concealed offset on ω_a (or ω_p). Once the independent analyses of ω_a appeared to be mutually consistent, an offset common to all independent ω_a analyses was adopted, and a similar step was taken for the independent analyses of ω_p . The ω_a offsets were kept strictly concealed, especially from the ω_p analyzers. Similarly, the ω_p offsets were kept strictly concealed, especially from the ω_a analyzers. The nominal values of ω_a and ω_p were known at best to many ppm error, much larger than the eventual result, and could not be guessed with any precision. No one person was allowed to know about both offsets, and it was therefore impossible to calculate the value of a_μ until the offsets were publicly revealed, after all analyses were declared to be complete.

For each of the four yearly data sets, 1998–2001, there were between four and five largely independent analyses of ω_a and two independent analyses of ω_p . Typically, on ω_a there were one or two physicists conducting independent analyses in two successive years and one on ω_p , providing continuity between the analysis of the separate data sets. Each of the fit parameters, and each of the potential sources of systematic error, were studied in great detail. For the high statistics data sets, 1999, 2000 and 2001 it was necessary in the ω_a analysis to modify the five-parameter function given in equation (8) to account for a number of small effects. Often different approaches were developed to account for a given effect, although there were common features between some of the analyses.

All intermediate results for ω_a were presented in terms of \mathfrak{R} , which is defined by

$$\omega_a = 2\pi \cdot 0.2291 \text{ MHz} \cdot [1 \pm (\mathfrak{R} \pm \Delta\mathfrak{R}) \times 10^{-6}], \quad (43)$$

where $\pm\Delta\mathfrak{R}$ is the concealed offset. Similarly, the ω_p analyzers maintained a constant offset which was strictly concealed from the rest of the collaboration.

To obtain the value of $\mathcal{R} = \omega_a/\omega_p$ to use in equation (16), the pitch and radial electric-field corrections discussed in section 3.2 were added to the measured frequency ω_a obtained from the least-squares fit to the time spectrum. Once these two corrections were made, the value of a_μ was obtained from equation (16) and published with no other changes.

4.1. The average magnetic field: the ω_p analysis

The magnetic field data consist of three separate sets of measurements: the calibration data taken before and after each running period, maps of the magnetic field obtained with the NMR

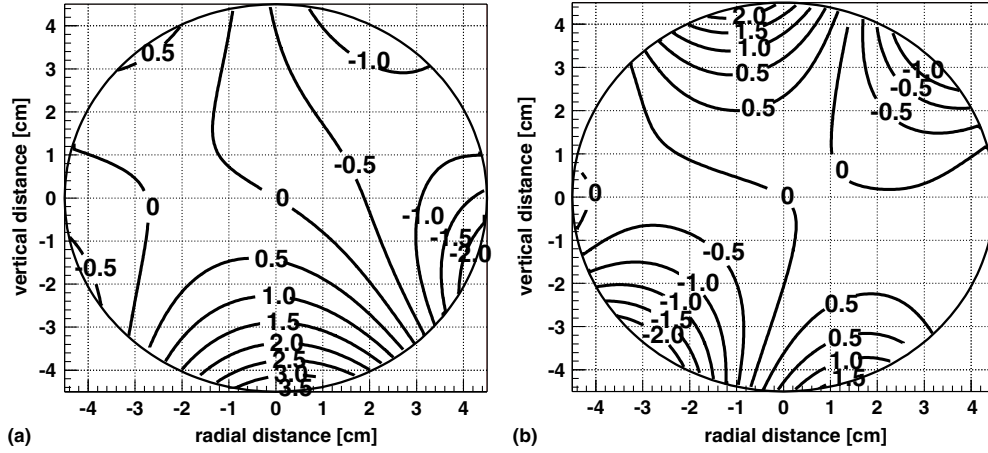


Figure 23. Homogeneity of the field (a) at the calibration position and (b) for the azimuthal average for one trolley run during the 2000 period. In both figures, the contours correspond to 0.5 ppm field differences between adjacent lines.

trolley at intervals of a few days and the field measured by each of the fixed NMR probes placed outside the vacuum chamber. It was these latter measurements, taken concurrent with the muon spin precession data and then tied to the field mapped by the trolley, which were used to determine the average magnetic field in the storage ring, and subsequently the value of ω_p to be used in equation (16).

4.1.1. Calibration of the trolley probes. The errors arising from the cross-calibration of the trolley probes with the plunging probes are caused both by the uncertainty in the relative positioning of the trolley probe and the plunging probe, and by the local field inhomogeneity. At this point in azimuth, trolley probes are fixed with respect to the frame that holds them and to the rail system on which the trolley rides. The vertical and radial positions of the trolley probes with respect to the plunging probe are determined by applying a sextupole field and comparing the change of field measured by the two probes. The field shimming at the calibration location minimizes the error caused by the relative-position uncertainty, which in the vertical and radial directions has an inhomogeneity less than 0.2 ppm cm^{-1} , as shown in figure 23(b). The full multipole components at the calibration position are given in table 4, along with the multipole content of the full magnetic field averaged over azimuth. For the estimated 1 mm position uncertainty, the uncertainty on the relative calibration is less than 0.02 ppm.

The absolute calibration utilizes a probe with a spherical water sample (see figure 14(a)) [75]. The Larmor frequency of a proton in a spherical water sample is related to that of the free proton through [64,66]

$$f_L(\text{sph} - \text{H}_2\text{O}, T) = [1 - \sigma(\text{H}_2\text{O}, T)] f_L(\text{free}), \quad (44)$$

where $\sigma(\text{H}_2\text{O}, T)$ is from the diamagnetic shielding of the proton in the water molecule, determined from [65]

$$\sigma(\text{H}_2\text{O}, 34.7^\circ\text{C}) = 1 - \frac{g_p(\text{H}_2\text{O}, 34.7^\circ\text{C})}{g_J(H)} \frac{g_J(H)}{g_p(H)} \frac{g_p(H)}{g_p(\text{free})} \quad (45)$$

$$= 25.790(14) \times 10^{-6}. \quad (46)$$

Table 4. Multipoles at the outer edge of the storage volume (radius = 4.5 cm). The left-hand set is for the plunging station where the plunging probe and the calibration are inserted. The right-hand set is the multipoles obtained by averaging over azimuth for a representative trolley run during the 2000 period.

Multipole (ppm)	Calibration		Azimuthal averaged	
	Normal	Skew	Normal	Skew
Quadrupole	-0.71	-1.04	0.24	0.29
Sextupole	-1.24	-0.29	-0.53	-1.06
Octupole	-0.03	1.06	-0.10	-0.15
Decupole	0.27	0.40	0.82	0.54

The g -factor ratio of the proton in a spherical water sample to the electron in the hydrogen ground state ($g_J(H)$) is measured to 10 parts per billion (ppb) [65]. The ratio of electron to proton g -factors in hydrogen is known to 9 ppb [67]. The bound-state correction relating the g -factor of the proton bound in hydrogen to the free proton are calculated in [68, 69]. The temperature dependence of σ is corrected for using $d\sigma(\text{H}_2\text{O}, T)/dT = 10.36(30) \times 10^{-9} \text{ }^\circ\text{C}^{-1}$ [70]. The free proton frequency is determined to an accuracy of 0.05 ppm.

The fundamental constant $\lambda_+ = \mu_{\mu^+}/\mu_p$ (see equation (16)) can be computed from the hyperfine structure of muonium (the μ^+e^- atom) [66] or from the Zeeman splitting in muonium [39]. The latter experiment used the same calibration probe as was used in our ($g - 2$) experiment; however the magnetic environments of the two experiments were different, so that perturbations of the probe materials on the surrounding magnetic field differed by a few ppb between the two experiments.

The errors in the calibration procedure result both from the uncertainties on the positions of the water samples inside the trolley and the calibration probe, and from magnetic field inhomogeneities. The precise location of the trolley in azimuth and the location of the probes within the trolley are not known better than a few mm. The uncertainties in the relative calibration resulting from position uncertainties are 0.03 ppm. Temperature and power-supply voltage dependences contribute 0.05 ppm, and the paramagnetism of the O_2 in the air-filled trolley causes a 0.037 ppm shift in the field.

4.1.2. Mapping the magnetic field. During a trolley run, the value of B is measured by each probe at approximately 6000 locations in azimuth around the ring. The magnitude of the field measured by the central probe is shown as a function of azimuth in figure 24 for one of the trolley runs. The insert shows that the fluctuations in this map that appear quite sharp are in fact quite smooth, and are not noisy. The field maps from the trolley are used to construct the field profile averaged over azimuth. This contour plot for one of the field maps is shown in figure 23(b). Since the storage ring has weak focusing, the average over azimuth is the important quantity in the analysis. Because NMR is only sensitive to the magnitude of B and not to its direction, the multipole distributions must be determined from azimuthal magnetic field averages, where the field can be written as

$$B(r, \theta) = \sum_{n=0}^{n=\infty} r^n (c_n \cos n\theta + s_n \sin n\theta), \quad (47)$$

where in practice the series is limited to five terms.

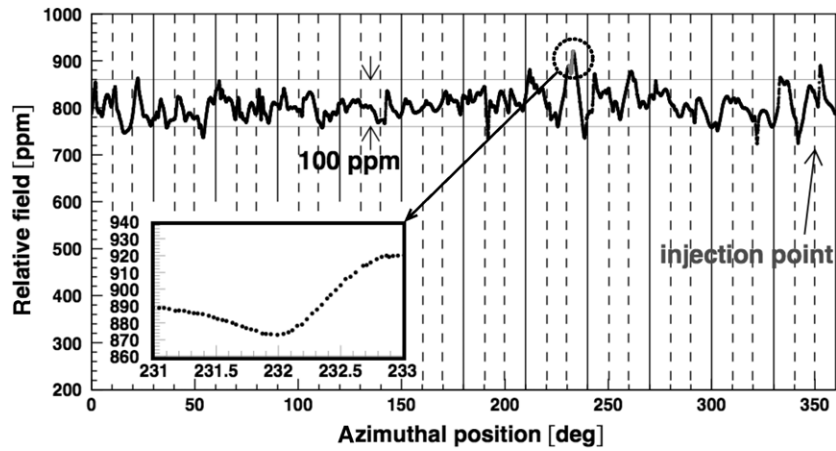


Figure 24. The magnetic field measured at the centre of the storage region versus azimuthal position. Note that while the sharp fluctuations appear to be noise, when the scale is expanded the variations are quite smooth and represent true variations in the field.




Figure 25. The difference between the average magnetic field measured by the trolley and that inferred from tracking the magnetic field with the fixed probes between trolley maps. The vertical lines show when the magnet was powered down and then back up. After each powering of the magnet, the field does not exactly come back to its previous value, so that only trolley runs taken between magnet powerings can be compared directly.

4.1.3. Tracking the magnetic field in time. During data-collection periods the field is monitored with the fixed probes. To determine how well the fixed probes permitted us to monitor the field felt by the muons, the measured field and that predicted by the fixed probes are compared for each trolley run. The results of this analysis for the 2001 running period is shown in figure 25. The rms distribution of these differences is 0.10 ppm.

4.1.4. Determination of the average magnetic field: ω_p . The value of ω_p entering into the determination of a_μ is the field profile weighted by the muon distribution. The multipoles of the field, equation (47), are folded with the muon distribution

$$M(r, \theta) = \sum [\gamma_m(r) \cos m\theta + \sigma_m(r) \sin m\theta], \quad (48)$$

Table 5. Systematic errors for the magnetic field for the different run periods.

Source of errors	1999 (ppm)	2000 (ppm)	2001 (ppm)
Absolute calibration of standard probe	0.05	0.05	0.05
Calibration of trolley probes	0.20	0.15	0.09
Trolley measurements of B_0	0.10	0.10	0.05
Interpolation with fixed probes	0.15	0.10	0.07
Uncertainty from muon distribution	0.12	0.03	0.03
Inflector fringe field uncertainty	0.20	—	—
Other ^a	0.15	0.10	0.10
Total systematic error on ω_p	0.4	0.24	0.17
Muon-averaged field (Hz) $\omega_p/2\pi$	61 791 256	61 791 595	61 791 400

^a Higher multipoles, trolley temperature and its power-supply voltage response and eddy currents from the kicker.

to produce the average field,

$$\langle B \rangle_{\mu\text{-dist}} = \int M(r, \theta) B(r, \theta) r \, dr \, d\theta, \quad (49)$$

where the moments in the muon distribution couple moment-by-moment to the multipoles of \vec{B} . Computing $\langle B \rangle$ is greatly simplified if the field is quite uniform (with small higher multipoles), and the muons are stored in a circular aperture, thus reducing the higher moments of $M(r, \theta)$. This worked quite well in E821, and the uncertainty on $\langle B \rangle$ weighted by the muon distribution was ± 0.03 ppm.

The weighted average was determined both by a tracking calculation that used a field map and calculated the field seen by each muon, and also by using the quadrupole component of the field and the beam centre determined from a fast-rotation analysis to determine the average field. These two agreed extremely well, vindicating the choice of a circular aperture and the ± 1 ppm specification on the field uniformity, that were set in the design stage of the experiment [26].

4.1.5. Summary of the magnetic field analysis. The limitations on our knowledge of the magnetic field come from measurement issues, not statistics, so in E821 the systematic errors from each of these sources had to be evaluated and understood. The results and errors are summarized in table 5.

4.2. The muon spin precession frequency: the ω_a analysis

To obtain the muon spin precession frequency ω_a given in equation (14),

$$\vec{\omega}_a = -a_\mu \frac{q\vec{B}}{m}, \quad (50)$$

which is observed as an oscillation of the number of detected electrons with time

$$N(t, E_{\text{th}}) = N_0(E_{\text{th}}) e^{-t/\gamma\tau} [1 + A(E_{\text{th}}) \cos(\omega_a t + \phi(E_{\text{th}}))], \quad (51)$$

it is necessary to

- modify the five-parameter function above to include small effects such as the coherent betatron oscillations (CBO), pulse pile-up, muon losses and gain changes, without adding so many free parameters that the statistical power for determining ω_a is compromised,

- obtain an acceptable χ_R^2 per degree of freedom in all fits, i.e. consistent with 1, where $\sigma(\chi_R^2) = \sqrt{2/\text{NDF}}$ and
- insure that the fit parameters are stable independent of the starting time of the least-square fit. This was found to be a very reliable means of testing the stability of fit parameters as a function of the time after injection.

In general, fits are made to the data for about $640 \mu\text{s}$, about 10 muon lifetimes.

4.2.1. Distribution of decay electrons. Decay electrons with the highest laboratory energies, typically $E > 1.8 \text{ GeV}$, are used in the analysis for ω_a , as discussed in section 1.1. In the (excellent) approximations that $\vec{\beta} \cdot \vec{B} = 0$ and the effect of the electric field on the spin is small, the average spin direction of the muon ensemble (i.e. the polarization vector) precesses, relative to the momentum vector, in the plane perpendicular to the magnetic field, $\vec{B} = B_y \hat{y}$, according to

$$\hat{s} = (s_{\perp} \sin(\omega_a t + \phi) \hat{x} + s_y \hat{y} + s_{\perp} \cos(\omega_a t + \phi) \hat{z}). \quad (52)$$

The unit vectors \hat{x} , \hat{y} and \hat{z} are directed along the radial, vertical and azimuthal directions respectively. The (constant) components of the spin parallel and perpendicular to the B-field are $S_y \ll 1$ and S_{\perp} , with $\sqrt{(s_{\perp}^2 + s_y^2)} = 1$. The polarization vector precesses in the plane perpendicular to the magnetic field at a rate independent of the ratio s_y/s_{\perp} . Since $a_{\mu} = (g - 2)/2 > 0$, the spin vector rotates in the same plane, but slightly faster than the momentum vector. Note that the present experiment is, apart from small detector acceptance effects, insensitive to whether the spin vector rotates faster or slower than the momentum vector rotation, and therefore it is insensitive to the sign of a_{μ} . There are small geometric acceptance effects in the detectors which demonstrate that our result is consistent with $a_{\mu} > 0$.

The value for ω_a is determined from the data using a least-square χ^2 minimization fit to the time spectrum of electron decays, $\chi^2 = \sum_i (N_i - N(t_i))^2 / N(t_i)$, where the N_i are the data points and $N(t_i)$ is the fitting function. The statistical uncertainty, in the limit where data are taken over an infinite number of muon lifetimes, is given by equation (11). The statistical figure of merit is $F_M = A\sqrt{N}$, which reaches a maximum at about $y = 0.8$, or $E \approx 2.6 \text{ GeV}/c$ (see figure 2). If all the electrons are taken above some minimum energy threshold, $F_M(E_{\text{th}})$ reaches a maximum at about $y = 0.6$, or $E_{\text{thresh}} = 1.8 \text{ GeV}/c$ (figure 3).

The spectra to be fit are in the form of histograms of the number of electrons detected versus time (see figure 26), which in the ideal case follow the five-parameter distribution function, equation (8). While this is a fairly good approximation for the E821 data sets, small modifications, mainly due to detector acceptance effects, must be made to the five-parameter function to obtain acceptable fits to the data. The most important of these effects are described in the following section.

The five-parameter function has an important, well-known invariance property. A sum of arbitrary time spectra, each obeying the five-parameter distribution and having the same λ and ω , but different values for N_0 , A and ϕ , also has the five-parameter functional form with the *same* values for λ and ω . That is,

$$\sum_i B_i e^{-\lambda(t_i - t_{i0})} (1 + A_i \cos(\omega(t_i - t_{i0}) + \phi_i)) = B e^{-\lambda t} (1 + A \cos(\omega t + \phi)). \quad (53)$$

This invariance property has significant implications for the way in which data are handled in the analysis. The final histogram of electrons versus time is constructed from a sum over the ensemble of the time spectra produced in individual spills. It extends in time from less than a few tens of microseconds after injection to $640 \mu\text{s}$, a period of about 10 muon lifetimes. To a




Figure 26. Histogram of the total number of electrons above 1.8 GeV versus time (modulo $100 \mu\text{s}$) from the 2001 μ^- data set. The bin size is the cyclotron period, $\approx 149.2 \text{ ns}$, and the total number of electrons is 3.6 billion.

very good approximation, the spectrum from each spill follows the five-parameter probability distribution. From the invariance property, the $t = 0$ point and the gains from one spill to the next do not need to be precisely aligned.

Pulse shape and gain stabilities are monitored primarily using the electron data themselves rather than laser pulses, or some other external source of pulses. The electron times and energies are given by fits to standard pulse shapes, which are established for each detector by taking an average over many pulses at late times. The variations in pulse shapes in all detectors are found to be sufficiently small as a function of energy and decay time, and contribute negligibly to the uncertainty in ω_a . In order to monitor the gains, the energy distributions integrated over one spin precession period and corrected for pile-up are collected at various times relative to injection. The high-energy portion of the energy distribution is well described by a straight line between the energy points at heights of 20% and 80% of the plateau in the spectrum (see figure 27). The position of the x-axis intercept is taken to be the endpoint energy, 3.1 GeV. It is found that the energy stability of the detectors on the ‘quiet’ side of the ring stabilize earlier in the spill than those on the noisy side of the ring. Some of the gain shift is due to the PMT gating operation. Since the noisy detectors are gated on later in the spill than the quiet ones, their gains tend to stabilize later. The starting times for the detectors are chosen so that most of their gains are calibrated to better than 0.2%. On the quiet side of the ring, data fitting can begin as early as a couple of microseconds after injection; however it is necessary to delay at least until the beam-scraping process is completed. For the noisiest detectors, just downstream of the injection point, the start of fitting may be delayed to $30 \mu\text{s}$ or more. The time histograms are then accumulated after applying the gain correction to the energy of each electron. An uncertainty in the gain stability on average over a fill affects N , λ , ϕ and A to a small extent. The result is a systematic error on ω_a on the order of 0.1 ppm.

The cyclical motion of the bunched beam around the ring (fast-rotation structure, section 3.1) results in an additional multiplicative oscillation superimposed on the usual five-parameter spectrum, which is clearly visible in figure 19. The amplitude of the modulation, which we refer to as the ‘fast-rotation’ structure, dies away with a lifetime $\approx 26 \mu\text{s}$ as the muon bunch spreads out around the ring (see section 3.1). In order to filter the fast-rotation structure

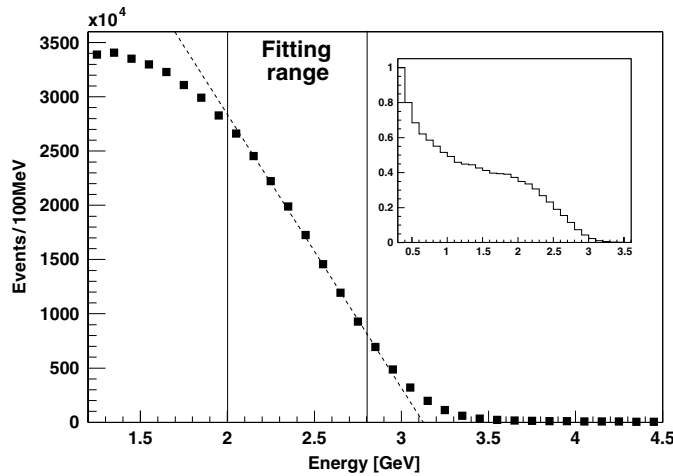


Figure 27. Typical calorimeter energy distribution, with an endpoint fit superimposed. The inset shows the full range of reconstructed energies from 0.3 to 3.5 GeV.

from the time spectrum, two steps are taken. First, the $t = 0$ point for electron data from each spill is offset by a random time chosen uniformly over one cyclotron period of 149 ns. Second, the data are formed into histograms with a bin width equal to the cyclotron period. The fast-rotation structure is also partially washed out when data are summed over all detectors, since the fast-rotation phase varies uniformly from 0 to 2π in going from one detector to the next around the ring. These measures suppress the fast-rotation structure from the spectrum by better than a factor of 100. It is confirmed in extensive simulations that these fast rotation filters have no significant effect on the derived value for ω_a . We note that the accidental overlap of pulses ('pile-up') is enhanced by the fast-rotation structure and is accounted for in the process of pile-up reconstruction, as described later in this section.

The five-parameter functional shape is unaffected by the size of the bin width in the histogram. The number of counts in each bin of the time histogram is equal to the integral of the five-parameter function over the bin width. The integrated five-parameter function has the same λ and ω as the differential five-parameter function by the invariance property. However, we note that unduly large values for the bin width will reduce the asymmetry and perhaps introduce undesirable correlations between ω_a and the other parameters of the fit. A bin width equal to the cyclotron period of about 149 ns is chosen in most of the analyses to help suppress the fast-rotation structure. This is narrow enough for binning effects to be minimal.

4.2.2. Modification of the five-parameter fitting function. In order to successfully describe the functional form of the spectrum of *detected* electrons versus time, the five-parameter function must be modified to include additional small effects of detector acceptance, muon losses, electron drift time from the point of muon decay to the point of detection, pile-up, etc. Given the desire to maintain as nearly as possible the invariance property of the five-parameter function, and the fact that some of the necessary corrections to the function turn out to be imprecisely known, every effort has been made to design the apparatus so as to minimize the needed modifications. Fortunately, most of the necessary modifications to the five-parameter function lead to additional parameters having little correlation with ω_a and therefore contribute minimally to the systematic error. The following discussion will concentrate on those modifications that contribute most to the systematic error on ω_a .

Two separate decay electrons arriving at one calorimeter with a time separation less than the time resolution of the calorimeters ($\sim 3\text{--}5$ ns, depending on the pulse heights) can be mistaken for a single pulse (pile-up). The registered energy and time in this case will be incorrect, and if not properly accounted for, it can be a source of a shift in ω_a . The registered energy and time of a pile-up pulse is approximately the sum of the energies and the energy-weighted average time, respectively, of the two pulses. The phase of the ($g - 2$) precession depends on energy, and since the pile-up pulse has an incorrectly registered energy, it has an incorrect phase. Such a phase shift would not be a problem if pile-up were the same at early and late decay times. However, the rate of pile-up is approximately proportional to the *square* of the total decay rate, and therefore the ratio of its rate to the normal data rate varies as $\sim e^{-t/\gamma\tau}$. The early-to-late change in relative rates leads to an early-to-late shift in the pile-up phase $\phi_{\text{pu}}(t)$. When the oscillating term $\cos(\omega_a t + \phi_{\text{pu}}(t) + \phi)$ is fit to $\cos(\omega_a t + \phi)$, a shift in the value for ω_a can result.

To correct for pile-up-induced phase shifts, an average pile-up spectrum is constructed and then subtracted from the data spectrum. It is constructed as follows. When a pulse exceeds a pre-assigned energy threshold, it triggers the readout of a regular electron pulse, which consists of about 80 ns of the WFD information with a pulse height reading every 2.5 ns in that time interval. This time segment is large enough to include the trigger pulse ($S_T(E, t)$) plus a significant period beyond it, which usually contains only pedestal. Occasionally, a second ('shadow', $S_S(E, t)$) pulse from another electron falls in the pedestal region. If it falls into a pre-selected time interval after the trigger pulse, it is used to construct the pile-up spectrum by combining the S_T and S_S pulses into a single pile-up pulse, D . The width of the pre-selected time interval is chosen to be equal to the minimum time separation between pulses which the detectors are capable of resolving (typically around 3 ns, however this depends slightly on the relative energies in the two peaks). The fixed delay from the trigger pulse and the window is subtracted from the time of S_S and then the pulse D is formed from the sum of the WFD samples of S_T and S_S . The pile-up spectrum is formed from $P(E, t) = D - S_T - S_S$. By construction, the rate of shadow pulses per unit time is properly normalized to the trigger rate to a good approximation. The single pulses are explicitly subtracted when forming $P(E, t)$ because two single pulses are lost for every pile-up pulse produced. To properly account for the impact on the pile-up due to the rapid variation in data rate produced by the fast-rotation structure, the time difference between the trigger pulse and the shadow pulse's time window is kept small compared with the fast-rotation period of 149 ns.

The resulting pile-up spectrum is then subtracted from the total spectrum, on average eliminating the pile-up. As figure 28 demonstrates, the constructed pile-up spectrum is in excellent agreement with the data above the endpoint energy of 3.1 GeV, where only pile-up events can occur. Below about 2.5 GeV, the pile-up spectrum becomes negative because the number of single pulses, S_T and S_S lost exceeds the number of pile-up D pulses gained.

The magnitude of the pile-up is reconstructed with an estimated uncertainty of about $\pm 8\%$. There is also an uncertainty in the phase of the pile-up. Further, the pulse-fitting procedure only recognizes shadow pulses above about 250 MeV. Based on simulations, the error due to 'unseen pile-up' (pulses below 250 MeV) was about 0.03 ppm in the 2001 data set. The systematic uncertainties were 0.014 ppm and 0.028 ppm from uncertainties in the pile-up amplitude and phase, respectively, and a total contribution of 0.08 ppm came from pile-up effects in the 2001 data set.

The coherent betatron oscillations, or oscillations in the average position and width of the stored beam, (section 3) cause unwanted oscillations in the muon decay time spectrum. These effects are generally referred to as CBO. Some of the important CBO frequencies are given in table 3. They necessitate small modifications to the five-parameter functional form of the




Figure 28. Absolute value of the constructed pile-up spectrum (solid line) and actual data spectrum (dashed line). The two curves agree well at energies above the electron endpoint at 3.1 GeV, as expected. The constructed spectrum is negative below 2.5 GeV because the number of single electrons lost to pile-up exceeds the number of pile-up pulses.

spectrum, and, like pile-up, can cause a shift in the derived value of ω_a if they are not properly accounted for in the analysis.

In particular, the functional form of the distribution of detector acceptance versus energy changes as the beam oscillates, affecting the number and average energy of the detected electrons. The result is that additional parameters need to be added to the five-parameter function to account for CBO oscillations. The effect of vertical betatron oscillations is small and dies away much faster than the horizontal oscillations, so that it can be neglected at the fit start times used in most of the analyses.

The horizontal CBO affects the spectra mainly in two ways: it causes oscillations in the number of particles because of an oscillation in the acceptance of the detectors, and it causes oscillations in the average detected energy. For a time spectrum constructed from the decay electrons in a given energy band, oscillation in the parameter N is due primarily to the oscillation in the detector acceptance. Oscillations induced in A and ϕ , on the other hand, depend primarily on the oscillation in the average energy. In either case, each of the parameters N , A and ϕ acquire small CBO-induced oscillations of the general form

$$P_i = P_{i0}[1 + B_i e^{-\lambda_{\text{CBO}}} \cos(\omega_{\text{CBO}}t + \theta_i)]. \quad (54)$$

The presence of the CBO introduces f_{CBO} , its harmonics, and the sum and difference frequencies associated with beating between f_{CBO} and $f_a = \omega_a/2\pi \approx 229.1$ kHz. If the CBO effects are *not* included in the fitting function, it will pull the value of ω_a in the fit by an amount related to how close f_{CBO} is to the second harmonic of f_a , introducing a serious systematic error (see table 3). This effect is shown qualitatively in figure 29.

The problems posed by the CBO in the fitting procedure were solved in a variety of ways in the many independent analyses. All analyses took advantage of the fact that the CBO phase varied fairly uniformly from 0 to 2π around the ring in going from one detector to the next; the CBO oscillations should tend to cancel when data from all detectors are summed together. In the 2001 data set, where the field index n was chosen to move the difference frequency $f_- = f_{\text{CBO}} - f_a$ well away from f_a , one of the analyses relied only on the factor of 9 reduction in the CBO amplitude in the sum of data over all detectors, and no CBO-induced modification to the five-parameter function was necessary.




Figure 29. The relative shift in the value obtained for ω_a as a function of the CBO frequency, when the CBO effects are neglected in the fitting function. The vertical line is at $2f_a$, and the operating points for the different data-collection periods is indicated on the curve.

The cancellation of CBO from all detectors combined would be perfect if all the detector acceptances were identical, or even if opposite pairs of detectors at 180° in the ring were identical. Imperfect cancellation is due to the reduced performance of some of the detectors and to slight asymmetries in the storage-ring geometry. This was especially true of detector number 20, where there were modifications to the vacuum chamber and whose position was displaced to accommodate the traceback chambers. The 180° symmetry is broken for detectors near the kicker because the electrons pass through the kicker plates. Also the fit start times for detectors near the injection point are inevitably later than for detectors on the other side of the ring because of the presence of the injection flash.

In addition to relying on the partial CBO cancellation around the ring, all of the other analysis approaches use a modified function in which all parameters except ω_a and λ oscillate according to equation (54). In fits to the time spectra, the CBO parameters ω_{CBO} and $\lambda_{\text{CBO}} \approx 100 \mu\text{s}$ (the frequency and lifetime of the CBO oscillations, respectively) are typically held fixed to values determined in separate studies. They were established in fits to time spectra formed with independent data from the FSDs and calorimeters, in which the amplitude of CBO modulation is enhanced by aligning the CBO oscillation phases of the individual detectors and then adding all the spectra together.

Using these fixed values for ω_{CBO} and λ_{CBO} , fits to the regular spectra are made with P_{i0} , B_i and ϕ_i in equation (54) (one set of these three for each of N , A , ϕ) as free parameters. With the addition of ω_a and λ as parameters, one obtains a total of 11 free parameters in the global fit to data. In practice, in some of the analyses, the smaller CBO parameters are held fixed or eliminated altogether.

Another correction to the fitting function is required to account for muons being lost from the storage volume before they have had a chance to decay. Such losses lead to a distortion of the spectrum which can result in an incorrect fitted value for the lifetime and a poor value for the reduced χ^2 of the fit. The correlation between the lifetime and the precession frequency is quite small, so that the loss in number of counts does not really affect the value of the precession frequency. The major concern is that the average spin phase of muons lost might be different from that of the muons which remain stored. If this were the case, the average

phase of the stored muons would shift as a function of time, leading directly to an error in the fitted value of ω_a . This uncertainty in the phase forces an inflation in the systematic error due to muon losses. Muon losses from the ring are believed to be induced by the coupling between the higher moments of the muon and the field distributions, driving resonances that result in an occasional muon striking a collimator and leaving the ring. The losses are as large as 1% per muon lifetime at early decay times, but typically settle down to less than 0.1% 200 μ s after injection. Such losses require the modification $N_0 \rightarrow N_0(1 - A_{\text{loss}}n_{\text{loss}}(t))$ in the fitting function, with A_{loss} being an additional parameter in the fit. The function n_{loss} , which represents the time distribution of lost muons, is obtained from the muon loss monitors.

An alternative analysis method to determine ω_a utilizes the so-called ‘ratio method’. Each electron event is randomly placed with equal probability into one of four time histograms, $N_1 - N_4$, each looking like the usual time spectrum, figure 26. A spectrum based on the ratio of combinations of the histograms is formed:

$$r(t_i) = \frac{N_1(t_i + \frac{1}{2}\tau_a) + N_2(t_i - \frac{1}{2}\tau_a) - N_3(t_i) - N_4(t_i)}{N_1(t_i + \frac{1}{2}\tau_a) + N_2(t_i - \frac{1}{2}\tau_a) + N_3(t_i) + N_4(t_i)}. \quad (55)$$

For a pure five-parameter distribution, keeping only the important large terms, this reduces to

$$r(t) = A \cos(\omega_a t + \phi) + \frac{1}{16} \left(\frac{\tau_a}{\tau_\mu}\right)^2, \quad (56)$$

where $\tau_a = 2\pi/\omega_a$ is an estimate (~ 10 ppm is easily good enough) of the spin precession period, and the small constant offset produced by the exponential decay is $\frac{1}{16} \left(\frac{\tau_a}{\tau_\mu}\right)^2 = 0.000287$. Construction of independent histograms N_3 and N_4 simplifies the estimates of the statistical uncertainties in the fitted parameters.

The advantage of fitting $r(t)$, as opposed to the standard technique of fitting $N(t)$, is that any slowly varying (e.g. with a period much larger than τ_a) multiplicative modulation of the five-parameter function, which includes the exponential decay of the muon itself, largely cancels in this ratio. Other examples of slow modulation are muon losses and small shifts in PMT gains due to the high rates encountered at early decay times. There are only 3 parameters, equation (56), compared with 5, equation (8), in the regular spectrum. Note that faster-varying effects, such as the CBO, will not cancel in the ratio and must be handled in ways similar to the standard analyses. One of the ratio analyses of the 2001 data set used a fitting function formed from the ratio of functions h_i , consisting of the five-parameter function modified to include parameters to correct for acceptance effects such as the CBO:

$$r_{fit} = \frac{h_1(t + \frac{1}{2}\tau_a) + h_2(t - \frac{1}{2}\tau_a) - h_3(t) - h_4(t)}{h_1(t + \frac{1}{2}\tau_a) + h_2(t - \frac{1}{2}\tau_a) + h_3(t) + h_4(t)}. \quad (57)$$

The systematic errors for three yearly data sets, 1999 and 2000 for μ^+ and 2001 for μ^- , are given in table 6.

5. The permanent electric dipole moment of the muon

If the muon were to possess a permanent electric dipole moment (see equation (2)), an extra term would be added to the spin precession equation (equation (20)):

$$\omega_{\text{EDM}} = -\frac{q}{m} \frac{\eta}{2} \left[\vec{\beta} \times \vec{B} + \frac{1}{c} \left(\vec{E} - \frac{\gamma}{\gamma+1} (\vec{\beta} \cdot \vec{E}) \vec{\beta} \right) \right], \quad (58)$$

where the dimensionless constant η is proportional to the muon electric dipole moment, $\vec{d} = \eta \frac{q}{2mc} \vec{s}$. The dominant $\vec{\beta} \times \vec{B}$ term is directed radially in the storage ring, transverse

Table 6. Systematic errors for ω_a in the 1999, 2000 and 2001 data periods. In 2001, systematic errors for the AGS background, timing shifts, E -field and vertical oscillations, beam de-bunching/randomization, binning and fitting procedure together equalled 0.11 ppm and this is indicated by ‡ in the table.

$\sigma_{\text{syst}} \omega_a$	1999 (ppm)	2000 (ppm)	2001 (ppm)
Pile-up	0.13	0.13	0.08
AGS background	0.10	0.01	‡
Lost muons	0.10	0.10	0.09
Timing shifts	0.10	0.02	‡
E -field and pitch	0.08	0.03	‡
Fitting/binning	0.07	0.06	‡
CBO	0.05	0.21	0.07
Gain changes	0.02	0.13	0.12
Total for ω_a	0.3	0.31	0.21

to $\vec{\omega}_a$. The total precession vector, $\vec{\omega} = \vec{\omega}_a + \vec{\omega}_{\text{EDM}}$, is tipped from the vertical direction by the angle $\delta = \tan^{-1} \frac{\eta\beta}{2a}$. Equation (52), with the simplification that initially $s_y = 0$, $s_{\perp} = 1$ and $\phi = 0$, is modified to

$$\hat{s} = (\cos \delta \sin \omega_a t \hat{x} + \sin \delta \sin \omega_a t \hat{y} + \cos \omega_a t \hat{z}). \tag{59}$$

The tipping of the precession plane produces an oscillation in the average vertical component of the spin which, because of the correlation between the spin and electron momentum directions, in turn causes oscillation in the average vertical component of the electron momentum. The average vertical position of the electrons, at the entrance face of the calorimeters, oscillates at angular frequency ω , 90° out of phase with the number oscillation depicted in figure 26, with an amplitude proportional to the EDM. Additionally, the magnitude of the precession frequency is increased by the EDM,

$$\omega = \sqrt{\omega_a^2 + \left(\frac{q\eta\beta B}{2m}\right)^2}. \tag{60}$$

A measure (or limit) of the muon EDM, independent of the value of a_μ , will be produced from an on-going analysis of the vertical electron motion using the FSD, PSD and traceback data from E821. Furthermore, a difference between the experimental and Standard-Model values of a_μ could be generated by an electric dipole moment.

No intrinsic electric dipole moment (EDM) has ever been experimentally detected in the muon or in any other elementary particle or atom. The existence of a permanent EDM in an elementary particle would violate both T (time reversal) and P (parity) symmetries. This is in contrast to the magnetic dipole moment (MDM), which is allowed by both of these symmetries. A non-vanishing EDM (or MDM) is consistent with C , charge conjugation symmetry, and with the combined symmetry, CPT , provided that the magnitudes of the EDM (or MDM) for a particle and its anti-particle be equal in magnitude but opposite in sign. No violation of CPT symmetry has ever been observed, and it is strictly invariant in the Standard Model. Assuming CPT invariance, the T and P violations of a non-vanishing EDM imply a violation of CP and CT , respectively. While P violation in the weak interactions is maximal and has been observed in many reactions, CP violation has been observed only in decays of neutral K and B mesons, with very small amplitudes. The violation of T invariance has been observed only in the decays of neutral K mesons [77]. An EDM large enough to be detected would require new physics beyond the Standard Model.

It is widely believed that CP violation is one of the ingredients needed to explain the baryon asymmetry of the universe, however the type and size of CP violation which has been

experimentally observed to date is much too small to account for it. Searches for EDMs are the subject of many past, current and future experiments since its detection would herald the presence of new sources of CP violation beyond the Standard Model. Indeed many proposed extensions to the Standard Model predict or at least allow the possibility of EDMs and the failure to observe them has historically served as a powerful constraint on these models.

The current experimental limit on the muon EDM, a by-product of the third CERN ($g-2$) experiment, is $d_\mu < 1 \times 10^{-18}$ e-cm (90% C.L.) [63]. This is based on the non-observation of an oscillation in the number of electrons above compared with below the calorimeter mid-planes at the precession frequency, f_a , 90° out of phase with the number oscillation.

A new muon EDM limit will be set using E821 data, by searching for an oscillation in the average vertical positions of electrons from the FSD and PSD data. The new data have the advantage over the CERN data of containing information on the shape of the vertical distribution of electrons, not just the number above or below the detector mid-plane, and the sizes of the data samples are significantly larger. Analyses of these data are in progress.

A major systematic error in the EDM measurements by CERN, and by E821 using the PSD and FSD information, arises from any misalignment between the vertical centre of the stored muons and the vertical positions of the detectors. As the spin precesses, the vertical distribution of electrons entering the calorimeters changes. In the case of the μ^- decays, high energy (in the MRF) electrons are more likely to be emitted opposite to rather than along the muon spin direction. When the spin has an inward (outward) radial component in the storage ring, the average electron will have a positive (negative) radial component of laboratory momentum, and will have to travel a greater (lesser) distance to get to the detectors. When the spin points inward, more electrons are emitted outwards, and on average the electrons travel further and spread out more in the vertical direction before getting to the detectors. The change in vertical distribution combined with a misalignment between the beam and the detectors, inevitably leads to a vertical oscillation which appears as a false EDM signal. This is the major limitation to this approach for measuring the EDM, and substantial improvements in the EDM limit in the future will require a dedicated experiment to circumvent this problem [78].

In E821, the traceback chambers are able to measure oscillation of the vertical *angles* of the electron trajectories as well as average position, providing a limit on the muon EDM which is complementary to that obtained from the FSD and PSD data. The angle information from the traceback data is less susceptible to the systematic error from beam misalignment than up-down data, however the statistical sample is smaller than the FSD or PSD data. It is expected that the final EDM result from the E821 the FSD, PSD, and traceback data will improve upon the present limit on the muon EDM by a factor of 4–5, to the level of $\approx \text{few} \times 10^{-19}$ e-cm.

In the process of deducing the experimental value of a_μ (see equation (21)) from the measured values of ω_a and ω_p , it is assumed that the muon EDM is negligibly small and can be neglected. It is found under this assumption and discussed in a later section of this manuscript that the experimental value of the anomaly differs from the Standard-Model prediction (equation (162)),

$$\Delta a_\mu = 295(88) \times 10^{-11}. \quad (61)$$

The extreme view could be taken that this apparent difference, Δa_μ , is not due to a shift in the magnetic anomaly, but rather is entirely due to a shift of the precession frequency, ω_a , caused by a non-vanishing EDM. From equation (60) the EDM would have to have the value $d_\mu = 2.4(0.4) \times 10^{-19}$ e-cm. This is a factor of $\approx 10^8$ larger than the current limit on the electron EDM. Such a large muon EDM is not predicted by even the most speculative models outside of the Standard Model, but cannot be excluded by the previous CERN experimental limit on the EDM, and also may not be definitively excluded by the eventual E821 experimental

Table 7. The frequencies ω_a and ω_p obtained from the three major data-collection periods. The radial electric field and pitch corrections applied to the ω_a values are given in the second column. Total uncertainties for each quantity are shown. The right-hand column gives the values of \mathcal{R} , where the tilde indicates the muon spin precession frequency corrected for the radial electric field and the pitching motion. The error on the average includes correlations between the systematic uncertainties of the three measurement periods.

Period	$\omega_a/(2\pi)$ (Hz)	E/pitch (ppm)	$\omega_p/(2\pi)$ (Hz)	$\mathcal{R} = \tilde{\omega}_a/\omega_p$
1999 (μ^+)	229 072.8(3)	+0.81(8)	61 791 256(25)	0.003 707 2041(51)
2000 (μ^+)	229 074.11(16)	+0.76(3)	61 791 595(15)	0.003 707 2050(25)
2001 (μ^-)	229 073.59(16)	+0.77(6)	61 791 400(11)	0.003 707 2083(26)
Average	—	—	—	0.003 707 2063(20)

result from the PSD, FSD and traceback data. Of course, if the change in the precession frequency were due to an EDM rather than a shift in the anomaly, then this would also be a very interesting indication of new physics [79].

6. Results and summary of E821

The values obtained in E821 for a_μ are given in table 1. However, the experiment measures $\mathcal{R} = \tilde{\omega}_a/\omega_p$, not a_μ directly, where the tilde over ω_a means that the pitch and radial electric field corrections have been included (see section 3.2). The fundamental constant $\lambda_+ = \mu_{\mu^+}/\mu_p$ (see equation (16)) connects the two quantities. The values obtained for \mathcal{R} are given in table 7.

The results are

$$\mathcal{R}_{\mu^-} = 0.003\,707\,2083(26) \quad (62)$$

and

$$\mathcal{R}_{\mu^+} = 0.003\,707\,2048(25). \quad (63)$$

The *CPT* theorem predicts that the magnitudes of \mathcal{R}_{μ^+} and \mathcal{R}_{μ^-} should be equal. The difference is

$$\Delta\mathcal{R} = \mathcal{R}_{\mu^-} - \mathcal{R}_{\mu^+} = (3.5 \pm 3.4) \times 10^{-9}. \quad (64)$$

Note that it is the quantity \mathcal{R} that must be compared for a *CPT* test, rather than a_μ , since the quantity λ_+ which connects them is derived from measurements of the hyperfine structure of the μ^+e^- atom (see equation (16)).

Using the latest value $\lambda_+ = \mu_{\mu^+}/\mu_p = 3.183\,345\,39(10)$ [39] gives the values for a_μ shown in table 1. Assuming *CPT* invariance, we combine all measurements of a_μ^+ and a_μ^- to obtain the ‘world average’ [26]

$$a_\mu(\text{Expt}) = 11\,659\,208.0(6.3) \times 10^{-10} \text{ (0.54 ppm)}. \quad (65)$$

The values of a_μ obtained in E821 are shown in figure 30. The final combined value of a_μ represents an improvement in precision of a factor of 13.5 over the CERN experiments. The final error of 0.54 ppm consists of a 0.46 ppm statistical component and a 0.28 systematic component.

7. The theory of the muon $g - 2$

As discussed in the introduction, the g -factor of the muon is the quantity which relates its spin \vec{s} to its magnetic moment $\vec{\mu}$ in appropriate units:

$$\vec{\mu} = g_\mu \frac{q}{2m_\mu} \vec{s} \quad \text{and} \quad \underbrace{g_\mu}_{\text{Dirac}} = 2(1 + a_\mu). \quad (66)$$

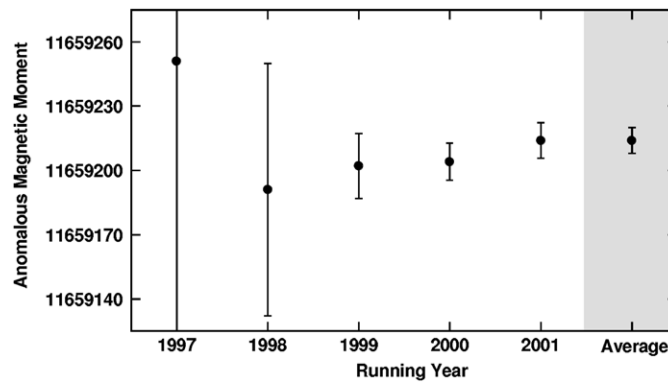


Figure 30. Results for the E821 individual measurements of a_μ by running year, together with the final average.




Figure 31. Lowest-order QED contribution. The solid blue line represents the muon in an external magnetic field (\mathbf{X} in the figure). The wavy black line represents the virtual photon.

In the Dirac theory of a charged spin-1/2 particle, $g = 2$. Quantum electrodynamics (QED) predicts deviations from the Dirac exact value, because in the presence of an external magnetic field the muon can emit and re-absorb virtual photons. The correction a_μ to the Dirac prediction is called the anomalous magnetic moment. As we have seen in the previous section, it is a quantity directly accessible to experiment.

In the following sections, we shall present a review of the various contributions to a_μ in the Standard Model, with special emphasis on the evaluations of the hadronic contributions.

7.1. The QED contributions

In QED of photons and leptons alone, the Feynman diagrams which contribute to a_μ at a given order in the perturbation theory expansion (powers of α/π) can be divided into four classes.

7.1.1. Diagrams with virtual photons and muon loops. Examples are the lowest-order contribution in figure 31 and the two-loop contributions in figure 32. In full generality, this class of diagrams consists of those with virtual photons only (wavy black lines), and those with virtual photons and internal fermion loops (solid blue loops) restricted to be of the same flavour as the external line (solid blue line) in an external magnetic field (\mathbf{X} in the diagrams). Since a_μ is a dimensionless quantity, and there is only one type of fermion mass in these graphs—with no other mass scales—these contributions are purely numerical, and they are the same for the three charged leptons: $l = e, \mu, \tau$. Indeed, $a_l^{(2)}$ from figure 31 is the celebrated Schwinger result [16], mentioned previously,

$$a_l^{(2)} = \frac{1}{2} \frac{\alpha}{\pi}; \quad (67)$$




Figure 32. Two-loop QED Feynman diagrams with the same lepton flavour.




Figure 33. A few Feynman diagrams of the three-loop type. In this class the flavour of the internal fermion loops is the same as the external fermion.

while, $a_l^{(4)}$ from the seven diagrams in figure 32 (the factor of 2 in a diagram corresponds to the contribution from the *mirror* diagram) gives the result [80, 81]

$$a_l^{(4)} = \left\{ \frac{197}{144} + \frac{\pi^2}{12} - \frac{\pi^2}{2} \ln 2 + \frac{3}{4} \zeta(3) \right\} \left(\frac{\alpha}{\pi} \right)^2 \tag{68}$$

$$= -0.328\,478\,965 \dots \left(\frac{\alpha}{\pi} \right)^2. \tag{69}$$

At the three-loop level there are 72 Feynman diagrams of this type. Only a few representative examples are shown in figure 33. Quite remarkably, their total contribution is also known analytically (see [82] and references therein). They bring in transcendental numbers like $\zeta(3) = 1.202\,0569 \dots$, the Riemann zeta-function of argument 3 which already appears at the two-loop level in equation (68), as well as transcendentals of higher complexity:⁹

$$\begin{aligned} a_l^{(6)} = & \left\{ \frac{28259}{5184} + \frac{17101}{810} \pi^2 - \frac{298}{9} \pi^2 \ln 2 + \frac{139}{18} \zeta(3) - \frac{239}{2160} \pi^4 \right. \\ & + \frac{100}{3} \left[\text{Li}_4(1/2) + \frac{1}{24} (\ln^2 2 - \pi^2) \ln^2 2 \right] \\ & \left. + \frac{83}{72} \pi^2 \zeta(3) - \frac{215}{24} \zeta(5) \right\} \left(\frac{\alpha}{\pi} \right)^3 \\ = & 1.181\,241\,456 \dots \left(\frac{\alpha}{\pi} \right)^3, \end{aligned} \tag{70}$$

⁹ There is in fact an interesting relationship between the appearance of these transcendentals in perturbative quantum field theory and mathematical structures like knot theory and non-commutative geometry, which is under active study (see, e.g. [83, 84] and references therein).

where $\text{Li}_4(1/2) = 0.517479\dots$, is a particular case of the polylogarithm-function (see, e.g. [85, 86]) which often appears in loop calculations in perturbation theory, in the form of the integral representation:

$$\text{Li}_k(x) = \frac{(-1)^{k-1}}{(k-2)!} \int_0^1 \frac{dt}{t} \ln^{k-2} t \ln(1-xt), \quad k \geq 2 \quad (71)$$

$$= \sum_{n=1}^{\infty} \left(\frac{1}{n}\right)^k x^n, \quad |x| \leq 1, \quad (72)$$

while

$$\zeta(k) = \text{Li}_k(1) = \sum_{n=1}^{\infty} \frac{1}{n^k}, \quad k \geq 2. \quad (73)$$

At the four-loop level, there are 891 Feynman diagrams of this type. Some of them are already known analytically, but in general one has to resort to numerical methods for a complete evaluation. This impressive calculation, which requires many technical skills (see the chapter ‘Theory of the anomalous magnetic moment of the electron—numerical approach’ in [87] for an overall review), is under constant updating due to advances in computing technology. The most recent published value for the whole four-loop contribution, from fermions with the same flavour, gives the result (see [88] for details and references therein)

$$a_l^{(8)} = -1.7283(35) \left(\frac{\alpha}{\pi}\right)^4, \quad (74)$$

where the error is due to the present numerical uncertainties.

Notice the alternating sign of the results from the contributions of one loop to four loops, a simple feature which is not yet *a priori* understood. Also, the fact that the sizes of the $(\alpha/\pi)^n$ coefficients for $n = 1, 2, 3, 4$ remain rather small is an interesting feature, allowing one to expect that the order of magnitude of the five-loop contribution, from a total of 12 672 Feynman diagrams [88, 89], is likely to be of $\mathcal{O}(\alpha/\pi)^5 \simeq 7 \times 10^{-14}$. This magnitude is well beyond the accuracy required to compare with the present experimental results on the muon anomaly, but eventually needed for a more precise determination of the fine-structure constant α from the electron anomaly [28].

7.1.2. Vacuum polarization diagrams from electron loops. Vacuum polarization contributions result from the replacement

$$-i \frac{g_{\alpha\beta}}{q^2} \Rightarrow i \left(g_{\alpha\beta} - \frac{q_\alpha q_\beta}{q^2} \right) \frac{\Pi^{(f)}(q^2)}{q^2}, \quad (75)$$

whereby a free-photon propagator (here in the Feynman gauge) is dressed with the renormalized proper photon self-energy induced by a loop with fermion f . Formally, with $J_{\text{em}}^{(f)\alpha}(x)$ the electromagnetic current generated by the f -fermion field at a space-time point x , the proper photon self-energy is defined by the correlation function

$$i \int d^4x e^{iq \cdot x} \langle 0 | T \{ J_{\text{em}}^{(f)\alpha}(x) J_{\text{em}}^{(f)\beta}(0) \} | 0 \rangle = - (g^{\alpha\beta} q^2 - q^\alpha q^\beta) \Pi^{(f)}(q^2). \quad (76)$$

In full generality, the renormalized photon propagator involves a summation over all fermion loops as indicated in figure 34 and is given by the expression

$$D_{\alpha\beta}(q) = -i \left(g_{\alpha\beta} - \frac{q_\alpha q_\beta}{q^2} \right) \frac{1}{q^2} \sum_f \frac{1}{1 + \Pi^{(f)}(q^2)} - ia \frac{q_\alpha q_\beta}{q^4}, \quad (77)$$




Figure 34. Diagrammatic representation of the full photon propagator in equation (77).

$$\begin{aligned} \text{Im} \left(\text{wavy line with loop } f \right) &= \text{wavy line} \langle \text{triangle} \rangle \text{wavy line} \\ &= \left| \text{wavy line} \langle \text{triangle} \rangle \right|^2 \end{aligned}$$

Figure 35. Diagrammatic relation between the spectral function and the cross-section in equation (79).




Figure 36. Vacuum polarization contribution from a small internal mass.

where a is a parameter reflecting the gauge freedom in the free-field propagator ($a = 1$ in the Feynman gauge).

Since the photon self-energy is transverse in the q -momenta, the replacement in equation (75) is unaffected by the possible gauge dependence of the free-photon propagator.

On the other hand, the on-shell renormalized photon self-energy obeys a dispersion relation with a subtraction at $q^2 = 0$ (associated with the on-shell renormalization); therefore

$$\frac{\Pi^{(f)}(q^2)}{q^2} = \int_0^\infty dt \frac{1}{t} \frac{1}{t - q^2} \frac{1}{\pi} \text{Im} \Pi^{(f)}(t), \tag{78}$$

where $\frac{1}{\pi} \text{Im} \Pi^{(f)}(t)$ denotes the f -spectral function, related to the one-photon e^+e^- annihilation cross-section into f^+f^- (see the illustration in figure 35) as follows:

$$\sigma(t)_{e^+e^- \rightarrow f^+f^-} = \frac{4\pi^2\alpha}{t} \frac{1}{\pi} \text{Im} \Pi^{(f)}(t). \tag{79}$$

More specifically, at the one-loop level in perturbation theory

$$\frac{1}{\pi} \text{Im} \Pi^{(f)}(t) = \frac{\alpha}{\pi} \frac{1}{3} \sqrt{1 - \frac{4m_f^2}{t}} \left(1 + 2 \frac{m_f^2}{t} \right) \theta(t - 4m_f^2). \tag{80}$$

The simplest example of this class of Feynman diagrams is the one in figure 36, which is the only contribution of this type at the two-loop level, with the result [90, 80]

$$a_\mu^{(4)}(m_\mu/m_e) = \left[\underbrace{\left(\frac{2}{3} \right)}_{\beta_1} \left(\frac{1}{2} \right) \ln \frac{m_\mu}{m_e} - \frac{25}{36} + \mathcal{O} \left(\frac{m_e}{m_\mu} \right) \right] \left(\frac{\alpha}{\pi} \right)^2. \tag{81}$$

Vacuum polarization contributions from fermions with a mass smaller than that of the external line ($m_e \ll m_\mu$) are enhanced by the QED short-distance logarithms of the ratio of the two

masses (the muon mass to the electron mass in this case), and are therefore very important, since $\ln(m_\mu/m_e) \simeq 5.3$. As shown in [91], these contributions are governed by a Callan–Symanzik-type equation

$$\left(m_e \frac{\partial}{\partial m_e} + \beta(\alpha)\alpha \frac{\partial}{\partial \alpha}\right) \tilde{a}_\mu\left(\frac{m_\mu}{m_e}, \alpha\right) = 0, \quad (82)$$

where $\beta(\alpha)$ is the QED-function associated with charge renormalization, and $\tilde{a}_\mu(\frac{m_\mu}{m_e}, \alpha)$ denotes the asymptotic contribution to a_μ from powers of logarithms of $\frac{m_\mu}{m_e}$ and constant terms only. This renormalization group equation is at the origin of the simplicity of the result in equation (81) for the leading term: the factor $2/3$ in front of $\ln \frac{m_\mu}{m_e}$ comes from the first term in the expansion of the QED β -function in powers of $\frac{\alpha}{\pi}$ [92]:

$$\beta(\alpha) = \frac{2}{3} \left(\frac{\alpha}{\pi}\right) + \frac{1}{2} \left(\frac{\alpha}{\pi}\right)^2 - \frac{121}{144} \left(\frac{\alpha}{\pi}\right)^3 + \mathcal{O}\left[\left(\frac{\alpha}{\pi}\right)^4\right], \quad (83)$$

while the factor $1/2$ in equation (81) is the lowest-order coefficient of α/π in equation (67), which fixes the boundary condition (at $m_\mu = m_e$) to solve the differential equation in equation (82) at the first non-trivial order in perturbation theory, i.e. $\mathcal{O}(\frac{\alpha}{\pi})^2$. Knowing the QED β -function at three loops, and a_l (from the universal class of diagrams discussed above) also at three loops, allows one to sum analytically the leading, next-to-leading and next-to-next-to-leading powers of $\ln m_\mu/m_e$ to all orders in perturbation theory [91]. Of course, these logarithms can be re-absorbed in a QED running fine-structure coupling at the scale of the muon mass $\alpha_{\text{QED}}(m_\mu)$. Historically, the first experimental evidence for the running of a coupling constant in quantum field theory comes precisely from the anomalous magnetic moment of the muon in QED, well before QCD and well before the measurement of $\alpha(M_Z)$ at LEP¹⁰ (a fact which unfortunately has been largely forgotten).

The exact analytic representation of the vacuum polarization diagram in figure 36 in terms of the Feynman parameters is rather simple:

$$a_\mu^{(4)}(m_\mu/m_e) = \left(\frac{\alpha}{\pi}\right)^2 \int_0^1 \frac{dx}{x} (1-x)(2-x) \int_0^1 dy y(1-y) \frac{1}{1 + \frac{m_e^2}{m_\mu^2} \frac{1-x}{x^2} \frac{1}{y(1-y)}}. \quad (84)$$

This simplicity (rational integrand) is a characteristic feature of the Feynman parametric integrals, which makes them rather useful for numerical integration. It has also been recently shown [93] that the Feynman parametrization when combined with a Mellin–Barnes integral representation is very well suited to obtain asymptotic expansions in ratios of mass parameters in a systematic way.

The integral in equation (84) was first computed in [94]. A compact form of the analytic result is given in [95], which we reproduce below to illustrate the typical structure of an exact analytic expression ($\rho = m_e/m_\mu$):

$$a_\mu^{(4)}(m_\mu/m_e) = \left\{ -\frac{25}{36} - \frac{1}{3} \ln \rho + \rho^2(4 + 3 \ln \rho) + \rho^4 \left[\frac{\pi^2}{3} - 2 \ln \rho \ln \left(\frac{1}{\rho} - \rho \right) - \text{Li}_2(\rho^2) \right] \right. \\ \left. + \frac{\rho}{2} (1 - 5\rho^2) \left[\frac{\pi^2}{2} - \ln \rho \ln \left(\frac{1-\rho}{1+\rho} \right) - \text{Li}_2(\rho) + \text{Li}_2(-\rho) \right] \right\} \left(\frac{\alpha}{\pi}\right)^2. \quad (85)$$

¹⁰ LEP Electroweak working group. The December 2005 version of this report can be found at <http://arxiv.org/abs/hep-ex/0511027> but the reader should note that this report is periodically updated. See <http://lepewwg.web.cern.ch/LEPEWWG> for the latest information.




Figure 37. Three-loop contribution from vacuum polarization corrections due to two electron loops.




Figure 38. Three-loop contribution from the proper fourth-order vacuum polarization corrections due to an electron loop.

Numerically, using the latest CODATA [96] and PDG [97] recommended value for the mass ratio $m_\mu/m_e = 206.768\,2838(54)$, one finds

$$a_\mu^{(4)}(m_\mu/m_e) = 1.094\,258\,3111\ (84) \left(\frac{\alpha}{\pi}\right)^2, \tag{86}$$

where the error in the last two significant figures is the one induced by the present error in the mass ratio.

At the three-loop level, there is a diagram in figure 37 generated by the insertion in the Schwinger diagram in figure 31 of the vacuum polarization due to two electron loops [98], as well as the three diagrams in figure 38 generated by the proper fourth-order vacuum polarization with one electron loop [99].

At the same order, one has to add the 14 Feynman diagrams generated by the insertion of the lowest-order vacuum polarization correction due to an electron loop in the fourth-order vertex Feynman diagrams in figure 32. This generates the diagrams of figure 39, which we have collected in five subclasses of gauge-invariant diagrams. The history and references of the earlier evaluation of these contributions can be found in the review article in [100]. The exact analytic expression from these vacuum polarization graphs, for arbitrary values of the masses, was completed in 1993 by Laporta and Remiddi [101, 103]. The expression is so lengthy that it is not even reproduced in the original papers, where instead, the asymptotic expansion in the mass ratio m_μ/m_e is given up to very high order. The first few terms in that expansion are

$$\begin{aligned} a_\mu^{(6)}(m_\mu/m_e)_{\text{vp}} = & \left(\frac{\alpha}{\pi}\right)^3 \left\{ \frac{2}{9} \ln^2\left(\frac{m_\mu}{m_e}\right) + \left[\frac{31}{27} + \frac{\pi^2}{9} - \frac{2}{3}\pi^2 \ln 2 + \zeta(3) \right] \ln\left(\frac{m_\mu}{m_e}\right) \right. \\ & + \frac{1075}{216} - \frac{25}{18}\pi^2 + \frac{5}{3}\pi^2 \ln 2 - 3\zeta(3) + \frac{11}{216}\pi^2 - \frac{1}{9}\ln^4 2 - \frac{2}{9}\pi^2 \ln^2 2 - \frac{8}{3}\text{Li}_4(1/2) \\ & \left. + \frac{m_e}{m_\mu} \left[\frac{3199}{1080}\pi^2 - \frac{16}{9}\pi^2 \ln 2 - \frac{13}{18}\pi^3 \right] + \mathcal{O}\left[\left(\frac{m_e}{m_\mu}\right)^2 \ln^2\left(\frac{m_\mu}{m_e}\right)\right] \right\}. \tag{87} \end{aligned}$$

The leading, next-to-leading and next-to-next-to-leading terms coincide with the earlier renormalization group calculation in [91]. Numerically, using the exact analytic expression




Figure 39. Three-loop contributions from vacuum polarization corrections due to an electron loop; the diagrams in each line define a subclass of gauge-invariant contributions.




Figure 40. Examples of four-loop contributions of group I.

one finds [95]

$$a_{\mu}^{(6)}(m_{\mu}/m_e)_{\text{vp}} = 1.920\,455\,130\,(33) \left(\frac{\alpha}{\pi}\right)^3, \quad (88)$$

while using the asymptotic expression in equation (87) one gets

$$a_{\mu}^{(6)}(m_{\mu}/m_e)_{\text{vp}} \simeq 1.920\dots \left(\frac{\alpha}{\pi}\right)^3, \quad (89)$$

which gives an indication of the size of the terms neglected in equation (87).

The contributions at the four-loop level with the insertion of vacuum polarization diagrams from electron loops have been carefully analysed by Kinoshita and Nio in a series of papers, the most recent being [106] where earlier references can also be found. They have classified all the possible Feynman diagrams in to three groups of gauge-invariant contributions.

- (i) Group I consists of all the possible electron-loop-type vacuum polarization insertions in the photon propagator of the Schwinger graph in figure 31. All together, this group consists of 49 diagrams. Typical examples of the Feynman graphs of this group are shown in figure 40.




Figure 41. Examples of four-loop contributions of group II.

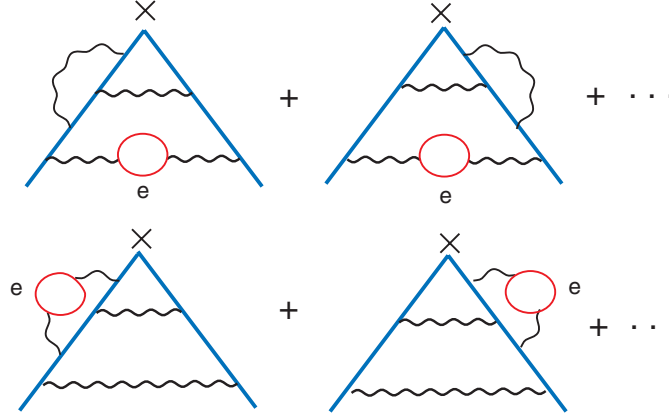


Figure 42. Examples of four-loop contributions of group III.

- (ii) Group II consists of the 90 diagrams obtained from all the possible electron-loop-type vacuum polarization insertions in the proper two-loop vertex diagrams in figure 32. Examples of the Feynman graphs of this group are shown in figure 41.
- (iii) Group III consists of the 150 diagrams obtained by the lowest-order electron-loop vacuum polarization insertion in all possible photon propagators of the proper sixth-order vertex diagrams. Examples of the Feynman graphs of this group are shown in figure 42.

All integrals of these three groups have been evaluated numerically, although some of them have also been evaluated using their asymptotic expansion in m_μ/m_e [101, 102]. The numerical results obtained in [106] are

$$a_\mu^{(8)}(m_\mu/m_e)_{\text{vpI}} = 16.720\,359\,(20) \left(\frac{\alpha}{\pi}\right)^4, \tag{90}$$

$$a_\mu^{(8)}(m_\mu/m_e)_{\text{vpII}} = -16.674\,591\,(68) \left(\frac{\alpha}{\pi}\right)^4, \tag{91}$$

$$a_\mu^{(8)}(m_\mu/m_e)_{\text{vpIII}} = 10.793\,43\,(414) \left(\frac{\alpha}{\pi}\right)^4. \tag{92}$$

The strong cancellation between the vpI and the vpII contributions, as well as their individual size, can be qualitatively understood using the Callan–Symanzik framework which describes their respective asymptotic contributions, as discussed in [91].

7.1.3. *Vacuum polarization diagrams from tau loops.* The simplest example of this class of diagrams is the Feynman graph in figure 43, which gives a contribution [99]

$$a_\mu^{(4)}(m_\mu/m_\tau) = \left[\frac{1}{45} \left(\frac{m_\mu}{m_\tau}\right)^2 + \mathcal{O}\left(\frac{m_\mu^4}{m_\tau^4} \log \frac{m_\tau}{m_\mu}\right) \right] \left(\frac{\alpha}{\pi}\right)^2. \tag{93}$$




Figure 43. Vacuum polarization contribution from a large internal mass.

In full generality, internal heavy masses in the vacuum polarization loops (heavy with respect to the external lepton line) decouple; i.e. they give a contribution which vanishes in the limit of an infinitely heavy mass. This is the reason why the muon anomaly is more sensitive to new physics than the electron anomaly, roughly by a factor $(m_\mu/m_e)^2$. This is also the reason why the QED theory for the electron anomalous magnetic moment is more precise, since the errors induced by the masses of the heavy leptons appear at a much smaller level.

From a structural point of view, the Feynman diagram in figure 43 has much in common with the *hadronic vacuum polarization* contribution which we shall discuss later and therefore it deserves some attention. In particular, one would like to understand how the simple leading behaviour in equation (93) arises. For this purpose it is convenient to reconsider the equations (75)–(80) for the case where the reference fermion is a tauon, i.e. $f = \tau$. The contribution to the anomalous magnetic moment of the muon from a τ -loop vacuum polarization insertion can then be viewed as the convolution of the τ -spectral function $1/\pi \text{Im}\Pi^{(\tau)}(t)$ with the contribution to the muon anomaly, induced by a *fictitious massive photon* with a free-propagator:

$$-i \left(g_{\alpha\beta} - \frac{q_\alpha q_\beta}{q^2} \right) \frac{1}{q^2 - t}. \quad (94)$$

The overall contribution to the muon anomaly from figure 43 can then be written as

$$a_\mu^{(4)}(m_\mu/m_\tau) = \frac{\alpha}{\pi} \int_{4m_\tau^2}^{\infty} \frac{dt}{t} \frac{1}{\pi} \text{Im}\Pi^{(\tau)}(t) \int_0^1 dx \frac{x^2(1-x)}{x^2 + \frac{t}{m_\mu^2}(1-x)}, \quad (95)$$

explicitly showing the fact that the integrand is positive and monotonically decreasing in the integration region $4m_\tau^2 \leq t \leq \infty$. (This is why the result in equation (93) is positive.) Notice also that, for large values of t/m_μ^2 , the integral over the Feynman parameter x behaves as $\frac{1}{3} \frac{m_\mu^2}{t} + \mathcal{O}[(m_\mu^2/t)^2 \log(t/m_\mu^2)]$. Since the threshold in the τ -spectral function is much larger than the muon mass ($4m_\tau^2 \gg m_\mu^2$), we can approximate the function a_μ (figure 43) by its leading asymptotic behaviour, which results in

$$a_\mu^{(4)}(m_\mu/m_\tau) \simeq \left(\frac{\alpha}{\pi} \right)^2 \frac{1}{3} \int_{4m_\tau^2}^{\infty} \frac{dt}{t} \frac{m_\mu^2}{t} \frac{1}{\pi} \text{Im}\Pi^{(\tau)}(t) = \left(\frac{\alpha}{\pi} \right)^2 \frac{1}{3} \frac{1}{15} \frac{m_\mu^2}{m_\tau^2}, \quad (96)$$

where the remaining integral over the τ -spectral function is just the slope of the τ -vacuum polarization at the origin, which can be easily calculated inserting the explicit spectral function in equation (80). This is how the leading behaviour in equation (93) arises. Numerically, using the latest CODATA- [96] and PDG-recommended [97] value for the mass ratio $m_\mu/m_\tau = 5.945\,92(97) \times 10^{-2}$, one finds

$$a_\mu^{(4)}(m_\mu/m_\tau) \simeq 0.000\,078\,564\,(26) \left(\frac{\alpha}{\pi} \right)^2, \quad (97)$$




Figure 44. Mixed vacuum polarization contribution at the three-loop level.

using the simple asymptotic result in equation (96); while using the exact expression in equation (85) (with $\rho = m_\tau/m_\mu$), one gets

$$a_\mu^{(4)}(m_\mu/m_\tau) = 0.000\,078\,064\,(25) \left(\frac{\alpha}{\pi}\right)^2, \quad (98)$$

showing that the leading asymptotic behaviour already reproduces correctly the first two non-zero digits of the exact $(\alpha/\pi)^2$ coefficient.

At the three-loop level, the τ -vacuum polarization diagrams are those of figures 37–39 with the electron replaced by the tau and hence: $m_e \Rightarrow m_\tau$. The numerical contribution, from the exact analytic calculation of Laporta and Remiddi [101, 103], gives the result

$$a_\mu^{(6)}(m_\mu/m_\tau)_{\text{vp}} = -0.001\,782\,33\,(48) \left(\frac{\alpha}{\pi}\right)^3. \quad (99)$$

To this one has to add the mixed vacuum polarization diagrams in figure 44, *a priori* enhanced because of the electron loop.

A recent analytic evaluation of the asymptotic contributions up to terms of $\mathcal{O}\left[\left(\frac{m_\mu^2}{m_\tau^2}\right)^5 \ln \frac{m_\tau^2}{m_\mu^2} \ln \frac{m_\tau^2 m_\mu^2}{m_e^4}\right]$ and $\mathcal{O}\left(\frac{m_e^2 m_\mu^2}{m_\tau^2 m_\tau^2}\right)$ has been made in [93] using a new technique, which appears to be very powerful, to obtain asymptotic expansions of the Feynman graphs. The result agrees with the asymptotic terms calculated in [104], using the more complex *method of regions* developed by Smirnov [105]. Numerically, this results in a contribution

$$a_\mu^{(6)}(m_\mu/m_e, m_\mu/m_\tau)_{\text{vp}} = 0.000\,527\,66\,(17) \left(\frac{\alpha}{\pi}\right)^3. \quad (100)$$

At the four-loop level, the vacuum polarization contributions from τ -loops that are still potentially important are those enhanced by electron vacuum polarization loops. They have been calculated numerically by Kinoshita and Nio [106, 107] with the result

$$a_\mu^{(8)}(m_\mu/m_e, m_\mu/m_\tau)_{\text{vp}} = -0.046\,188\,(37) \left(\frac{\alpha}{\pi}\right)^4. \quad (101)$$

7.1.4. Light-by-light scattering diagrams from lepton loops. It is well known that the light-by-light diagrams in QED, once the full set of gauge-invariant combinations is considered, do not require new renormalization counter-terms. Because of that, historically, it came as a big surprise to find out that the set of diagrams in figure 45, when the lepton l in the loop is an electron, produced a contribution proportional to $\log(m_\mu/m_e)$; and, in fact, with a large coefficient [108], which was first evaluated numerically. Much progress has been made since then, and this contribution is now known analytically for arbitrary values of the lepton masses [109].

From a theoretical point of view, there are some general basic features of the leading light-by-light contributions to the muon anomaly, which will also be useful later for the discussion

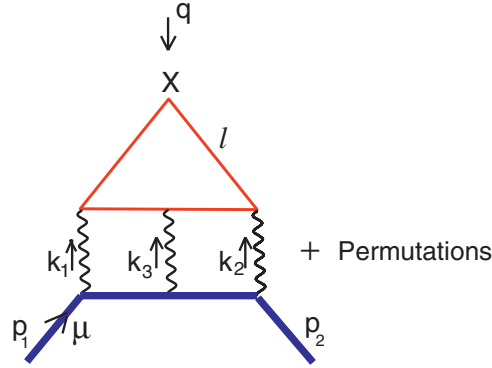


Figure 45. Light-by-light scattering contribution from an internal lepton l .

of the *hadronic light-by-light contribution*, and which we wish to recall here. Basically, one is confronted with a vertex function (see figure 45 for reference to the routing of momenta¹¹):

$$\Gamma_{\mu}^{(l)}(p_2, p_1) = ie^6 \int \frac{d^4 k_1}{(2\pi^4)} \int \frac{d^4 k_2}{(2\pi^4)} \frac{\Pi_{\mu\nu\rho\sigma}^{(l)}(q, k_1, k_3, k_2)}{k_1^2 k_2^2 k_3^2} \times \gamma^{\nu} (\not{p}_2 + \not{k}_2 - m_{\mu})^{-1} \gamma^{\rho} (\not{p}_1 - \not{k}_1 - m_{\mu})^{-1} \gamma^{\sigma}, \quad (102)$$

where $\Pi_{\mu\nu\rho\sigma}^{(l)}(q, k_1, k_3, k_2)$ denotes the off-shell photon–photon scattering amplitude ($q + k_1 + k_2 + k_3 = 0$) induced by the lepton l -loop, and $q = p_2 - p_1$. An ingenious use of gauge-invariance, first proposed by the authors of [108], makes the extraction of the anomalous magnetic moment much simpler, at the same time explicitly removing the spurious logarithmic ultraviolet divergence inherent to each individual photon–photon scattering amplitude. Current conservation (or external gauge-invariance) gives rise to the Ward identity

$$q^{\mu} \Pi_{\mu\nu\rho\sigma}^{(l)}(q, k_1, k_3, k_2) = 0, \quad (103)$$

from which one obtains by differentiation

$$\Pi_{\mu\nu\rho\sigma}^{(l)}(q, k_1, k_3, k_2) = -q^{\lambda} \frac{\partial}{\partial q^{\mu}} \Pi_{\lambda\nu\rho\sigma}^{(l)}(q, k_1, k_3, k_2). \quad (104)$$

Therefore, one can write

$$\Gamma_{\mu}^{(l)}(p_2, p_1) = q^{\lambda} \Gamma_{\mu\lambda}^{(l)}(p_2, p_1), \quad (105)$$

with

$$\Gamma_{\mu\lambda}^{(l)}(p_2, p_1) = -ie^6 \int \frac{d^4 k_1}{(2\pi^4)} \int \frac{d^4 k_2}{(2\pi^4)} \frac{\frac{\partial}{\partial q^{\mu}} \Pi_{\lambda\nu\rho\sigma}^{(l)}(q, k_1, k_3, k_2)}{k_1^2 k_2^2 k_3^2} \times \gamma^{\nu} (\not{p}_2 + \not{k}_2 - m_{\mu})^{-1} \gamma^{\rho} (\not{p}_1 - \not{k}_1 - m_{\mu})^{-1} \gamma^{\sigma}. \quad (106)$$

The muon anomaly is then extracted from this expression, via the (formally) simple projection

$$a_{\mu}^{(6)}(m_{\mu}/m_l)_{1 \times 1} = \frac{1}{48} \frac{1}{m_{\mu}} \text{tr} \{ (\not{p} + m_{\mu}) [\gamma^{\rho}, \gamma^{\sigma}] (\not{p} + m_{\mu}) \Gamma_{\rho\sigma}^{(l)}(p, p) \}. \quad (107)$$

¹¹ We use the following conventions: $\{\gamma_{\mu}, \gamma_{\nu}\} = 2g_{\mu\nu}$, with $g_{\mu\nu}$ the Minkowski metric tensor of signature $(+, -, -, -)$; $\sigma_{\mu\nu} = (i/2)[\gamma_{\mu}, \gamma_{\nu}]$, $\gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3$, and with the totally antisymmetric tensor $\epsilon_{\mu\nu\rho\sigma}$ chosen so that $\epsilon_{0123} = +1$.

The first few terms in the asymptotic expansion of $a_\mu^{(6)}(m_\mu/m_l)_{1\times 1}$ when the lepton in the loop is an electron are

$$a_\mu^{(6)}(m_\mu/m_e)_{1\times 1} = \left\{ \frac{2}{3}\pi^2 \ln \frac{m_\mu}{m_e} + \frac{2}{3} - \frac{10}{3}\pi^2 - 3\zeta(3) + \frac{59}{270}\pi^4 \right. \quad (108)$$

$$\left. + \frac{m_e}{m_\mu} \left[\frac{4}{3}\pi^2 \ln \frac{m_\mu}{m_e} + \frac{424}{9}\pi^2 - \frac{196}{3}\pi^2 \ln 2 \right], \quad (109)$$

$$\left. + \mathcal{O} \left[\left(\frac{m_e}{m_\mu} \right)^2 \ln^3 \left(\frac{m_e}{m_\mu} \right) \right] \right\} \left(\frac{\alpha}{\pi} \right)^3, \quad (110)$$

explicitly showing the large coefficient of the $\ln \frac{m_\mu}{m_e}$ term ($2/3\pi^2 = 6.58$), first noticed numerically by the authors of [108]. Numerically, using the exact expression of Laporta and Remiddi [109], one gets

$$a_\mu^{(6)}(m_\mu/m_e)_{1\times 1} = 20.947\,924\,89\,(16) \left(\frac{\alpha}{\pi} \right)^3. \quad (111)$$

The contribution when the lepton in the loop in figure 45 is heavy decouples, as it does for the case of the τ . The first few terms in the asymptotic expansion are

$$a_\mu^{(6)}(m_\mu/m_\tau)_{1\times 1} = \left\{ \left(\frac{m_\mu}{m_\tau} \right)^2 \left[-\frac{19}{16} + \frac{3}{2}\zeta(3) - \frac{161}{810} \left(\frac{m_\mu}{m_\tau} \right)^2 \ln^2 \frac{m_e}{m_\mu} \right] \right. \quad (112)$$

$$\left. + \mathcal{O} \left[\left(\frac{m_\mu}{m_\tau} \right)^4 \ln \frac{m_\mu}{m_\tau} \right] \right\} \left(\frac{\alpha}{\pi} \right)^3. \quad (113)$$

Numerically, from the exact expression [109], one obtains a contribution

$$a_\mu^{(6)}(m_\mu/m_\tau)_{1\times 1} = 0.002\,142\,83\,(69) \left(\frac{\alpha}{\pi} \right)^3. \quad (114)$$

At the four-loop level there are 180 muon vertex diagrams containing a light-by-light-scattering electron-loop sub-diagram with second-order radiative corrections. They have been calculated numerically [106] using two different techniques for the more complicated subclasses, with the total result

$$a_\mu^{(8)}(m_\mu/m_e)_{1\times 1} = 121.8431\,(59) \left(\frac{\alpha}{\pi} \right)^4. \quad (115)$$

At the same four-loop level there are still diagrams with mixed electron loops and tau loops, with one of them of the light-by-light type. Although suppressed by a $(m_\mu/m_\tau)^2$ factor, they are still enhanced by a $\ln \frac{m_\mu}{m_e}$ factor. They have also been calculated numerically by Kinoshita and Nio [106], with the result

$$a_\mu^{(8)}(m_\mu/m_e, m_\mu/m_\tau)_{1\times 1} = 0.083\,782\,(75) \left(\frac{\alpha}{\pi} \right)^4. \quad (116)$$

All together, the purely QED contributions to the muon anomalous magnetic moment, including electron- and tau-loops of the vacuum polarization type and/or of the light-by-light scattering type, are known to four loops, to an accuracy which is certainly good enough for the present comparison between theory and experiment. In table 8 we have collected all the results of the QED contributions that we have discussed. The numerical results have been obtained using the latest recommended CODATA value for the fine-structure constant [96],

$$\alpha^{-1} = 137.035\,999\,11(46) [3.3 \text{ ppb}], \quad (117)$$

Table 8. QED contributions to the muon anomalous magnetic moment. The two numerical values in the same entry correspond to the CODATA value [96] and to the more recent [28] value of α . The preliminary $(\frac{\alpha}{\pi})^5$ estimate is not listed (see text) but it is used in the final comparison of theory with experiment.

Contribution	Result in powers of $\frac{\alpha}{\pi}$	Numerical value in 10^{-11} units
$a_\mu^{(2)}$ equation (67)	$0.5 \left(\frac{\alpha}{\pi}\right)$	116 140 973.27 (0.39) 116 140 972.76 (0.08)
$a_\mu^{(4)}$ equation (69)	$-0.328 478 965 (00) \left(\frac{\alpha}{\pi}\right)^2$	
$a_\mu^{(4)}(m_\mu/m_e)$ equation (86)	$1.094 258 311 (08) \left(\frac{\alpha}{\pi}\right)^2$	
$a_\mu^{(4)}(m_\mu/m_\tau)$ equation (98)	$0.000 078 064 (26) \left(\frac{\alpha}{\pi}\right)^2$	
$a_\mu^{(4)}$ (total)	$0.765 857 410 (27) \left(\frac{\alpha}{\pi}\right)^2$	413 217.62 (0.015)
$a_\mu^{(6)}$ equation (70)	$1.181 241 46 (00) \left(\frac{\alpha}{\pi}\right)^3$	
$a_\mu^{(6)}(m_\mu/m_e)_{\text{vp}}$ equation (88)	$1.920 455 13 (03) \left(\frac{\alpha}{\pi}\right)^3$	
$a_\mu^{(6)}(m_\mu/m_\tau)_{\text{vp}}$ equation (99)	$-0.001 782 33 (48) \left(\frac{\alpha}{\pi}\right)^3$	
$a_\mu^{(6)}(m_\mu/m_e, m_\mu/m_\tau)_{\text{vp}}$ equation (100)	$0.000 527 66 (17) \left(\frac{\alpha}{\pi}\right)^3$	
$a_\mu^{(6)}(m_\mu/m_e)_{1\times 1}$ equation (111)	$20.947 924 89 (16) \left(\frac{\alpha}{\pi}\right)^3$	
$a_\mu^{(6)}(m_\mu/m_\tau)_{1\times 1}$ equation (111)	$0.002 142 83 (69) \left(\frac{\alpha}{\pi}\right)^3$	
$a_\mu^{(6)}$ (total)	$24.050 509 64 (87) \left(\frac{\alpha}{\pi}\right)^3$	30,141.90(0.001)
$a_\mu^{(8)}$ equation (74)	$-1.728 3 (35) \left(\frac{\alpha}{\pi}\right)^4$	
$a_\mu^{(8)}(m_\mu/m_e)_{\text{vp}}$ equations (90), (91), (92)	$10.839 2 (41) \left(\frac{\alpha}{\pi}\right)^4$	
$a_\mu^{(8)}(m_\mu/m_e, m_\mu/m_\tau)_{\text{vp}}$ equation (101)	$-0.046 2 (00) \left(\frac{\alpha}{\pi}\right)^4$	
$a_\mu^{(8)}(m_\mu/m_e)_{1\times 1}$ equation (115)	$121.843 1 (59) \left(\frac{\alpha}{\pi}\right)^4$	
$a_\mu^{(8)}(m_\mu/m_e, m_\mu/m_\tau)_{1\times 1}$ equation (116)	$0.083 8 (01) \left(\frac{\alpha}{\pi}\right)^4$	
$a_\mu^{(8)}$ (total)	$130.991 6 (80) \left(\frac{\alpha}{\pi}\right)^4$	381.33 (0.023)
$a_\mu^{(2+4+6+8)}$ (QED)		116 584 714.12 (0.39)(0.04) 116 584 713.61 (0.08)(0.04)

which, as discussed in the introduction, is dominated by the electron anomalous magnetic moment measurements. In fact, the most recent determination of α from the comparison between QED and a new measurement of the electron g -factor [20] gives the result [28]

$$\alpha^{-1} = 137.035 999 710(96) [0.70 \text{ ppb}], \quad (118)$$

which is quite a remarkable improvement. The numerical results using this more precise value of α as an input are also given in table 8. Use of this new value only affects the determination of the lowest-order term $a_\mu^{(2)}$ (the second-line value) and hence the total at the bottom in table 8.




Figure 46. Hadronic vacuum polarization contribution.

Notice that the errors in $a_\mu^{(4)}$ (total), $a_\mu^{(6)}$ (total) and $a_\mu^{(8)}$ (total) are very small and dominated by the error in the τ -mass, and in the μ -mass to a lesser degree. The first error in the total sum $a_\mu^{(2+4+6+8)}$ (QED), which is still the dominant error, is the one induced by the error in the determination of the fine-structure constant; the second error is the one induced by the lepton mass ratios.

The most recent estimate of the five-loop contribution by Kinoshita and Nio [88, 107] gives a result

$$a_\mu^{(10)}(m_\mu/m_e, m_\mu/m_\tau)_{\text{vp and } 1\times 1} \simeq 663 (20) \left(\frac{\alpha}{\pi}\right)^5 = 4.48 (0.14) \times 10^{-11}. \quad (119)$$

Because this estimate is not yet at the level of rigour of the lowest-order determinations, we have not included it in table 8; however, we shall take it into account when making a final comparison between theory and experiment.

The question that naturally arises is whether or not the discrepancy between the experimental result on the one hand, and the QED contribution from leptons alone which we have discussed, can be understood in terms of the extra hadronic and electroweak contributions predicted by the Standard Model. This will be the subject of the following subsections.

7.2. Contributions from hadronic vacuum polarization

All calculations of the lowest-order hadronic vacuum polarization contribution to the muon anomaly (see figure 46) are based on the spectral representation [110]

$$a_\mu^{(4)}(\text{H})_{\text{vp}} = \frac{\alpha}{\pi} \int_0^\infty \frac{dt}{t} \frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t) \int_0^1 dx \frac{x^2(1-x)}{x^2 + \frac{t}{m_\mu^2}(1-x)}, \quad (120)$$

with the hadronic spectral function $\frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t)$ related to the one photon e^+e^- annihilation cross-section into hadrons as follows:

$$\sigma(t)_{\{e^+e^- \rightarrow (\gamma) \rightarrow \text{hadrons}\}} = \frac{4\pi^2\alpha}{t} \frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t). \quad (121)$$

As already explained in the previous subsection when discussing the contribution from vacuum polarization due to a heavy lepton, equation (120) results from the replacement

$$-i \frac{g_{\alpha\beta}}{q^2} \Rightarrow i \left(g_{\alpha\beta} - \frac{q_\alpha q_\beta}{q^2} \right) \frac{\Pi^{(\text{H})}(q^2)}{q^2} \quad (122)$$

in the free-photon propagator of the lowest-order QED diagram in figure 31 by the one corrected by the proper hadronic photon self-energy in figure 46. Since the photon self-energy is transverse in the q -momenta, the replacement is unaffected by the gauge dependence in the free-photon propagator. The on-shell renormalized photon self-energy obeys a dispersion

relation, with a subtraction at $q^2 = 0$ associated with the on-shell charge renormalization; therefore

$$\frac{\Pi^{(\text{H})}(q^2)}{q^2} = \int_0^\infty \frac{dt}{t} \frac{1}{t - q^2} \frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t). \quad (123)$$

The fact that only one subtraction is needed in this dispersion relation follows from the QCD short-distance behaviour of $\Pi^{(\text{H})}(q^2)$ in the deep Euclidean region ($-q^2 \rightarrow \infty$). The rhs of equation (120) can thus be viewed as the convolution (the integral over t) of the hadronic spectral function with the contribution to the muon anomaly, induced by a *fictitious massive photon* ($m^2 = t$) with a free-propagator:

$$-i \left(g_{\alpha\beta} - \frac{q_\alpha q_\beta}{q^2} \right) \frac{1}{q^2 - t}. \quad (124)$$

This *massive photon* contribution to the muon anomaly [111], which we shall denote by $(\alpha/\pi)K^{(2)}(t/m_\mu^2)$, results in a simple Feynman parametric integral (the integral over x in equation (120)):

$$K^{(2)}\left(\frac{t}{m_\mu^2}\right) = \int_0^1 dx \frac{x^2(1-x)}{x^2 + \frac{t}{m_\mu^2}(1-x)}. \quad (125)$$

The integrand on the rhs of equation (125) is positive and monotonically decreasing in the integration region $0 \leq t \leq \infty$. This is why the lowest-order hadronic vacuum polarization contribution to the muon anomaly is positive. This integral was first evaluated analytically in [111]. A convenient representation for numerical (and analytical) evaluations, which has often been used in the literature, is the one given in [113].

For $t \geq 4m_\mu^2$, and with $\beta_\mu = \sqrt{1 - \frac{4m_\mu^2}{t}}$ and $x = \frac{1-\beta_\mu}{1+\beta_\mu}$,

$$K^{(2)}\left(\frac{t}{m_\mu^2}\right) = \frac{1}{2}x^2(2-x^2) + \frac{(1+x)^2(1+x^2)}{x^2} [\ln(1+x) - x + \frac{1}{2}x^2] + \frac{1+x}{1-x}x^2 \ln x. \quad (126)$$

Let us also remark that, for large values of t/m_μ^2 , the function $K^{(2)}(t/m_\mu^2)$ decreases as m_μ^2/t :

$$\lim_{t \rightarrow \infty} K^{(2)}(t/m_\mu^2) = \frac{1}{3} \frac{m_\mu^2}{t} + \mathcal{O}\left[\left(m_\mu^2/t\right)^2 \log(t/m_\mu^2)\right], \quad (127)$$

and this is why the bulk of the contribution to a_μ (figure 10) comes from the low-energy region—in particular from the prominent lowest ρ -resonance. The shape of the function $K^{(2)}(t/m_\mu^2)$, as well as its leading asymptotic behaviour in equation (127), are plotted in figure 47. One can see from this figure that, already at the ρ -mass value $t = M_\rho^2$, the leading asymptotic behaviour is a rather good approximation to the exact function. In fact, as first pointed out in [112], the leading asymptotic behaviour of the function $K^{(2)}(t/m_\mu^2)$ provides an *upper bound* to the size of a_μ (figure 10):

$$a_\mu^{(4)}(\text{H})_{\text{vp}} \leq \frac{\alpha}{\pi} \frac{1}{3} m_\mu^2 \int_0^\infty dt \frac{1}{t^2} \frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t) = \frac{\alpha}{\pi} \frac{1}{3} m_\mu^2 \left(\frac{\partial \Pi^{(\text{H})}(q^2)}{\partial q^2} \right)_{q^2=0}. \quad (128)$$

In other words, to lowest-order in powers of the fine-structure constant, the hadronic vacuum polarization contribution to the muon anomaly has an upper bound which is governed by the slope at the origin of the hadronic photon self-energy, a quantity which is one of the low-energy constants (a constant of $\mathcal{O}(p^6)$) of the effective chiral Lagrangian of QCD.

From a theoretical point of view, it is more convenient to convert equation (120) into an integral over the hadronic photon self-energy in the Euclidean region, i.e. an integral over




Figure 47. Behaviour of the function $K^{(2)}(t/m_\mu^2)$ in equation (125). The dotted blue line represents its leading asymptotic behaviour in equation (127).

$Q^2 = -q^2 \geq 0$ instead of an integral over the Minkowski region $t \geq 4m_\pi^2$. This is easily achieved combining equation (120) and the dispersion relation in equation (78), with the result [100, 113]:

$$a_\mu^{(4)}(\text{H})_{\text{vp}} = \frac{\alpha}{\pi} \int_0^1 dx (1-x) \left[-\Pi^{(\text{H})} \left(-\frac{x^2}{1-x} m_\mu^2 \right) \right]. \quad (129)$$

Another useful representation, which follows from this one by partial integration, is the one in terms of the Adler function

$$\mathcal{A}^{(\text{H})}(Q^2) = \int_0^\infty dt \frac{Q^2}{(t+Q^2)^2} \frac{1}{\pi} \text{Im} \Pi^{(\text{H})}(t), \quad (130)$$

with the result

$$a_\mu^{(4)}(\text{H})_{\text{vp}} = \frac{\alpha}{\pi} \frac{1}{2} \int_0^1 \frac{dx}{x} (1-x)(2-x) \mathcal{A}^{(\text{H})} \left(\frac{x^2}{1-x} m_\mu^2 \right). \quad (131)$$

These representations are better adapted to theoretical analysis and, eventually, to lattice QCD numerical evaluations. In fact, some exploratory lattice QCD work in this direction has already started [117].

It is possible to make a quick analytic estimate of $a_\mu^{(4)}(\text{H})_{\text{vp}}$ using a large- N_c QCD framework (see, e.g. [114, 115] for review papers and [116]). The spectral function in the *minimal hadronic approximation* (MHA) to large- N_c QCD consists of a lowest narrow state (the ρ) plus a perturbative QCD continuum starting at a threshold s_0 , yielding a simple parametrization for the Adler function (for three flavours):

$$\mathcal{A}^{(\text{H})}(Q^2) = \frac{2}{3} e^2 \left\{ 2f_\rho^2 M_\rho^2 \frac{Q^2}{(Q^2 + M_\rho^2)^2} + \frac{N_c}{16\pi^2} \frac{4}{3} \left(1 + \frac{3}{8} \frac{N_c \alpha_s(s_0)}{\pi} + \dots \right) \frac{Q^2}{Q^2 + s_0} \right\}, \quad (132)$$

where f_ρ denotes the ρ to vacuum coupling related to the electronic width of the ρ as follows:

$$\Gamma_{\rho \rightarrow e^+e^-} = \frac{4\pi\alpha^2}{3} f_\rho^2 M_\rho, \quad (133)$$

Table 9. Compilation of results on $a_\mu^{(4)}(\text{H})_{\text{vp}}$

Authors	Contribution to $a_\mu \times 10^{11}$
Gourdin de Rafael [118] (1969)	$6500 \pm 500 \pm (?)$
Bailey <i>et al</i> [37] (1975)	7300 ± 1000
Calmet–Narison–Perrottet–de Rafael [119] (1976)	7020 ± 800
...	...
Davier <i>et al</i> [e^+e^- -data] [122] (2003)	$6963 \pm 62_{\text{exp}} \pm 36_{\text{rad}}$
Jegerlehner [123] (2004)	6948 ± 86
de Trocóniz–Ynduráin [124] (2004)	$6944 \pm 48_{\text{exp}} \pm 10_{\text{rad}}$
Hagiwara <i>et al</i> [125] (2004)	$6924 \pm 59_{\text{exp}} \pm 24_{\text{rad}}$
Davier <i>et al</i> [τ -data] [122] (2003)	$7110 \pm 50_{\text{exp}} \pm 8_{\text{rad}} \pm 28_{SU(2)}$
de Trocóniz–Ynduráin [124] (2004)	$7027 \pm 47_{\text{exp}} \pm 10_{\text{rad}}$
Davier [e^+e^- -data] [129] (2006)	$6908 \pm 30_{\text{exp}} \pm 19_{\text{rad}} \pm 7_{\text{QCD}}$
Hagiwara <i>et al</i> [e^+e^- -data] [130] (2006)	$6894 \pm 42_{\text{exp}} \pm 18_{\text{rad}}$

and the dots stand for higher-order terms in powers of α_s . The short-distance scale s_0 denotes the threshold of the perturbative QCD (pQCD) continuum. The actual value of s_0 is fixed from the requirement that, in the chiral limit, there is no $1/Q^2$ -term in the short-distance behaviour of the Adler function, a property which follows from the operator product expansion (OPE). This leads to the constraint

$$2f_\rho^2 M_\rho^2 = \frac{N_c}{16\pi^2} \frac{4}{3} s_0 \left(1 + \frac{3}{8} \frac{N_c \alpha_s(s_0)}{\pi} + \dots \right). \quad (134)$$

Inserting the large- N_c ansatz given in (132), with $Q^2 = \frac{x^2}{1-x} m_\mu^2$ and the constraint (134) in equation (131), results in the simple estimate

$$a_\mu^{(4)}(\text{H})_{\text{vp}} \sim (57.2 \pm 15.0) \times 10^{-9}, \quad (135)$$

where we have set $f_\rho^2 = 2 \frac{F_0^2}{M_\rho^2}$, as predicted by the MHA to large- N_c QCD, with F_0 the pion to vacuum coupling in the chiral limit with: $F_0 \simeq 87 \text{ MeV}$ and $M_\rho \simeq 750 \text{ MeV}$. The error in equation (135) is a generous estimate of the systematic error of the approximations involved. One can, in principle, obtain a much more refined theoretical evaluation, but not with the high precision which is nowadays required. Regrettably, the present lattice QCD estimates [117] are not more precise than this simple large- N_c QCD estimate.

The phenomenological evaluation of the hadronic vacuum polarization to the muon anomaly, using $\sigma(e^+e^- \rightarrow \text{hadrons})$ data, goes back to earlier papers in [37, 118, 119] with the results quoted in the first part of table 9. (The question-mark error in [118] is due to ignorance at that time of the high-energy contributions, since this determination was prior to QCD.) The advent of more and more precise data on $e^+ - e^-$ annihilations, as well as the new $g_\mu - 2$ results from the Brookhaven E821 experiment, have motivated more refined re-evaluations of these hadronic contributions, which also incorporate QCD asymptotic behaviour properties. In the second part of table 9 we collect recent detailed determinations including, in particular, the re-analysis of earlier measurements by the CMD-2 detector at the VEPP-2M collider in Novosibirsk [120]. It is also possible to measure $\frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t)$ in the region $0 \leq t \leq M_\tau^2$ using hadronic τ -decays [121]. Two representative results, which partly use τ -data, and which also take into account corrections to the conserved-vector-current-symmetry (CVC) limit where the spectral functions from τ -decays and e^+e^- are directly related, are quoted in the third part of

table 9. As a test of the CVC hypothesis, the e^+e^- data can be used to predict the τ branching fractions which, unfortunately, leads to branching ratios in disagreement with the data from LEP and CLEO [122]. As additional e^+e^- and τ data have become available, this discrepancy has increased [129].

A preliminary update of the hadronic vacuum polarization, contributions reported by Davier [129] at the Tau06 conference incorporates the new CMD-2 results [126, 127], new SND results [128], as well as BaBar data for some exclusive channels other than $\pi^+\pi^-$. Hagiwara *et al* [130] have also included these new data, as well as the KLOE data, in a recent re-analysis and obtained a similar result. We quote these recent evaluations at the bottom part of table 9. The overall results quoted in this Table illustrate the progress which has been made in this field, especially in the last few years. Note that the new data allow for a significant reduction of the experimental error as well as the error due to radiative corrections.

Several remarks, however, concerning the results in table 9 are in order.

- There still remains a problematic discrepancy between the most recent e^+e^- determinations and those using τ -hadronic data. This prevents one from using a straightforward average of these different determinations.
- Various sources of isospin breaking effects, when transforming τ -hadronic spectral functions to the required e^+e^- spectral function, have been identified. A straightforward correction is the one due to the electroweak short-distance contribution to the effective four-fermion coupling $\tau \rightarrow \nu_\tau d\bar{u}$ [131–133]. Effects due to the mass differences between the π^+ and the π^0 , as well as mass and width differences between the charged and the neutral ρ , have also been considered [121, 134].
- A more problematic issue is the one of radiative corrections for a specific channel. Calculations concerning the mode $\nu_\tau\pi^-\pi^0$ have been reported in [135, 136].
- The most precise $e^+e^- \rightarrow \pi^+\pi^-$ data come from the CMD-2 detector, which are now in agreement with the data from SND, but not with the recent data from KLOE at Frascati. In fact, the two most recent compilations [129, 130] differ on the use of KLOE's results, which agree with CMD on the integral but not the shape of the pion form factor. The CMD-2 results have been corrected for leptonic and hadronic vacuum polarization effects, so that the measured final state corresponds to $\pi^+\pi^-$, including pion-radiated photons and virtual final-state QED effects.
- By contrast, hadronic τ -data are available from the four LEP detectors, CLEO, OPAL, L3 and ALEPH, and show good consistency when measuring the $\tau^- \rightarrow \nu_\tau\pi^-\pi^0$ branching ratio.
- One can hope that the forthcoming results from the KLOE detector [137] and BaBar [138], using the radiative return technique, will eventually provide the necessary consistency check of the e^+e^- input. The BELLE τ -data, with one hundred times the statistics of the LEP experiments, should also provide a test of the LEP and CLEO results.
- Finally, let us comment on the fact that the contribution from the high-energy tail to $a_\mu^{(4)}(\text{H})_{\text{vp}}$ in equation (120) is well under control because of the asymptotic freedom property of QCD. The perturbative QCD prediction is known to next-to-next-to-leading order with second-order quark mass corrections included. (See [139] and references therein for earlier calculations.)

7.2.1. Higher-order hadronic vacuum polarization contributions. There are three classes of hadronic vacuum polarization contributions of an overall $\mathcal{O}[(\alpha/\pi)^3]$ which were first considered in [119]. A first class is the set of Feynman diagrams in figure 48. They are obtained by inserting one hadronic vacuum polarization correction into each of the virtual




Figure 48. Hadronic vacuum polarization corrections to the two-loop graphs in figure 32.

photons of the purely QED fourth-order diagrams in figure 32. The diagrams in each line in figure 48 are well-defined gauge-invariant subsets. One can easily see that their overall contribution, by analogy to our discussion of the lowest-order hadronic vacuum polarization, will be governed by an integral

$$a_{\mu}^{(6)}(\text{H})_{\text{vp}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^{\infty} \frac{dt}{t} \frac{1}{\pi} \text{Im}\Pi^{(\text{H})}(t) K^{(4)}\left(\frac{t}{m_{\mu}^2}\right), \quad (136)$$

where $(\alpha/\pi)^2 K^{(4)}(t/m_{\mu}^2)$ is now the contribution to the muon anomaly from a *fictitious massive photon* in the presence of $\mathcal{O}(\alpha/\pi)$ QED corrections. The exact form of the function $K^{(4)}(t/m_{\mu}^2)$, which is a rather complicated expression, was obtained by Barbieri and Remiddi [140] and has not been fully checked since then. We reproduce in figure 49 the shape of this function and compare it with its leading asymptotic behaviour, which agrees with the one that has been extracted [141] from an independent effective field-theory calculation of anomalous dimensions of composite operators in [142],

$$\lim_{t \rightarrow \infty} K^{(4)}\left(\frac{t}{m_{\mu}^2}\right) = -\frac{23}{18} \frac{m_{\mu}^2}{t} \ln\left(\frac{t}{m_{\mu}^2}\right) + \mathcal{O}\left(\frac{m_{\mu}^2}{t}\right). \quad (137)$$

The fact that the function $K^{(4)}(t/m_{\mu}^2)$ is negative for $4\mu^2 \leq t \leq \infty$ implies that, contrary to the lowest-order case, the contribution to $a_{\mu}^{(6)}(\text{H})_{\text{vp}}$ must be *negative*.

A second class of potentially important contributions are the mixed vacuum polarization insertions in figure 50 due to an electron loop and a hadronic loop. (The similar contribution where the electron loop is replaced by a τ loop decouples as m_{μ}^2/m_{τ}^2 and is negligible at the present required level of accuracy.) This contribution, can be expressed as a double




Figure 49. Behaviour of the function $K^{(4)}(t/m_\mu^2)$ in equation (136). The dotted blue line represents its leading asymptotic behaviour in equation (137).




Figure 50. Second class of higher-order hadronic vacuum polarization contributions.

convolution [119]:

$$a_\mu^{(6)}\left(\frac{m_\mu}{m_e}, H\right)_{\text{vp}} = \frac{\alpha}{\pi} \int_{4M_\pi^2}^\infty \frac{dt}{t} \frac{1}{\pi} \text{Im}\Pi(t) \underbrace{\int_0^1 \frac{x^2(1-x)}{x^2 + \frac{t}{m_\mu^2}(1-x)} dx}_{\sim \frac{1}{3} \frac{m_\pi^2}{t}} \underbrace{(-2)\Pi\left(\frac{x^2}{1-x} \frac{m_\mu^2}{m_e^2}\right)}_{\sim (-2) \frac{\alpha}{\pi} \left(-\frac{1}{3}\right) \log \frac{m_\mu^2}{m_e^2}}, \quad (138)$$

where, in underbraces, we have indicated the leading asymptotic behaviours in m_μ^2/t and m_e^2/m_μ^2 —indicating that this contribution must be positive. The question remains: which of the two contributions from figure 48 (which, we concluded, is *negative*) and from figure 50 (which is *positive*) dominates? A simple answer to this can be found, without doing the full calculation [143], by combining the asymptotic behaviours in equations (137) and (138):

$$a_\mu^{(6)}\left(H \text{ and } \frac{m_\mu}{m_e}, H\right)_{\text{vp}} \simeq \left(\frac{\alpha}{\pi}\right)^2 \int_{4M_\pi^2}^\infty \frac{dt}{t} \frac{m_\mu^2}{t} \frac{1}{\pi} \text{Im}\Pi(t) \left[-\frac{23}{18} \ln \frac{t}{m_\mu^2} + \frac{2}{9} \ln \frac{m_\mu^2}{m_e^2} \right], \quad (139)$$

$$\sim \left(\frac{\alpha}{\pi}\right)^3 \frac{m_\mu^2}{M_\rho^2} \frac{2}{3} \frac{16\pi^2 F_0^2}{M_\rho^2} \left[-\frac{23}{18} \ln \frac{M_\rho^2}{m_\mu^2} + \frac{2}{9} \ln \frac{m_\mu^2}{m_e^2} \right], \quad (140)$$

where in the second line we have used the QCD large- N_c dominance of the ρ contribution. The overall result is indeed negative, because of the larger weight in the hadronic term. This cancellation, however, also shows the need for a more refined estimate. Indeed, the most recent evaluation, using the same e^+e^- data as in the lowest-order calculation, gives [144]

$$a_\mu^{(6)}\left(H \text{ and } \frac{m_\mu}{m_e}, H\right)_{\text{vp}} = (-97.9 \pm 0.9_{\text{exp}} \pm 0.3_{\text{rad}}) \times 10^{-11}. \quad (141)$$




Figure 51. Further contributions from higher-order hadronic vacuum polarization.




Figure 52. Hadronic light-by-light Contributions.

This determination goes down to $a_{\mu}^{(6)}(\text{H and } m_{\mu}/m_e, \text{H})_{\text{VP}} = -101 \pm 1 \times 10^{-11}$ if the τ data mentioned previously are used.

As first mentioned in [119], at $\mathcal{O}[(\alpha/\pi)^3]$, there are still two extra contributions from hadronic vacuum polarization which have to be considered. They are shown in figure 51. The evaluation of the diagram in figure 51(a) can, in principle, be achieved by lumping together in the lowest-order hadronic vacuum polarization, the effect of replacing the *one-photon* exchange cross-section in equation (121) by the *physical* cross-section which incorporates the full hadronic vacuum polarization (i.e. not just the proper hadronic self-energy). The contribution from the diagram in figure 51(b) can also be incorporated in the lowest-order evaluation, if the one-photon-exchange hadronic cross-section also takes into account final states with one-photon emission *as well as* the one photon radiative corrections to the hadronic vertex. Unfortunately, extracting the latter from the full data, involves a certain amount of model dependence, with uncertainties which are difficult to ascertain from the discussions in the literature. In principle, these uncertainties should be reflected in the *radiative correction* error quoted in equation (141).

7.3. Contributions from hadronic light-by-light scattering

These are the contributions illustrated by the diagrams in figure 52. All the estimates of these contributions made so far are model dependent. There has been progress, however, in identifying the dominant regions of virtual momenta, and in using models which incorporate some of the required features of the underlying QCD theory. The combined frameworks of QCD in the $1/N_c$ -expansion and of chiral perturbation theory [145] have been very useful in providing a guideline to these estimates.

So far, the only rigorous result in this domain has come from the observation [146] that, in large- N_c QCD and to leading-order in the chiral expansion, the dominant contribution to the muon $g-2$ from hadronic light-by-light scattering comes from the contribution of the diagrams which are one-particle (Goldstone-like) reducible. The diagrams in figure 53(a) produce a $\ln^2(\mu/m)$ -term modulated by a coefficient which has been calculated analytically [146], with




Figure 53. One Goldstone reducible diagrams in chiral perturbation theory.




Figure 54. Hadronic light-by-light from a π^0 intermediate state.

the result:

$$a_\mu[\text{figure 53}] = \left(\frac{\alpha}{\pi}\right)^3 \left\{ \frac{N_c^2}{48\pi^2} \frac{m_\mu^2}{F_\pi^2} \ln^2\left(\frac{\mu}{m}\right) + \mathcal{O}\left[\ln\left(\frac{\mu}{m}\right) + \kappa(\mu)\right] \right\}. \quad (142)$$

Here, F_π denotes the pion coupling constant in the chiral limit ($F_\pi \sim 90 \text{ MeV}$); the μ -scale in the logarithm is an arbitrary ultraviolet (UV)-scale, and the m -scale is an infrared (IR) mass, either m_μ or m_π if only the leading term in equation (142) is known.

The arbitrariness on the UV and IR scales in equation (142) would be removed, if one knew the terms linear in $\log \mu$ from figure 53(b), as well as the constant $\kappa(\mu)$ from the local counter-terms generated by figure 53(c). Unfortunately, neither the determination of the coefficient of the $\log \mu$ -term nor the constant $\kappa(\mu)$ -term, can be made in a completely model-independent way¹².

The nice feature about equation (142) is that, as discussed in [146], it plays a fundamental role in fixing the overall sign of the hadronic light-by-light scattering contribution to the muon $g - 2$. Indeed, in the various hadronic model calculations of this contribution, there appear hadronic scales (like the ρ -mass), which act as a UV-regulator and play the role of μ in equation (142). Therefore, letting the hadronic leading mass-scale become large, and provided that the model incorporates correctly the basic chiral properties of the underlying QCD theory, must reproduce the characteristic universal $\ln^2(\mu)$ behaviour with the same coefficient as in equation (142). This test, when applied to the most recent earlier calculations [148, 149] (prior to the Knecht–Nyffeler calculation in [150]), failed to reproduce the sign of the coefficient of the $\ln^2(\mu)$ -term in equation (142). The authors of [148, 149] have later found errors in their calculations which, when corrected, reproduce the effective field-theory test. Their results now agree with the Knecht–Nyffeler calculation [150] on which we next report.

The Knecht–Nyffeler calculation [150]. In full generality, the pion pole contribution to the muon anomaly has an hadronic structure, as represented by the shaded blobs in figure 54. The authors of [150] have shown that, for a large class of off-shell $\pi^0\gamma\gamma$ form factors (which includes the large- N_c QCD class), this contribution has an integral representation over two Euclidean invariants Q_1^2 and Q_2^2 associated with the two loops in figure 54:

$$a_\mu[\text{figure 54}] = \int_0^\infty dQ_1^2 \int_0^\infty dQ_2^2 \mathcal{W}(Q_1^2, Q_2^2) \mathcal{H}(Q_1^2, Q_2^2), \quad (143)$$

¹² Reference [146] provides a discussion on this point using a renormalization group approach. Essentially the same arguments have been subsequently emphasized in [147].




Figure 55. Hadronic light-by-light in the constituent quark mode.

where $\mathcal{W}(Q_1^2, Q_2^2)$ is a skeleton kernel which they calculate explicitly, and $\mathcal{H}(Q_1^2, Q_2^2)$ is a convolution of two generic $\mathcal{F}_{\pi^0\gamma^*\gamma^*}(q_1^2, q_2^2)$ form factors. In large- N_c QCD,

$$\mathcal{F}_{\pi^0\gamma^*\gamma^*}(q_1^2, q_2^2)|_{N_c \rightarrow \infty} = \sum_{ij} \frac{c_{ij}(q_1^2, q_2^2)}{(q_1^2 - M_i^2)(q_2^2 - M_j^2)}, \quad (144)$$

with the sum extended to an infinite set of narrow states. In practice, the calculation in [150] has been done by restricting the sum in equation (144) to one- and two-vector states, and fixing the polynomial $c_{ij}(q_1^2, q_2^2)$ from short-distance and long-distance QCD properties. This way, they obtained the result

$$a_\mu^{(6)}(\pi^0)_{1 \times 1} = (5.8 \pm 1.0) \times 10^{-10}, \quad (145)$$

where the error also includes an estimate of the hadronic approximation. Further inclusion of the η and η' states results in a final estimate

$$a_\mu^{(6)}(\pi^0 + \eta + \eta')_{1 \times 1} = (8.3 \pm 1.2) \times 10^{-10} \quad (146)$$

in agreement with the earlier calculations by Bijnens *et al* [148] and Hayakawa and co-workers *et al* [149] (after correcting for the overall sign).

Comment on the constituent quark model (CQM). We would like to comment on an argument which is often used in favour of the *constituent quark model* as a ‘simple’ way to fix the sign and size of the hadronic light-by-light scattering contribution to the muon $g - 2$. Since the argument has even been advocated in some publications, we feel obliged to refute it here, with the hope that it will stop further confusion. The constituent quark model contribution from the diagram in figure 55, can be easily extracted from the work of Laporta and Remiddi [109] for a heavy lepton contribution (see equation (112) above), with the result

$$a_\mu^{(6)}(\text{CQM})_{1 \times 1} = \left(\frac{\alpha}{\pi}\right)^3 N_c \frac{2}{9} \left\{ \underbrace{\left[\frac{3}{2} \zeta(3) - \frac{19}{16} \right]}_{0.616} \left(\frac{m_\mu}{M_Q}\right)^2 + \mathcal{O}\left[\left(\frac{m_\mu}{M_Q}\right)^4 \log^2\left(\frac{M_Q}{m_\mu}\right)\right] \right\}. \quad (147)$$

Seen from a low-energy-effective-field-theory point of view, the constituent quark mass M_Q in the CQM should provide the UV-regulating scale. However, the model *is not* a good effective theory of QCD and, therefore, it fails to reproduce the characteristic QCD $\ln^2 M_Q$ behaviour when M_Q is allowed to become arbitrarily large; the CQM contribution decouples in the large M_Q -limit. Incidentally, the contribution from free u and d quarks with a small mass $m \ll \mu$ goes as $\sim (\alpha/\pi)^3 N_c \frac{7}{27} \times \frac{2}{3} \pi^2 \ln \frac{m_\mu}{m}$, which is also incompatible with the QCD result in equation (142). Therefore, arguments based on the fact that the CQM (and/or pQCD) gives a *positive* contribution are certainly ‘simple’, but also *incorrect*. Notice however, that,




Figure 56. Weak interactions at the one-loop level.

contrary to the naive CQM, the constituent chiral quark model of Georgi and Manohar [151] (see also [152]) does indeed reproduce the correct $\ln^2 M_Q$ behaviour in the $M_Q \rightarrow \infty$ limit. This is because, in these models, the Goldstone particles couple with the constituent quarks in a way which respects chiral symmetry, and the pion pole diagram appears then explicitly. The same happens in the extended version of the Nambu–Jona-Lasinio model [153]. These models, however, suffer from other diseases [154], and therefore they are not fully reliable to compute the hadronic light-by-light scattering contribution.

Hopefully, hadronic models of the light-by-light scattering contribution which respect the QCD constraint in equation (142) will be progressively improved, so as to incorporate further and further QCD features; in particular, short-distance constraints, following the lines discussed in [114, 115, 155]. An interesting contribution in this direction has been reported in [156]. Unfortunately, as recently discussed in [157] their numerical evaluation is incomplete and model dependent with largely underestimated errors. We believe that, at present, one can only claim to know the hadronic light-by-light scattering contribution with a *cautious* error, which takes into account the uncertainties from the integration regions which remain model dependent. While awaiting further improvement, the educated value one can quote at present, based on the combined work of [150] and [148, 149] (appropriately corrected) as well as [156], is

$$a_\mu^{(6)}(\text{H})_{1 \times 1} = (11 \pm 4) \times 10^{-10}. \quad (148)$$

7.4. Electroweak contributions

The leading contribution to the anomalous magnetic moment of the muon from the electroweak Lagrangian of the Standard Model, originates at the one-loop level. The relevant Feynman diagrams, which for simplicity we draw in the unitary gauge, are shown in figure 56, where we also indicate the size of their respective contributions.

The analytic evaluation of the overall contribution gives the result [22]

$$\begin{aligned} a_\mu^w &= \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \left[\frac{5}{3} + \frac{1}{3} (1 - 4 \sin^2 \theta_W)^2 + \mathcal{O} \left(\frac{m_\mu^2}{M_Z^2} \log \frac{M_Z^2}{m_\mu^2} \right) + \frac{m_\mu^2}{M_H^2} \int_0^1 dx \frac{2x^2(2-x)}{1-x + \frac{m_\mu^2}{M_H^2} x^2} \right] \\ &= 19.48 \times 10^{-10}, \end{aligned} \quad (149)$$

where the weak mixing angle is defined by $\sin^2 \theta_W = 1 - M_W^2/M_Z^2 = 0.223$ (for a Higgs mass $M_H = 150 \text{ GeV}$), and where $G_F = 1.16639(1) \times 10^{-5}$ is the Fermi constant. Notice that the contribution from the Higgs boson, shown in parametric form in the second line, is of $\mathcal{O} \left(\frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{4\pi^2} \frac{m_\mu^2}{M_H^2} \ln \frac{M_H^2}{m_\mu^2} \right)$, which is already very small for the present lower bound of M_H . Closed analytic expressions of the electroweak one-loop contribution can be found in [158] and references therein.

Let us recall that the present world average experimental error in the determination of the muon anomaly is $\Delta a_\mu|_{\text{Exp.}} = \pm 6.3 \times 10^{-10}$ [26], and, hoping for a continuation of the BNL experiment, it is expected to be further reduced. The *a priori* possibility that the two-loop electroweak corrections may bring in enhancement factors due to large logarithmic factors, like $\ln(M_Z^2/m_\mu^2) \simeq 13.5$, has motivated a thorough theoretical effort for their evaluation, which has been quite a remarkable achievement, and which we discuss next.

It is convenient to separate the two-loop electroweak contributions into two sets of the Feynman graphs: those containing closed fermion loops, which we denote by $a_\mu^{\text{EW}(2)}(\text{ferm})$, and the others which we denote by $a_\mu^{\text{EW}(2)}(\text{bos})$. With this notation, the electroweak contribution to the muon anomalous magnetic moment becomes

$$a_\mu^{\text{EW}} = a_\mu^{\text{W}} + a_\mu^{\text{EW}(2)}(\text{bos}) + a_\mu^{\text{EW}(2)}(\text{ferm}), \quad (150)$$

with a_μ^{W} the one-loop contribution in equation (149).

Bosonic contributions to $a_\mu^{\text{EW}(2)}(\text{bos})$. The leading logarithmic terms of the two-loop electroweak bosonic corrections have been extracted using asymptotic-expansion techniques. In fact, these contributions have now been evaluated analytically, in a systematic expansion in powers of $\sin^2 \theta_W$, up to $\mathcal{O}[(\sin^2 \theta_W)^3]$, where $\log \frac{M_W^2}{m_\mu^2}$ terms, $\log \frac{M_H^2}{M_W^2}$ terms, $\frac{M_W^2}{M_H^2} \log \frac{M_H^2}{M_W^2}$ terms, $\frac{M_W^2}{M_H^2}$ terms and constant terms are kept [159]. Using $\sin^2 \theta_W = 0.224$ and $M_H = 250$ GeV, the authors in [159] find

$$a_\mu^{\text{EW}(2)}(\text{bos}) = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \frac{\alpha}{\pi} \times \left[-5.96 \log \frac{M_W^2}{m_\mu^2} + 0.19 \right] = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \left(\frac{\alpha}{\pi} \right) \times (-79.3). \quad (151)$$

Fermionic contributions to $a_\mu^{\text{EW}(2)}(\text{ferm})$. The discussion of the two-loop electroweak fermionic corrections is more delicate. Because of the $U(1)$ anomaly cancellation between lepton loops and quark loops in the electroweak theory, one cannot separate hadronic from leptonic effects any longer in diagrams like the ones shown in figure 57, where a triangle with two-vector currents and an axial-vector current appears (the so-called VVA-triangle). Individually, the lepton-loop and quark-loop contributions are each gauge-dependent; depending on the gauge choice, they can even lead to UV-divergent contributions. Only the sum of contributions within each family of leptons and quarks is free from these ambiguities. As first discussed in [160, 161], it is this anomaly cancellation which eliminates some of the large logarithms that were incorrectly kept in a previous calculation in [162]. It is therefore appropriate to separate the two-loop electroweak fermionic corrections into two classes. One is the class arising from the Feynman diagrams like the ones in figure 57, where a subgraph with a VVA-triangle of leptons and quarks appears, including the graphs where the Z lines are replaced by Φ^0 lines, if the calculation is done in the ξ_Z -gauge. We denote this class by $a_\mu^{\text{EW}(2)}(l, q)$. The other class is defined by the rest of the diagrams, where quark loops and lepton loops can be treated separately, which we call $a_\mu^{\text{EW}(2)}(\text{ferm-rest})$ i.e.

$$a_\mu^{\text{EW}(2)}(\text{ferm}) = a_\mu^{\text{EW}(2)}(l, q) + a_\mu^{\text{EW}(2)}(\text{ferm-rest}).$$

The contribution from $a_\mu^{\text{EW}(2)}(\text{ferm-rest})$ brings in m_t^2/M_W^2 factors. It has been estimated, to a very good approximation, in [161] with the result,

$$a_\mu^{\text{EW}(2)}(\text{ferm-rest}) = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \frac{\alpha}{\pi} \times (-21 \pm 4), \quad (152)$$




Figure 57. Two-loop electroweak diagrams generated by the $\gamma\gamma Z$ -triangle. There are similar diagrams corresponding to the μ ; c , s and τ ; t , b generations.

where the error here is the one induced by diagrams with Higgs propagators with an allowed Higgs mass in the range $114 \text{ GeV} < M_H < 250 \text{ GeV}$.

With regard to the contributions to $a_\mu^{\text{EW}(2)}(l, q)$, it is convenient to treat the three generations separately. The contribution from the third generation can be calculated in a straightforward way, because all the fermion masses in the triangle loop are large with respect to the muon mass, with the result [160, 161]

$$a_\mu^{\text{EW}(2)}(\tau, t, b) = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \frac{\alpha}{\pi} \times \left[-3 \log \frac{M_Z^2}{m_\tau^2} - \log \frac{M_Z^2}{m_b^2} - \frac{8}{3} \log \frac{m_t^2}{M_Z^2} + \frac{8}{3} + \mathcal{O}\left(\frac{M_Z^2}{m_t^2} \log \frac{m_t^2}{M_Z^2}\right) \right] = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \frac{\alpha}{\pi} \times (-30.6). \quad (153)$$

The leading terms can be easily obtained using effective field-theory techniques [160]; but terms of $\mathcal{O}(M_Z^2/m_t^2 \log m_t^2/M_Z^2)$ and $\mathcal{O}(M_Z^2/m_t^2)$ have also been calculated in [161]. There are, in principle, QCD perturbative corrections to this estimate which have not been calculated, but the result in equation (153) is good enough for the accuracy required at present.

As emphasized in [160], an appropriate QCD calculation when the quark in the loop of figure 1 is a *light quark* should take into account the dominant effects of spontaneous chiral symmetry breaking. Since this involves the u and d quarks, as well as the second-generation s quark, it is convenient to lump together the contributions from the first and second generations. An evaluation of these contributions, which incorporates the QCD long-distance chiral realization [160, 163] as well as perturbative [142] and non-perturbative [163] short-distance constraints, gives the result

$$a_\mu^{\text{EW}(2)}(e, \mu, u, d, s, c) = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \frac{\alpha}{\pi} \times (-24.6 \pm 1.8). \quad (154)$$

From the theoretical point of view, this calculation has revealed surprising properties concerning the *non-anomalous* component of the VVA-triangle [164], with a new set of *non-renormalization theorems* in perturbation theory [164, 165] (see also [166–168]). The physical meaning of these perturbation theory results when considered in the full QCD vacuum is, however, still unclear [165].

Putting together the numerical results in equations (151)–(153) with the more recent result in equation (154), one finally obtains the value

$$a_\mu^{\text{EW}} = \frac{G_F}{\sqrt{2}} \frac{m_\mu^2}{8\pi^2} \left\{ \frac{5}{3} + \frac{1}{3} (1 - 4 \sin^2 \theta_W)^2 - \left(\frac{\alpha}{\pi}\right) [155.5(4)(1.8)] \right\}, \quad (155)$$

$$= 15.4(0.2)(0.1) \times 10^{-10}, \quad (156)$$

where the first error is essentially due to the Higgs mass uncertainty, while the second comes from higher-order hadronic effects in the VVA loop evaluation. The overall result shows indeed that the two-loop correction represents a sizable reduction of the one-loop result by an amount of 21%. This result has prompted the evaluation of the electroweak leading-log three-loop effects of $\mathcal{O}[G_F/\sqrt{2}m_\mu^2/8\pi^2(\alpha/\pi)^2 \ln M_Z/m_\mu]$, using renormalization group arguments similar to the ones discussed in section 7.1.2. The results [142, 169] show that those effects are negligible [$\sim \mathcal{O}(10^{-12})$] for the accuracy needed at present.

7.5. Summary of the Standard Model contributions to the muon anomaly

The present status in the evaluation of the anomalous magnetic moment of the muon in the Standard Model can be summarized as follows.

- *Leptonic QED contributions*

$$a_\mu(\text{QED}) = (116\,584\,718.09 \pm 0.14_{5\text{loops}} \pm 0.08_\alpha \pm 0.04_{\text{masses}}) \times 10^{-11}. \quad (157)$$

This is the number one obtains by adding to the final result at the bottom of table 8 (using recent [28] value of α) to the estimate of the five-loop contribution in equation (119).

- *Hadronic contributions*

- *Hadronic vacuum polarization with EM-data*

From the two most recent determinations in table 9, which take into account the new data on e^+e^- annihilation into hadrons, we get

$$a_\mu[\text{HVP}(06)] = (6901 \pm 42_{\text{exp}} \pm 19_{\text{rad}} \pm 7_{\text{QCD}}) \times 10^{-11}, \quad (158)$$

where we have averaged the central values, but kept the largest uncertainties in the two determinations. Notice that this result is consistent with the average of the central values of the earlier determinations from e^+e^- data in the second part of table 9, which when keeping the largest error from experiment and from radiative corrections give $a_\mu[\text{HVP}] = (6945 \pm 62_{\text{exp}} \pm 36_{\text{rad}}) \times 10^{-11}$.

- *Higher-order hadronic vacuum polarization with EM-data*

$$a_\mu[\text{HVP h.o.}] = (-97.9 \pm 0.9_{\text{exp}} \pm 0.3_{\text{rad}}) \times 10^{-11}. \quad (159)$$

- *Hadronic light-by-light scattering*

$$a_\mu[\text{HLLS}] = (110 \pm 40) \times 10^{-11}. \quad (160)$$

- *Electroweak contributions*

$$a_\mu[\text{EW}] = (154 \pm 2 \pm 1) \times 10^{-11}. \quad (161)$$

The sum of these contributions, using the HVP06 result in equation (158) and adding experimental and theoretical errors in quadrature, gives then a total

$$a_\mu^{\text{SM}(06)} = 116\,591\,785 (61) \times 10^{-11}. \quad (162)$$

These determinations are to be compared with the experimental world average in equation (21)

$$a_\mu^{\text{exp}} = 116\,592\,080 (63) \times 10^{-11}. \quad (163)$$

Therefore, we conclude that, with the input for the Standard-Model contributions discussed above, one finds at present a 3.4σ discrepancy.

Table 10. Systematic and statistical errors in ppm for each of the e821 running periods.

Systematic uncertainty	1998	1999	2000	2001	E969 Goal
Magnetic field ω_p (ppm)	0.5	0.4	0.24	0.17	0.1
Anomalous precession ω_a (ppm)	0.8	0.3	0.31	0.21	0.1
Statistical uncertainty (ppm)	4.9	1.3	0.62	0.66	0.20
Total uncertainty (ppm)	5.0	1.3	0.73	0.72	0.25

8. Future experimental possibilities

The discrepancy between the experiment and the theory gives strong motivation to improve the experimental number. An upgraded experiment, E969¹³, with goals of $\sigma_{\text{sys}} = 0.14$ ppm and $\sigma_{\text{stat}} = 0.20$ ppm received scientific approval from the Brookhaven Laboratory. Despite the 3.4σ difference between theory and experiment, the new experiment remains unfunded. Such an experiment would stimulate additional theoretical work which would further decrease the theoretical error. The ultimate experimental sensitivity of a future experiment depends both on the available beam intensity and on the ability to further reduce the systematic errors.

The systematic and statistical errors by year for E821, along with the goals for E969, are given in table 10. There is a clear improvement in the systematic errors with time, and the improvement required for E969 is modest and should be readily achievable with planned upgrades in the detector system, electronics and the field mapping system. In order to achieve the required number of muons for E969, a plan to improve the beamline in order to increase the stored muon flux by a factor of five has been proposed.

The E821 experimental approach, which combines the ‘magic γ ’, electrostatic focusing, and a fast muon kicker has the potential to push the experimental error below the 0.1 ppm level, perhaps as far as about 0.05 ppm before too many technical barriers are reached. To improve the statistical error on ω_a the muon flux would have to be increased substantially, and presently such a source does not exist. As the muon flux is increased, systematic errors from pulse pile-up and detector instabilities generated by high rates become more severe. To improve the ω_p measurement, additional work would have to be done on the absolute calibration, which at present depends on the measurement of Phillips *et al* [65] to relate the magnetic moment of a proton in a spherical water sample to that of the free proton. An alternate calibration based on a ^3He probe has been suggested [170]. To go beyond 0.1 ppm, it would also be necessary to further refine the measurement of λ_+ with a new muonium hyperfine splitting measurement [170].

Progressing beyond 0.05 ppm might require a new technique. For example, increasing $\gamma\tau$ by going to a higher muon momentum and increasing f_a by increasing the magnetic field would reduce the statistical error without requiring more muons (see equation (11)). However, a higher momentum would require abandoning the advantages of the magic gamma. Increasing the magnetic field would require the elimination of the iron yoke and pole tips in favour of

¹³ A ($g - 2$) $_{\mu}$ Experiment to ± 0.2 ppm Precision. The proposal can be found at <http://g2pc1.bu.edu/~roberts/>. At this time the collaboration consists of Carey R M, Logashenko I, Lynch K R, Miller J P, Roberts B L, Bunce G, Meng W, Morse W M, Semertzidis Y K, Grigoriev D, Khazin B I, Redin S I, Shatunov Yuri M, Solodov E, Orlov Y, Crnkovic J, Debevec P, Hertzog D W, Kammel P, McNabb R, Giovanetti K L, Yamamoto A, Jungmann K P, Onderwater C J G, Gorringer T P, Korsch W, Gray F E, Lauss B, Sichtermann E P, Cushman P, Qian T, Arimoto Y, Aoki M, Kuno Y, Sato A, Yamada K, Dhawan S and Farley F J M.

a magnetic field driven entirely by superconducting coils. A proposal to use sector focusing dipoles has been made [171], but the need to know $\int \vec{B} \cdot d\vec{\ell}$ for the muons to a precision of 10 ppb with magnetic gradients present is challenging, to say the least.

9. Summary and conclusions

Measurement of the muon's anomalous magnetic moment has a history stretching back to the late 1950s. In the paper from Nevis Laboratory that presented evidence for parity violation in muon decay [30], it was reported that the data were consistent with $g = 2$. In a subsequent paper from Garwin *et al* [33], it was demonstrated that the muon anomaly was consistent with the QED calculation of Schwinger of $\sim \alpha/2\pi$, and thus provided the first evidence, confirmed by all existing subsequent data, that the muon behaves like a heavy electron.

The result from E821 at the Brookhaven National Laboratory is sensitive to extremely refined effects predicted by the Standard Model, such as hadronic light-by-light scattering and electroweak contributions. If these two effects were removed from the theoretical prediction obtained in equation (162), the discrepancy with the experimental result would be 7.1σ . If only the electroweak contribution is removed, the discrepancy is 5.1σ , clearly showing the sensitivity of the E821 experiment to the electroweak contribution. The largest uncertainties in the theoretical prediction of a_μ come from the systematic errors in the experimental data from $e^+ + e^-$ annihilation to hadrons (in spite of the fact that the accuracy in the present determination ($<1\%$) is quite remarkable); and from the theoretical error in the hadronic light-by-light contribution.

There is progress to be expected as additional e^+e^- data become available. The good news is that the $\pi\pi$ data from two independent experiments at the Novosibirsk collider, CMD-2 and SND, are in agreement. There seems to be some disagreement between the shape of the pion form factor obtained by the KLOE experiment at Frascati and that from the Novosibirsk experiments. Further experimental input in the near future is expected on $e^+e^- \rightarrow \pi^+\pi^-$ from a new KLOE analysis with a tagged photon in the detector, and from BaBar using initial-state radiation. At present the data from Novosibirsk dominate the determination of a_μ [HVP] in the ρ -meson region, making an independent confirmation of the $\pi\pi$ cross-section in this energy region clearly desirable.

On the other hand, the puzzle centred on hadronic τ -decay data versus e^+e^- data remains in spite of efforts to quantify the CVC corrections. Problems exist in comparisons between the two that are quite independent of the contribution to a_μ . For example the failure to predict the measured τ branching fractions from the e^+e^- data and the conserved vector current hypotheses, at present invalidate attempts to use these τ data to obtain the hadronic contribution [129]. Work on these issues in τ physics is continuing. We conclude that until the issues with τ -decay data are resolved, the e^+e^- data are preferred to evaluate the lowest-order hadronic contribution to a_μ .

One should stress the impressive achievement reached in the calculation of the QED contribution, where the dominant error at present comes from the five-loop contribution. Thanks to the dedicated efforts of Kinoshita and collaborators, this uncertainty is likely to go down significantly in the near future. Great progress has also been made in understanding the electroweak contribution, which is now fully calculated at the two-loop level with an accuracy of 1.5%, an uncertainty dominated by the unknown Higgs mass. The main effort should now be concentrated on a more accurate determination of the hadronic light-by-light contribution. A goal of 15% accuracy in its determination seems possible with a dedicated effort from theorists.

An improvement in the experimental value of the muon anomaly by at least a factor of two appears to be quite feasible, and is the goal of the proposed E969 experiment at the Brookhaven National Laboratory. Given the current 3.4-standard-deviation difference between theory and experiment, and the prospects for the uncertainty in the theoretical value to decrease over the next few years, such a measurement appears timely and highly desirable.

Historically the muon anomaly has provided an important constraint on speculative theories of physics beyond the Standard Model [23–25]. As particle physics moves into the large hadron collider era in the next few years at CERN, the muon anomaly will provide an independent constraint on the interpretation of the discoveries made at the LHC.

Acknowledgments

The Brookhaven based experimental work reported in this paper was carried out over a twenty-year period by the E821 collaboration, and the authors gratefully acknowledge the many contributions that each member [26, 47, 48–51] made to its success. In this review of E821 the authors have tried to report accurately the work that led to the final precision of 0.54 ppm on the muon anomaly, and they (JPM and BLR) accept full responsibility for any shortcomings of this exposition. Earlier reviews on aspects of muon ($g - 2$) have been written by some of their colleagues [172]. They thank their colleagues on E821 and E969 for many helpful conversations on various aspects of the $g-2$ experiments. The authors are especially grateful to R M Carey and K Ellis for their significant editorial suggestions. They also thank F J M Farley, D W Hertzog and K R Lynch for their comments on this manuscript. BLR acknowledges useful comments by Andrzej Czarnecki on non-standard model contributions. The authors acknowledge K J Jungmann and R Prigl for their helpful comments on the sections about the magnetic field. The preparation of this manuscript was supported in part by the US National Science Foundation. E821 was supported in part by the US Department of Energy, the US National Science Foundation, the German Bundesminister für Bildung und Forschung, the Russian Ministry of Science, the US–Japan Agreement in High Energy Physics and the NATO Office of Scientific Research. The work of E de R has been supported in part by TMR, EC-Contract No. HPRN-CT-2002-0031 (EURIDICE) and by the European Community’s Marie Curie Research Training Network under contract no. MRTN-CT-2006-035482 (FLAVIANet).

References

- [1] Stern O 1921 *Z. Phys.* **7** 249
- [2] Gerlach W and Stern O 1922 *Z. Phys.* **8** 110
Gerlach W and Stern O 1924 *Z. Phys.* **9** 349
Gerlach W and Stern O 1924 *Z. Phys.* **9** 353
- [3] Gerlach W and Stern O 1924 *Ann. Phys.* **74** 673
- [4] Phipps T E and Taylor J B 1927 *Phys. Rev.* **29** 309
- [5] Compton A K 1921 *J. Franklin Inst.* **192** 145
- [6] Uhlenbeck G E and Goudsmit S 1925 *Naturwissenschaften* **47** 953
Uhlenbeck G E and Goudsmit S 1926 *Nature* **117** 264
- [7] Dirac P A M 1928 *Proc. R. Soc. Lond. A* **117** 610
Dirac P A M 1928 *Proc. R. Soc. Lond. A* **118** 351
Dirac P A M 1958 *The Principles of Quantum Mechanics* 4th edn (London: Oxford University Press)
- [8] Pais A 1998 *Paul Dirac: The Man and His Work* ed P Goddard (New York: Cambridge University Press)
- [9] Bjorken J D and Drell S D 1964 *Relativistic Quantum Mechanics* (New York: McGraw-Hill) p 13
- [10] Frisch R and Stern O 1933 *Z. Phys.* **85** 4
Estermann I and Stern O 1933 *Z. Phys.* **85** 17
- [11] Alvarez L W and Bloch F 1940 *Phys. Rev.* **57** 111

- [12] Bethe H A and Salpeter E E 1957 *Quantum Mechanics of One- and Two-Electron Atoms* (Berlin: Springer) p 51
- [13] Nafe J E, Nelson E B and Rabi I I 1947 *Phys. Rev.* **71** 914
- [14] Nagel D E, Julian R S and Zacharias J R 1947 *Phys. Rev.* **72** 971
- [15] Kusch P and Foley H M 1947 *Phys. Rev.* **72** 1256
- [16] Schwinger J 1948 *Phys. Rev.* **73** 416L
Schwinger J 1949 *Phys. Rev.* **76** 790
- [17] Breit G 1947 *Phys. Rev.* **72** 984
- [18] Kusch P and Foley H M 1948 *Phys. Rev.* **73** 250
- [19] For a nice historical overview of the lepton g -factors see Rich A and Wesley J 1972 *Rev. Mod. Phys.* **44** 250
- [20] Odom B, Hanneke D, D'Urso B and Gabrielse G 2006 *Phys. Rev. Lett.* **97** 030801
- [21] Van Dyck R S *et al* 1987 *Phys. Rev. Lett.* **59** 26
Van Dyck R S *et al* 1990 *Quantum Electrodynamics (Directions in High Energy Physics vol 7)* ed T Kinoshita (Singapore: World Scientific) p 322
- [22] Bardeen W A, Gastmans R and Lautrup B E 1972 *Nucl. Phys. B* **46** 315
Altarelli G, Cabbibo N and Maiani L 1972 *Phys. Lett. B* **40** 415
Jackiw R and Weinberg S 1972 *Phys. Rev. D* **5** 2473
Bars I and Yoshimura M 1972 *Phys. Rev. D* **6** 374
Fujikawa M, Lee B W and Sanda A I 1972 *Phys. Rev. D* **6** 2923
- [23] Czarnecki A and Marciano W J 2001 *Phys. Rev. D* **64** 013014
- [24] Martin S P and Wells J D 2003 *Phys. Rev. D* **67** 015002
- [25] Stöckinger D 2006 *Preprint hep-ph/0609168*
Stöckinger D 2007 *J. Phys. G: Nucl. Part. Phys.* **34** R45
- [26] Bennett G *et al* (Muon ($g - 2$) Collaboration) 2006 *Phys. Rev. D* **73** 072003
- [27] Cladé P *et al* 2006 *Phys. Rev. Lett.* **96** 033001 and references therein
- [28] Gabrielse G, Hanneke D, Kinoshita T, Nio M and Odom B 2006 *Phys. Rev. Lett.* **97** 030802 and references therein
- [29] Wu C S, Ambler E, Hayward R W, Hoppes D D and Hudson R P 1957 *Phys. Rev.* **105** 1413
- [30] Garwin R L, Lederman L M and Weinrich M 1957 *Phys. Rev.* **105** 1415
- [31] Friedman J I and Telegdi V L 1957 *Phys. Rev.* **105** 1681
- [32] Konopinski E J 1959 *Annu. Rev. Nucl. Sci.* **9** 99
- [33] Garwin R L, Hutchinson D P, Penman S and Shapiro G 1960 *Phys. Rev.* **118** 271
- [34] Charpak G *et al* 1961 *Phys. Rev. Lett.* **6** 128
Charpak G *et al* 1961 *Nuovo Cimento* **22** 1043
Charpak G *et al* 1962 *Phys. Lett.* **1** 16
Charpak G *et al* 1965 *Nuovo Cimento* **37** 1241
- [35] Charpak G *et al* 1962 *Phys. Lett.* **1** 16
- [36] Bailey J *et al* 1968 *Phys. Lett. B* **28** 287
Additional details can be found in Bailey J *et al* 1972 *Nuovo Cimento A* **9** 369 and references therein
- [37] Bailey J *et al* 1975 *Phys. Lett.* **55** 420
- [38] Bailey J *et al* 1979 *Nucl. Phys. B* **150** 1
- [39] Liu W *et al* 1999 *Phys. Rev. Lett.* **82** 711
- [40] Yao W-M *et al* (Particle Data Group) 2006 *J. Phys. G: Nucl. Part. Phys.* **33** 1
- [41] Fliegel W and Krienen F 1973 *Nucl. Instrum. Methods* **143** 549
- [42] Thomas L H 1927 *Phil. Mag.* **3** 1
Bargmann V, Michel L and Telegdi V L 1959 *Phys. Rev. Lett.* **2** 435
- [43] Wiedemann H 1993 *Particle Accelerator Physics vol 1* (Berlin: Springer) p 54
- [44] Edwards D A and Syphers M J 1993 *An Introduction to the Physics of High Energy Accelerators* (New York: Wiley) p 75
- [45] Combley F and Picasso E 1974 *Phys. Rep.* **14** 1
- [46] Henry G R, Schrank G and Swanson R A 1969 *Nuovo Cimento A* **63** 995
- [47] Carey R M *et al* 1999 *Phys. Rev. Lett.* **82** 1632
- [48] Brown H N *et al* (Muon ($g - 2$) Collaboration) 2000 *Phys. Rev. D* **62** 091101
- [49] Brown H N *et al* (Muon ($g - 2$) Collaboration) 2001 *Phys. Rev. Lett.* **86** 2227
- [50] Bennett G W *et al* (Muon ($g - 2$) Collaboration) 2002 *Phys. Rev. Lett.* **89** 101804
- [51] Bennett G W *et al* (Muon ($g - 2$) Collaboration) 2004 *Phys. Rev. Lett.* **92** 161802
- [52] Orlov Y *et al* 2002 *Nucl. Instrum. Methods A* **482** 767
- [53] We follow the approach given by Paley J M 2004 *PhD Dissertation* Boston University

- [54] Farley F J M and Picasso E 1990 *Quantum Electrodynamics (Advanced Series on Directions in High Energy Physics vol 7)* ed T Kinoshita (Singapore: World Scientific) p 479
- [55] Farley F J M 1972 *Phys. Lett. B* **42** 66
Field J H and Fiorentini G 1974 *Nuovo Cimento A* **21** 297
- [56] Farley F J M, private communication
- [57] Drumm H *et al* 1979 *Nucl. Instrum. Methods* **158** 347
- [58] Krienen F, Loomba D and Meng W 1989 *Nucl. Instrum. Methods A* **283** 5
- [59] Yamamoto A *et al* 2002 *Nucl. Instrum. Methods Phys. Res. A* **491** 23–40
- [60] Danby G T *et al* 2001 *Nucl. Instrum. Methods Phys. Res. A* **457** 151–74
- [61] Sedyk S A *et al* 2000 *Nucl. Instrum. Methods A* **455** 346
- [62] Ouyang J *et al* 1996 *Nucl. Instrum. Methods* **374** 215
- [63] Bailey J *et al* 1978 *J. Phys. G: Nucl. Part. Phys.* **4** 345
- [64] See Abragam A 1961 *Principles of Nuclear Magnetism* (Oxford: Oxford University Press) pp 173–8
- [65] Phillips W D *et al* 1979 *Metrologia* **13** 179
- [66] Mohr P J and Taylor B H 2005 *Rev. Mod. Phys.* **77** 1
- [67] Winkler P F, Kleppner D, Myint T and Walther F G 1972 *Phys. Rev. A* **5** 83
- [68] Lamb W E Jr 1941 *Phys. Rev.* **60** 817
- [69] Grotch H and Hegstrom R A 1971 *Phys. Rev. A* **4** 59
- [70] Petley B W *et al* 1984 *Metrologia* **20** 81
- [71] Bennett G W *et al* (Muon ($g - 2$) Collaboration) Statistical equations and methods applied to the precision muon ($g-2$) experiment at BNL *Nucl. Instrum. Methods* submitted
- [72] Efstathiadis E *et al* 2002 *Nucl. Instrum. Methods Phys. Res. A* **496** 8–25
- [73] Semertzidis Y K 2003 *Nucl. Instrum. Methods Phys. Res. A* **503** 458–84
- [74] Prigl R *et al* 1996 *Nucl. Instrum. Methods Phys. Res. A* **374** 118
- [75] Fei X, Hughes V and Prigl R 1997 *Nucl. Instrum. Methods Phys. Res. A* **394** 349
- [76] *LORAN-C User's Handbook* OMDTPUB P16562.6, available at <http://www.navcen.uscg.gov/Loran/handbook/h-book.htm>
- [77] Angelopoulos A *et al* (CPLEAR Collaboration) 1998 *Phys. Lett. B* **444** 43
- [78] Farley F J M *et al* 2004 *Phys. Rev. Lett.* **92** 052001
- [79] Feng J L, Matchev K T and Shadmi Y 2003 *Phys. Lett. B* **555** 89
Feng J L, Matchev K T and Shadmi Y 2001 *Nucl. Phys. B* **613** 366
- [80] Petermann A 1957 *Phys. Rev.* **105** 1931
- [81] Sommerfield C M 1957 *Phys. Rev.* **107** 328
- [82] Laporta S and Remiddi E 1996 *Phys. Lett. B* **379** 283
- [83] Broadhurst D J and Kreimer D 1997 *Phys. Lett. B* **393** 403
- [84] Connes A and Kreimer D 1998 *Commun. Math. Phys.* **199** 203
- [85] Lewin L 1958 *Dilogarithms and Associated Functions* (London: Macdonald)
- [86] Kölbig K S, Mignaco J A and Remiddi E 1970 *BIT* **10** 38
- [87] Kinoshita T 1990 *Quantum Electrodynamics* ed T Kinoshita (Singapore: World Scientific)
- [88] Kinoshita T and Nio M 2006 *Phys. Rev. D* **73** 013003
- [89] Aoyama T, Hayakawa M, Kinoshita T and Nio M 2006 *Nucl. Phys. B* **740** 138
- [90] Suura H and Wichmann E 1957 *Phys. Rev.* **105** 1930
- [91] Lautrup B E and de Rafael E 1974 *Nucl. Phys. B* **70** 317
- [92] de Rafael E and Rosner J L 1974 *Ann. Phys.* **82** 369
- [93] Friot S, Greynat D and de Rafael E 2005 *Phys. Lett. B* **628** 73
- [94] Elend H H 1966 *Phys. Lett.* **20** 682
Elend H 1966 *Phys. Lett.* **21** 720 (erratum)
- [95] Passera M 2005 *J. Phys. G: Nucl. Part. Phys.* **31** R75
- [96] Mohr P J and Taylor B N 2005 CODATA recommended values of the fundamental physical constants: 2005 *Rev. Mod. Phys.* **75** 1
- [97] Eidelman S *et al* (Particle Data Group Collaboration) 2004 *Phys. Lett. B* **592** 1
- [98] Kinoshita T 1967 *Nuovo Cimento B* **51** 140
- [99] Lautrup B E and de Rafael E 1968 *Phys. Rev.* **174** 1835
- [100] Lautrup B E, Peterman A and de Rafael E 1972 *Phys. Rep. C* **3** 193–260
- [101] Laporta S 1993 *Nuovo Cimento A* **106** 675
- [102] Baikov P A and Broadhurst D J 1995 *Proc. Pisa-AIHENP Conf.* ed B Derby and D Perret-Gallix (Singapore: World Scientific) pp 016772 *Preprint hep-ph/9504398*
Broadhurst D J, Kataev A L and Tarasov O V 1993 *Phys. Lett. B* **298** 445

- [103] Laporta S and Remiddi E 1993 *Phys. Lett. B* **301** 440
- [104] Czarnecki A and Skrzypek M 1977 *Phys. Lett. B* **72** 114
- [105] Smirnov V 2002 *Applied Asymptotic Expansions in Momenta and Masses (Springer Tracts in Modern Physics)* (Heidelberg: Springer)
- [106] Kinoshita T and Nio M 2005 *Phys. Rev. D* **70** 113001
- [107] Kinoshita T 2005 *Nucl. Phys. (Proc. Suppl.) B* **144** 206
Kinoshita T and Nio M 2006 *Phys. Rev. D* **73** 053007
- [108] Aldins J, Brodsky S, Dufner A and Kinoshita T 1969 *Phys. Rev. Lett.* **23** 441
Aldins J, Brodsky S, Dufner A and Kinoshita T 1970 *Phys. Rev. D* **1** 2378
- [109] Laporta S and Remiddi E 1993 *Phys. Lett. B* **301** 440
- [110] Bouchiat C and Michel L 1961 *J. Phys. Radium* **22** 121
- [111] Brodsky S and de Rafael E 1968 *Phys. Rev.* **168** 1620
- [112] Bell J S and de Rafael E 1969 *Nucl. Phys.* **11** 611
- [113] Lautrup B E and de Rafael E 1971 *Nuovo Cimento A* **1** 238
- [114] de Rafael E 2003 Analytic approaches to kaon physics *Nucl. Phys. B (Proc. Suppl.)* **119** 71
- [115] Peris S 2002 Electroweak matrix elements at large- N_c : matching quarks to mesons *Preprint hep-ph/0204181*
- [116] Friot S, Greynat D and de Rafael E 2004 *J. High Energy Phys.* **JHEP10(2004)043**
- [117] Blum T 2003 *Phys. Rev. Lett.* **91** 052001
Aubin C and Blum T 2005 *Preprint hep-lat/0509064* (PoS LAT2005:089,2005)
Hayakawa M *et al* 2005 *Preprint hep-lat/0509016* (PoS LAT2005:353,2005)
- [118] Gourdin M and de Rafael E 1969 *Nucl. Phys. B* **10** 667
- [119] Calmet J, Narison S, Perrottet M and de Rafael E 1976 *Phys. Lett. B* **61** 283
- [120] Akhmetshin R R *et al* (CMD-2 Collaboration) 2004 *Phys. Lett. B* **578** 285
- [121] Alemany R, Davier M and Höcker M 1998 *Eur. Phys. J. C* **2** 123
- [122] Davier M, Eidelman S, Höcker A and Zhang Z 2003 *Eur. Phys. J. C* **31** 503
- [123] Jegerlehner F 2004 *Nucl. Phys. (Proc. Suppl.) B* **126** 325
Jegerlehner F 2004 *Nucl. Phys. (Proc. Suppl.) B* **131** 213
- [124] de Trocóniz J F and Ynduráin F J 2005 *Phys. Rev. D* **71** 073008 (*Preprint hep-ph/0402285*)
- [125] Hagiwara K, Martin A D, Nomura D and Teubner T 2004 *Phys. Rev. D* **69** 093003
- [126] Akhmetshin R R *et al* 2006 *JETP Lett.* **84** 413–7 (*Preprint hep-ex/0610016*)
Akhmetshin R R *et al* 2006 *Zh. Eksp. Teor. Fiz.* **84** 491–5
- [127] Akhmetshin R R *et al* 2007 *Phys. Lett. B* **648** 28 (*Preprint hep-ex/0610021*)
- [128] Achasov M N *et al* 2006 *JETP* **103** 380–4 (*Preprint hep-ex/0605013*)
Achasov M N *et al* 2006 *Zh. Eksp. Teor. Fiz.* **130** 437–41
- [129] Davier M 2007 *Preprint hep-ph/0701163v2*
- [130] Hagiwara K, Martin A D, Nomura D and Teubner T 2006 *Phys. Lett. B* at press *Preprint hep-ph/0611102*
- [131] Sirlin A 1982 *Nucl. Phys. B* **196** 83
- [132] Marciano W J and Sirlin A 1988 *Phys. Rev. Lett.* **61** 1815
- [133] Braaten E, Narison S and Pich A 1992 *Nucl. Phys. B* **373** 581
- [134] Ghozzi S and Jegerlehner F 2003 *Preprint hep-ph/0310181*
- [135] Cirigliano V, Ecker G and Neufeld H 2001 *Phys. Lett. B* **513** 361
- [136] Cirigliano V, Ecker G and Neufeld H 2002 *J. High Energy Phys.* **JHEP08(2002)002**
- [137] Aloisio A *et al* (KLOE collaboration) 2005 *Phys. Lett. B* **606** 12
- [138] Davier M 2003 Status of the BaBar R measurement using radiative return *Workshop on e^+e^- Hadronic Cross Section (Pisa, 8–10 October 2003)*
- [139] Chetyrkin K G, Kuhn J H and Steinhauser M 1996 *Nucl. Phys. B* **482** 213
- [140] Barbieri R and Remiddi E 1975 *Nucl. Phys. B* **90** 233
- [141] Friot S, Greynat D and de Rafael E *Lepton Moments 2003*(unpublished) reported by E. de Rafael available at <http://g2pc1.bu.edu/leptonmom/talks/deRafael.pdf>
- [142] Czarnecki A, Marciano W J and Vainshtein A 2003 *Phys. Rev. D* **67** 073006
- [143] Friot S 2005 *PhD Thesis* Université de la Méditerranée (Aix-Marseille II)
Greynat D 2005 *PhD Thesis* Université de la Méditerranée (Aix-Marseille II)
- [144] Hagiwara K, Martin A D, Nomura D and Teubner T 2004 *Phys. Rev. D* **69** 093003 (*Preprint hep-ph/0312250*)
- [145] de Rafael E 1994 *Phys. Lett. B* **322** 239
- [146] Knecht M, Nyffeler A, Perrottet M and de Rafael E 2002 *Phys. Rev. Lett.* **88** 071802
- [147] Ramsey-Musolf M J and Wise M B 2002 *Phys. Rev. Lett.* **89** 041601 (*Preprint hep-ph/0201297*)
- [148] Bijnens J, Pallante E and Prades J 1996 *Nucl. Phys. B* **474** 379
Bijnens J, Pallante E and Prades J 1995 *Phys. Rev. Lett.* **75** 1447

- Bijnens J, Pallante E and Prades J 1995 *Phys. Rev. Lett.* **75** 3781 (erratum)
Bijnens J, Pallante E and Prades J 2002 *Nucl. Phys.* B **626** 410
- [149] Hayakawa M and Kinoshita T 1998 *Phys. Rev. D* **57** 465
Hayakawa M and Kinoshita T 2002 *Phys. Rev. D* **66** 019902 (erratum)
Hayakawa M, Kinoshita T and Sanda A I 1995 *Phys. Rev. Lett.* **75** 790
- [150] Knecht M and Nyffeler A 2002 *Phys. Rev. D* **57** 465
- [151] Manohar A and Georgi H 1984 *Nucl. Phys.* B **234** 189
- [152] Espriu D, de Rafael E and Taron J 1990 *Nucl. Phys.* B **345** 22
- [153] Bijnens J, Bruno Ch and de Rafael E 1993 *Nucl. Phys.* **390** 501
- [154] Peris S, Perrottet M and de Rafael E 1998 *J. High Energy Phys.* [JHEP05\(1998\)011](#)
- [155] de Rafael E 2002 Large- N_c QCD and low energy interactions *QCDWORK, Martina Franca, Bari, Italy, AIP Conf. Proc.* vol 602 ed P Colangelo and G Nardulli p 14
- [156] Melnikov K and Vainshtein A 2004 *Phys. Rev.* **70** 113006
- [157] Bijnens J and Prades J 2007 *Preprint* [hep-ph/0701240](#)
- [158] Studenikin A I 1999 *Yad. Fiz.* **62** 2248
Studenikin A I 1999 *Phys. Atom. Nucl.* **62** 2071
- [159] Czarnecki A, Krause B and Marciano W J 1996 *Phys. Rev. Lett.* **76** 3267
- [160] Peris S, Perrottet M and de Rafael E 1995 *Phys. Lett.* B **355** 523
- [161] Czarnecki A, Krause B and Marciano W 1995 *Phys. Rev. D* **52** R2619
- [162] Kukhto T V, Kuraev E A, Schiller A and Silagadze Z K 1992 *Nucl. Phys.* B **371** 567
- [163] Knecht M, Peris S, Perrottet M and de Rafael E 2002 *J. High Energy Phys.* [JHEP11\(2002\)003](#)
- [164] Vainshtein A 2003 *Phys. Lett.* B **569** 187
- [165] Knecht M, Peris S, Perrottet M and de Rafael E 2004 *J. High Energy Phys.* [JHEP03\(2004\)035](#)
- [166] Jegerlehner F and Tarasov O V 2005 *Preprint* [hep-ph/0510308](#)
- [167] Pasechnik R S and Teryaev O V 2005 *Preprint* [hep-ph/0510290](#)
- [168] Dorokhov A E 2005 *Acta Phys. Pol.* B **36** 3751
- [169] Degradi G and Giudice G F 2003 *Phys. Rev. D* **58** 073006
- [170] Jungmann K P private communication
Jungmann K P 2001 Spectroscopy of the muonium atom *The Hydrogen Atom (Lecture Notes in Physics* vol 570 ed S Karshenboim *et al* pp 81–102
Jungmann 2000 Muonium—physics possibilities at new and intensive muon sources *High Intensity Muon Sources* ed Y Kuno and T Yokoi (Singapore: World Scientific) pp 26–40
- [171] One potential technique is discussed in Farley F J M 2004 *Nucl. Instrum. Methods A* **523** 251
- [172] Davier M and Marciano W J 2004 *Annu. Rev. Nucl. Part Phys.* **54** 115
Farley F J M and Semertzidis Y K 2004 *Prog. Nucl. Part. Phys.* **52** 1
Hertzog D W and Morse W M 2004 *Annu. Rev. Nucl. Part Phys* **54** 141