

Quantum noninvasive three-component beam-spin polarimetry in the Hadron Storage Ring of the Electron-Ion Collider

F. Rathmann

July 2026

Collider Accelerator Department
Brookhaven National Laboratory

U.S. Department of Energy
USDOE Office of Science (SC), Nuclear Physics (NP)

Notice: This technical note has been authored by employees of Brookhaven Science Associates, LLC under Contract No. with the U.S. Department of Energy. The publisher by accepting the technical note for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this technical note, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Quantum noninvasive three-component beam-spin polarimetry in the Hadron Storage Ring of the Electron-Ion Collider

Frank Rathmann

Brookhaven National Laboratory

June 28, 2026

Abstract

We propose a noninvasive SQUID-based polarimeter for the polarized proton beam in the Electron-Ion Collider (EIC) Hadron Storage Ring (HSR), exploiting the collective magnetic dipole moment of the bunches rather than scattering. The six-snake HSR lattice has synchronous-particle spin tune $\nu_s = 1/2$, placing the in-plane spin-precession signal at half the revolution frequency (~ 39 kHz), in the DC SQUID band. Three pickup channels (cosine- θ and sine- θ saddle loops for the transverse components, a coaxial axial gradiometer for the longitudinal one) reconstruct the full polarization vector (P_x, P_y, P_z) in two complementary modes. Static mode, the default for continuous noninvasive monitoring, reads all three components: P_y at the revolution frequency and the residual in-plane components at $\nu_s f_{\text{rev}}$, bunch by bunch over an hours-long fill, including P_z , inaccessible to single-spin scattering polarimetry by parity conservation. Dynamic mode gives a precise polarization-magnitude measurement: a longitudinal kicker tips a small fraction of the polarization into the horizontal (ring) plane to produce a free-induction-decay (FID) signal, and many phase-locked tip- π -echo-restore cycles are summed coherently via a matched filter across all bunches, with $O(\alpha^2/\pi^2) \sim 10^{-4}$ loss per cycle, negligible over a full $\delta P/P = 1\%$ measurement. For tipping angle $\alpha = 30$ mrad, polarization $P = 0.7$, and effective rms spin-tune spread $\sigma_{\nu_s}^{\text{eff}} = 10^{-3}$ (coherence time ~ 2 ms), the integration time to reach $\delta P/P = 1\%$ is about 18 s at injection and 5 min at flattop. The architecture extends to deuteron and ^3He beams via species-specific spin-magnetic factors, with applications to storage-ring EDM searches.

Contents

1	Introduction	3
2	Measurement concept and HSR spin lattice	6
2.1	Six-snake spin lattice	6
2.2	Transverse FID measurement	6
2.3	Local vector-pickup basis	7
2.4	Two-stage measurement strategy	7
3	Spin dynamics and signal geometry	8
3.1	Numerical input parameters	8
3.2	Snake locations and spin tune	9
3.3	Siberian snakes and the spin tune	9
3.4	Spin dynamics, stable spin axis, and spin tune	12
3.4.1	Thomas-BMT equation	12
3.4.2	Arc sector and snake rotation matrices	12
3.4.3	One-turn spin transfer matrix	13
3.4.4	Stable spin axis $\vec{n}_0$ and spin tune ν_s	13

3.4.5	Phasor derivation of $\nu_s = \frac{1}{2}$	13
3.4.6	Lee–Courant and Doubly Lee–Courant configurations for the HSR	14
3.4.7	Numerical results	16
3.5	Net magnetic dipole moment of a polarized bunch	16
3.6	Tipping the spin ensemble: the FID protocol	17
4	Spin-tune measurement: spectral search and phase lock	18
4.1	Spectral-search strategy	19
4.2	Detection feasibility and operating-point numbers	19
4.3	Iterative improvement loop	21
4.4	Transition to phase-locked operation	21
5	Spin-echo rephasing and the noninvasive measurement cycle	23
5.1	Reversible and irreversible decoherence	23
5.2	Echo rephasing and the π -pulse axis	24
5.3	Hardware reuse, phase-locked timing, and field budget	25
5.4	Restore pulse and the closed measurement cycle	27
5.5	The echo amplitude τ -scan and what is measured	29
6	Magnetic flux signal and matched-filter reconstruction	29
6.1	Relativistic field of a moving transverse magnetic dipole	29
6.2	Flux per bunch pass	30
6.3	Bunch spin pattern at injection	31
6.4	Lattice spin phase offsets	32
6.5	Naïve every-other-turn summation	32
6.6	Matched filter: phase-corrected coherent summation	33
6.7	In-plane coherence time and spin tune spread	35
6.8	Coherent integration	35
6.9	Numerical sensitivity estimates	36
7	Design overview: three-channel SQUID polarimeter	37
8	Full spin-vector extension and static vector polarimetry	40
8.1	Static bunch magnetic moment	40
8.2	Natural pickup basis on a cylindrical former	41
8.3	Three-channel single-passage response	41
8.4	Finite bunch-length convolution	42
8.5	Beam-synchronous static signal formation	42
8.6	Background rejection and cross-talk hierarchy	43
8.7	Systematic effects and cross-talk	44
9	Longitudinal component and axial gradiometer channel	44
9.1	Physical picture	44
9.2	Charge current rejection	45
9.3	Flux signal from the longitudinal magnetic moment	46
9.3.1	Point-dipole result	46
9.3.2	Finite bunch length	46
9.4	Gradiometer configuration	46
9.5	Signal-to-noise ratio	47
9.5.1	Signal accumulation: how SNR is built up	47
9.5.2	Strategy 1: all-bunch coherent summation	47
9.5.3	Strategy 2: single-bunch tracking	48
9.5.4	Coherence-time interaction and multi-cycle accumulation	48

9.5.5	Numerical performance and measurement strategies	48
9.5.6	SQUID readout bandwidth	49
9.6	Real-world effects not yet addressed	49
9.7	Implementation	50
9.8	Conclusions for the axial channel	50
10	Bunch-resolved vector polarimetry	51
10.1	Three-channel measurement modes	51
10.2	Comparison with existing polarimetry	52
10.3	Bunch-resolved polarization history	53
11	Extensions and outlook	54
11.1	Storage-ring electric dipole moment searches	54
11.2	Electron storage rings and self-polarized beams	56
12	Summary	56
A	Spin manipulation with an RF Wien filter	62
A.1	Field structure	63
A.2	Idealized spin reversal in the rotating frame	65
A.3	Implementation options	66
A.4	Pulsed-mode RF Wien filter	67

1 Introduction

Conventional proton beam polarimetry at high-energy colliders relies on nuclear scattering reactions: elastic proton-carbon (pC) scattering for fast relative measurements, or Coulomb-Nuclear Interference (CNI) scattering from a polarized atomic-hydrogen jet (HJET) for absolute calibration [1]. Both methods intercept the stored beam with a target (a thin carbon strip for pC, an atomic-hydrogen jet traversed by the full stored beam for HJET), causing measurable beam losses and heating. Beyond this beam-loss sense, pC polarimetry also relies on thin carbon-strip target technology with finite thermal, mechanical, and lifetime margins: a recent study of the thermal and electromechanical response of these targets in relativistic bunched beams indicates that the technology has operational limits and may become inadequate under the more demanding EIC beam conditions, strengthening the motivation for complementary noninvasive diagnostics [2].

The EIC physics program demands a relative polarization uncertainty $\delta P/P \leq 1\%$ (where P is the magnitude of the beam polarization vector and δP is its absolute uncertainty), bunch-by-bunch measurement capability, and full characterization of the polarization vector $\vec{P} = (P_x, P_y, P_z)$ [3]. Existing storage-ring polarimetry at RHIC measures the transverse projections (P_x, P_y) bunch by bunch via pC azimuthal sampling, with the absolute P_y scale calibrated by the HJET polarimeter operated with a vertical target polarization [1]. Dedicated procedures at RHIC combined the transverse pC/STAR measurements with controlled spin rotations and offline transport to recover a fill-averaged three-dimensional stable spin direction [4], but cannot follow $\vec{P}(t)$ bunch by bunch or continuously across a store. The upgraded HJET planned for the EIC will additionally measure P_x and P_z via target-polarization rotation, using techniques established for pp spin-correlation measurements [5]; even with this upgrade, however, no current or planned scattering polarimeter reconstructs the full polarization vector bunch by bunch and continuously throughout the fill. Given the tenfold reduction in bunch spacing relative to RHIC (11.0 ns vs. 106.6 ns) and the threefold increase in average beam current, noninvasive alternatives to scattering-based polarimetry are of particular interest. The three-channel SQUID polarimeter proposed in this paper addresses this gap, supplementing the established scattering-based

capabilities with continuous, noninvasive, bunch-resolved measurement of the full polarization vector $\vec{P}(t)$.

Earlier noninvasive beam-polarimetry concepts pursued two related routes for accessing the collective spin signal electromagnetically. One approach used RF-resonant Stern-Gerlach interactions, in which the magnetic moment of a polarized beam couples to selected cavity modes; this led to proposals for spin-state separation and absolute polarimetry in storage rings [6, 7]. A passive RF cavity for such a Stern-Gerlach polarimeter was carried as far as prototype fabrication and bench characterization for an eventual test with a transversely polarized electron beam at the MIT-Bates storage ring [8], and the relativistic Stern-Gerlach deflection of polarized electron beams has continued to be developed theoretically and in proposed experimental configurations [9, 10]; to date, however, none of these electromagnetic Stern-Gerlach schemes has been demonstrated on a beam. A second approach, closer in spirit to the present work, proposed direct SQUID detection of the magnetic moment of polarized RHIC beams using a pickup loop coupled to a SQUID readout, favoring a room-temperature loop for practical accelerator installation; that proposal already identified the central difficulty, namely that the ordinary beam-current field exceeds the spin field by many orders of magnitude and must be separated by pickup geometry and frequency-domain methods [11]. The present study revisits this direct magnetic-moment approach for the EIC HSR, using the six-snake $\nu_s = 1/2$ spin lattice, the EIC multi-bunch fill pattern, and matched-filter summation of the bunch-resolved SQUID signal.

A closely related and concurrent proposal also treats the stored-beam polarization as a continuous, phase-coherent observable read out noninvasively from the spin-dependent electromagnetic field of the beam [12]. That work couples symmetric pickup electrodes to a high- Q radio-frequency resonator tuned to a MHz bunch-train harmonic, and extracts the electric dipole moment of counter-propagating beams in a dedicated frozen-spin ring from the clockwise-minus-counter-clockwise difference of a controlled spin-wheel precession frequency. The present scheme differs in its detector, frequency regime, observable, and machine. We read the bunch magnetic moment directly with a SQUID flux pickup rather than an electrode-coupled resonator; we operate at the spin-precession sideband $f_s = f_{\text{rev}}/2 \simeq 39$ kHz set by the six-snake $\nu_s = 1/2$ lattice rather than at a MHz bunch harmonic; we reconstruct the full polarization vector (P_x, P_y, P_z) bunch by bunch as a routine beam diagnostic rather than a single electric-dipole-moment frequency difference; and we target the EIC HSR collider with a single circulating beam and matched-filter summation across the EIC fill, rather than a dedicated counter-rotating ring. The two approaches are complementary: the resonant electrode readout is favored at the MHz harmonics of that proposal, where SQUID magnetometry is less natural, whereas the SQUID is the natural sensor at the low-frequency $\nu_s = 1/2$ sideband exploited here. Ref. [12] also identifies polarized-proton and light-ion rings at the EIC as a target for its architecture; the present work develops the complementary SQUID-based, bunch-resolved, full-vector realization specific to the EIC HSR spin lattice.

A SQUID magnetometer can detect the net macroscopic magnetic dipole moment of a polarized bunch without any material interaction with the beam. This paper develops a quantitative design and feasibility analysis of this approach for the EIC HSR, taking full account of the spin dynamics imposed by the six Siberian snakes, the multi-bunch fill structure with its alternating spin-sign pattern, and the matched-filter reconstruction that recovers the full multi-bunch coherence from a SQUID time-stream sampled at the bunch rate. The detector concept uses a three-channel local pickup basis: a saddle cosine- θ loop sensitive to the horizontal transverse dipole component, a saddle sine- θ loop sensitive to the vertical component, and a coaxial axial gradiometer sensitive to the longitudinal component. The three channels are read out simultaneously by independent low-noise dc-SQUID front ends and reconstruct (P_x, P_y, P_z) . The polarimeter is designed to operate in two complementary modes. In *static* mode (no active spin manipulation), the stored vertical polarization P_y appears at the revolution frequency and any residual transverse and longitudinal components precess at $\nu_s f_{\text{rev}}$; all three are read continu-

ously and bunch by bunch throughout the fill, giving direct access to the longitudinal projection P_z that is inaccessible to single-spin scattering polarimetry by parity conservation. In *dynamic* mode, a longitudinal kicker tips the stored polarization by a controlled small angle to produce a free-induction-decay (FID) signal, and many phase-locked tip- π -echo-restore cycles are accumulated coherently via matched-filter summation across all bunches; the closing restore pulse returns the tipped in-plane component to the stable spin axis, so the cumulative cost in beam polarization across a full $\delta P/P = 1\%$ measurement remains negligible. The full engineering implementation including saddle-coil topology, mechanical mounting, magnetic shielding, and cryogenic readout chain is the subject of a dedicated companion paper; the present work establishes the principle, the sensitivity budget, and the operational concept. The same architecture extends to deuteron and ^3He beams via species-specific spin-magnetic factors, with storage-ring electric-dipole-moment applications treated in a dedicated follow-up paper.

The concept is related to the free-precession spin-tune measurement demonstrated by the JEDI collaboration at COSY [13], where the spin vector of a coherently precessing ensemble was sampled turn-by-turn using an internal scattering polarimeter. The JEDI program has produced a series of advances directly relevant to the present work: the establishment of long in-plane spin coherence times through chromaticity correction and sextupole tuning [14], spin-tune mapping as a precision diagnostic [15], the demonstration of phase-locking the spin precession [16], and a comprehensive framework for spin dynamics in storage rings [17]. The deuteron EDM program has further developed the tools for precision spin manipulation and noninvasive polarimetry [18–20]. Without dedicated machine development effort analogous to that at COSY, in particular sextupole tuning to correct chromaticity-driven spin decoherence, achieving long spin coherence times in the HSR cannot be assumed. The SQUID approach replaces the destructive internal polarimeter used at COSY with a passive magnetic pickup, making the measurement fully noninvasive.

Beyond the absolute calibration use case, the static-mode SQUID readout feeds back to the spin lattice and the operating point. With per-bunch matched filtering on the saddle and axial channels, drifts and localized anomalies in (P_x, P_y, P_z) can be resolved bunch by bunch with few-percent precision in measurement windows of order one minute, providing real-time feedback on the machine spin state. The EIC HSR will require tuning of the spin lattice, snake configuration, cooling state, emittance, momentum spread, and bunch intensity to preserve polarization over stores lasting several hours. RHIC experience shows that spin-tune spread is an operationally important and tunable quantity: matching the dispersion-function derivatives at the Siberian snakes reduced the spin-tune spread and substantially improved spin-flip efficiency [21]. At the same time, polarization transmission during acceleration was not determined by spin-tune spread alone, but also depended on deviations from $\nu_s = 1/2$, snake resonances, and interference between intrinsic and imperfection resonances [21]. High-intensity operation adds further dependence on bunch population, emittance growth, momentum tails, beam-beam tune spread, and intrabeam scattering, as observed in RHIC polarized-proton stores [22]. A noninvasive pickup that directly observes the free spin precession after a controlled tipping pulse would therefore provide a new diagnostic observable: the spin-precession frequency, linewidth, coherence envelope, and response to spin-echo or restore-pulse sequences. Such a measurement would be valuable not only for polarimetry, but also for tuning the HSR spin lattice and operating point.

The paper is organized as follows. Section 2 sets the measurement concept and the HSR six-snake spin lattice. Section 3 develops the spin dynamics, the bunch magnetic dipole moment, and the FID protocol. Section 4 treats the initial spin-tune diagnostic via spectral search and phase lock. Section 5 presents the spin-echo rephasing that closes the tip- π -restore cycle and makes the measurement noninvasive. Section 6 establishes the matched-filter reconstruction and the resulting per-bunch sensitivity. Section 7 overviews the three-channel SQUID polarimeter pickup. Sections 8 and 9 extend the treatment to static transverse vector polarimetry and to the longitudinal P_z via the axial gradiometer. Section 10 develops bunch-resolved (P_x, P_y, P_z)

reconstruction over the full fill as the operational deliverable. An appendix covers spin manipulation with an RF Wien filter, with an emphasis on the two-branch resonance degeneracy that arises at $\nu_s = 1/2$ in the HSR.

2 Measurement concept and HSR spin lattice

The SQUID polarimeter considered here measures the magnetic field produced by the collective magnetic dipole moment of the stored polarized bunches. It therefore does not require an internal target or a scattering reaction. The relevant signal is not the ordinary beam-current field, but the much smaller field associated with the aligned proton magnetic moments. The measurement concept uses the known HSR spin lattice to place this spin signal at a well-defined frequency and then uses bunch-resolved phase information to recover it coherently.

2.1 Six-snake spin lattice

In the EIC HSR, six full Siberian snakes are used to enforce the synchronous-particle spin tune

$$\nu_s = \frac{1}{2}. \quad (1)$$

For the Lee-Courant configuration considered here, the stable spin axis in suitable arc regions is approximately vertical,

$$\vec{n}_0 \simeq \vec{e}_y. \quad (2)$$

A spin component perpendicular to $\vec{n}_0$ therefore changes sign from one turn to the next. This turn-by-turn sign reversal is the central feature used by the SQUID measurement. It shifts the transverse spin signal to half the revolution frequency, while the ordinary beam-current field remains tied to the usual revolution harmonics and bunch pattern.

The detailed stable-spin-axis calculation and the corresponding six-snake layout are discussed below. In this section we only use the qualitative consequence of Eq. (1): after a controlled tipping pulse, the transverse component of the bunch magnetic moment oscillates with the spin phase and reverses sign on successive turns.

2.2 Transverse FID measurement

A short tipping pulse (produced by a longitudinal magnetic field B_z applied for a fraction of a microsecond, as detailed in Sec. 3.6; a rotation about B_z is a feasible option here, since it does not produce a transverse kick on the closed orbit) applied in a region with $\vec{n}_0 \simeq \vec{e}_y$ rotates a small fraction of the stored vertical polarization into the ring plane. Immediately after this pulse, the bunch carries a coherent transverse magnetic moment. As illustrated in Fig. 1, the transverse projection changes sign on successive turns for $\nu_s = 0.5$. This produces a free-induction-decay (FID) signal in the SQUID pickup at

$$f_s = \nu_s f_{\text{rev}} \simeq \frac{1}{2} f_{\text{rev}}. \quad (3)$$

For the EIC HSR this is near 39 kHz, in the useful bandwidth range of wideband DC SQUID readout.

The measurement sequence is therefore a tip-read-restore cycle. The tipping pulse creates the transverse spin component, the SQUID pickup records the FID signal over the available spin-coherence window, and a phase-locked restore pulse returns the in-plane component close to the stable-spin direction. The numerical baseline used later is a tipping angle of $\alpha = 30$ mrad and an effective rms spin-tune spread $\sigma_{\nu_s}^{\text{eff}} = 10^{-3}$ (Table 1), with these values entering only when the coherence time and signal-to-noise ratio are estimated.




Figure 1: Schematic of the spin-precession geometry for a guide field along the vertical direction $\vec{e}_y$. A tipping pulse creates a transverse magnetic moment $\vec{m}$ of the bunch. For $\nu_s = 0.5$, the transverse projection advances by π per turn, so the $\vec{e}_x$ component changes sign between turns n and $n + 1$. This produces the spin signal at $f_s = f_{\text{rev}}/2$.

2.3 Local vector-pickup basis

The transverse FID signal can be detected with two orthogonal pickup channels. A cosine- θ pickup couples to the transverse magnetic-field pattern associated with the local P_x component, while a sine- θ pickup couples to the corresponding P_y component. In this basis, the two transverse channels provide local sensitivity to

$$P_x \quad \text{and} \quad P_y. \quad (4)$$

The same pickup platform can be extended with an axial loop-gradiometer channel. This channel is sensitive to the longitudinal component

$$P_z, \quad (5)$$

and is geometrically insensitive to the ideal azimuthal magnetic field of a centered beam current. Figure 2 shows this three-channel pickup basis schematically on a common cylindrical support around the beam pipe.

In the arc locations emphasized here, the stored polarization is expected to be close to vertical before the tipping pulse, so the transverse FID channels are the primary measurement mode. The axial channel is nevertheless useful as part of a local vector-polarimetry basis and as a relative diagnostic for longitudinal spin components or spin leakage.

The three-channel pickup basis therefore groups the polarization measurement into two related parts. The cosine- θ and sine- θ channels measure the transverse FID signal and provide the main sensitivity to P_x and P_y . The axial gradiometer provides the corresponding P_z sensitivity. All three channels are gradiometric in the sense that they reject the dominant azimuthal beam-current field, but the rejection mechanism differs: the saddle pickups use their first-harmonic ($\cos \theta$ or $\sin \theta$) winding pattern to suppress the azimuthally symmetric current field, while the coaxial axial loop measures axial flux and is therefore intrinsically blind to that field by geometry. The detailed pickup geometry, field coupling, and common-mode rejection requirements are analyzed in later sections.

2.4 Two-stage measurement strategy

The SQUID FID signal supports two operating modes that address different measurement goals. In the initial mode (Sec. 4), a tipping pulse with a sufficiently large tipping angle is followed by




Figure 2: Conceptual schematic of the three-channel SQUID pickup basis on a common cylindrical support around the beam pipe. The beam direction is along the cylinder axis. The blue saddle-loop pair represents the cosine- θ transverse pickup, which is sensitive to the local P_x component. The red saddle-loop pair represents the sine- θ transverse pickup, rotated by 90° around the cylinder and sensitive to the local P_y component. The green pair of coaxial loops forms an axial gradiometer, sensitive to the longitudinal component P_z . The drawing is schematic and illustrates the pickup symmetries and common geometric center, not the detailed winding implementation.

a short observation within one coherence window, providing a signal-to-noise ratio of order unity in a single record before the spin coherence is known to last longer. This first record determines the spin-precession frequency $f_s = \nu_s f_{\text{rev}}$ and the spin-tune spread σ_{ν_s} from the spectral peak and its linewidth, respectively.

In the precision mode (Sec. 6.8), f_s and the available coherence window are taken as known from the initial mode. With the tipping angle $\alpha = 30 \text{ mrad}$, the closing restore pulse returns the in-plane component to the stable spin axis with a per-cycle loss of only $\mathcal{O}(\alpha^2/\pi^2) \sim 10^{-4}$, and many tip-read-restore cycles are accumulated coherently using matched-filter summation. Throughout each cycle the spin-precession phase is tracked continuously from a reference oscillator phase-locked to f_s [16], so that the closing restore pulse is applied at the precession phase required to rotate the in-plane component back to the stable spin axis. The relative statistical precision on the polarization magnitude P then improves as $T^{-1/2}$ in the total integration time T , with the beam polarization preserved across the measurement.

3 Spin dynamics and signal geometry

3.1 Numerical input parameters

We use the convention $\vec{e}_z$ = beam direction (longitudinal), $\vec{e}_y$ = vertical (along the arc dipole guide field), and $\vec{e}_x$ = horizontal transverse throughout. The relevant EIC Hadron Storage Ring parameters are summarised in Table 1. All numerical values are taken from Table II of Ref. [1]. The HSR reuses the RHIC Yellow ring tunnel (circumference $L = 3833.85 \text{ m}$). The revolution frequency at injection is slightly lower than at flattop because the proton velocity at 23.5 GeV ($\gamma = 25.05$) is marginally below c .

At injection the HSR is filled with 290 bunches at 44.1 ns spacing. After the energy ramp, RF bunch splitting increases the population to 1160 bunches at 11.0 ns spacing. Each parent bunch produces four daughter bunches that inherit its polarisation sign. There are no plans to change the spin pattern after injection. Although noninvasive spin-pattern manipulation using a radio-frequency Wien filter appears technically feasible, as demonstrated at COSY and discussed in the context of future collider experiments [19], such a device is not part of the baseline HSR design.

Table 1: Primary EIC HSR proton beam and restored-FID input parameters used in this analysis. Beam parameters are based on Table II of Ref. [1]. Derived quantities such as revolution frequency, spin-precession frequency, coherence time, sideband linewidth, and SNR integration times are computed from these inputs.

Parameter	Symbol	Unit	EIC injection	EIC flattop
Ring circumference	C	m	3833.85	same
Total beam energy	E	GeV	23.5	275
Lorentz factor	γ	—	25.05	293.1
Relativistic factor	β	—	0.9992	0.999 994
Protons per bunch	N_p	$\times 1 \times 10^{10}$	27.6	6.9
Number of bunches	N_b	—	290	1160
Bunch spacing	T_b	ns	44.1	11.0
RMS bunch width	σ_t	ns	0.801	0.200
RMS bunch length	σ_L	m	0.24	0.06
Average beam current	I_{avg}	A	1.002	1.003
Beam polarization	P	—	0.70	same
FID tipping angle	α	mrاد	30	same
Effective RMS spin-tune spread	$\sigma_{\nu_s}^{\text{eff}}$	—	1.0×10^{-3}	same
Spin tune	ν_s	—	0.5	same

The effective value $\sigma_{\nu_s}^{\text{eff}} = 1.0 \times 10^{-3}$ is used here as a working rms spin-tune-spread input for the FID coherence and matched-filter estimates. It should not be interpreted as a fixed machine constant: spin-tune spread and polarization transmission depend on the detailed HSR lattice, snake matching, momentum spread, and operational tuning state, as emphasized in HSR spin-tracking and snake-matching studies [23, 24].

3.2 Snake locations and spin tune

The EIC HSR has six interaction regions and six helical Siberian snakes distributed around the ring, as shown in Fig. 3. The figure shows the simplified hexagonal representation used in this note: six equal arc sectors separated by localized 180° snake rotations. The red lines indicate the Lee–Courant snake-axis pattern used for the analytical spin model.

3.3 Siberian snakes and the spin tune

The HSR employs six helical superconducting Siberian snakes, each producing a spin rotation of 180° about a horizontal axis. The six snakes are arranged as three pairs. Within each pair, the rotation axes are chosen to be perpendicular, so that each pair contributes an effective spin rotation by π . The product of the three pair rotations is again an odd multiple of π , which fixes




Figure 3: Schematic layout of the EIC HSR in hexagonal representation. The six interaction regions IR2–IR12 are located at the midpoints of the six straight sections. Each straight section hosts one Siberian snake. The red line through each snake box indicates the Lee–Courant rotation axis.

the one-turn spin tune of the synchronous particle at

$$\nu_s = \frac{1}{2} \quad (6)$$

independently of $G\gamma$ [24]. This removes the direct crossing of first-order imperfection resonances (at integer values of $G\gamma$) and of intrinsic spin resonances during acceleration.

The increase from two snakes in RHIC to six snakes in the HSR is driven by the stronger spin-resonance structure of the EIC lattice and by the need to accelerate additional polarized ion species. In particular, polarized ^3He has an anomalous magnetic moment $G = -4.19$ (Table 2), more than twice as large in magnitude as the proton value. Spin resonances are therefore correspondingly stronger for ^3He , which further motivates the six-snake scheme.

Because each snake produces a full 180° spin rotation, the six-snake configuration fixes $\nu_s = 1/2$ independently of energy, so the half-integer spin tune can be maintained throughout acceleration rather than only at collision energy. When the spin rotators reach their nominal

values at collision energy, the snake currents are adjusted simultaneously to compensate for the resulting spin-tune shift [25]. The detector solenoid does not change the spin tune, but it can perturb the local stable spin axis $\vec{n}_0(s)$. Matching the injected polarization direction from the AGS to the HSR stable spin direction is therefore still required.

It is important to distinguish the spin tune from the stable spin axis. The six-snake configuration fixes $\nu_s = 1/2$ for the synchronous particle, but the local stable spin axis $\vec{n}_0(s)$ is determined by the full sequence of snake axes, arc spin rotations, spin rotators, detector solenoids, and local lattice fields. Thus $\vec{n}_0(s)$ is a local lattice quantity, not simply a consequence of the value $\nu_s = 1/2$.

The snake-axis orientations $\{\phi_k\}$ are therefore free design parameters with two separate roles. They preserve the energy-independent spin tune, but they also determine the direction of $\vec{n}_0(s)$ and the suppression of higher-order spin resonances during the energy ramp. Of particular importance for the HSR is the Doubly Lee-Courant (DLC) configuration [24], in which the snake-axis pattern is chosen to provide local cancellation of spin-phase errors and robust polarization transmission. In the idealized DLC picture, the stable spin axis remains close to vertical in suitable arc regions,

$$\vec{n}_0(s) \simeq \vec{e}_y. \quad (7)$$

This vertical-axis approximation is the working assumption used in the simplified signal estimates below.

For the actual HSR implementation, however, the relevant quantity is the local $\vec{n}_0(s)$ map from the machine spin-tracking model. The placement and orientation of any spin-manipulation element, including the FID tipping element, must be chosen so that its spin-kick axis is transverse to the local $\vec{n}_0(s)$ at that location. In an arc region where $\vec{n}_0 \simeq \vec{e}_y$, a longitudinal field B_z provides such a transverse spin kick that rotates a small fraction of the vertical polarization into the ring plane, where it precesses and flips the sign of the transverse moment m_x on successive turns at $\frac{1}{2}f_{\text{rev}}$, as shown in Fig. 1. If $\vec{n}_0(s)$ has a significant horizontal component at a candidate location, the required field orientation must be adjusted accordingly.

Table 2 also indicates how the SQUID signal generalizes to other polarized ion species. The SQUID couples to the bunch net magnetic dipole moment, which is by construction a rank-1 (vector) observable; tensor polarizations of any rank are not accessible. For spin-1/2 particles such as the proton and ^3He , the polarization state is purely vectorial and the SQUID signal is proportional to P . For spin-1 species (deuteron, ^6Li), the polarization state contains both a vector polarization p_z and a rank-2 tensor polarization p_{zz} [26]; the SQUID measures p_z and is blind to p_{zz} . For spin-3/2 (^7Li), ranks 0–3 are present; the SQUID still picks up only the rank-1 component. Rank- k tensor components precess at $k\omega_s$, so higher ranks would alias to harmonics of the spin-precession sideband, but they produce no magnetic flux and therefore no SQUID signal at any harmonic. For ^3He , the analysis is identical to the proton case with $G = -4.19$ substituted throughout. For deuterons and ^6Li , $|G| \approx 0.14$ – 0.18 , so fewer spin resonances are crossed during acceleration, but the SQUID spin signal per nucleon is also reduced.

Table 2: Spin and magnetic properties of ion species relevant for the EIC, using the BMT definition $g = A\mu/(ZI\mu_N)$ and $G = (g - 2)/2$. Magnetic moments from Stone [26]; proton and deuteron values from CODATA [27].

Nucleus	Z	A	I	μ (μ_N)	$g = A\mu/(ZI\mu_N)$	$G = (g - 2)/2$
p	1	1	1/2	2.7928	5.5856	1.7928
^3He	2	3	1/2	-2.1276	-6.3828	-4.1914
d	1	2	1	0.8574	1.7148	-0.1426
^6Li	3	6	1	0.8220	1.6440	-0.1780
^7Li	3	7	3/2	3.2564	5.0655	1.5328

3.4 Spin dynamics, stable spin axis, and spin tune

3.4.1 Thomas-BMT equation

The spin vector $\vec{S}$ of a relativistic particle with charge q , mass m , and anomalous magnetic moment $G = (g - 2)/2$ obeys the Thomas-Bargmann-Michel-Telegdi (T-BMT) equation in the laboratory frame [28, 29]

$$\frac{d\vec{S}}{dt} = \vec{\Omega}_{\text{BMT}} \times \vec{S}, \quad (8)$$

with precession vector

$$\vec{\Omega}_{\text{BMT}} = -\frac{q}{m} \left[\left(G + \frac{1}{\gamma} \right) \vec{B} - \frac{G\gamma}{\gamma+1} (\vec{\beta} \cdot \vec{B}) \vec{\beta} - \left(G + \frac{1}{\gamma+1} \right) \frac{\vec{\beta} \times \vec{E}}{c} \right], \quad (9)$$

where $\vec{\beta} = \vec{v}/c$, $\gamma = (1 - \beta^2)^{-1/2}$, and $\vec{B}$, $\vec{E}$ are the laboratory magnetic and electric fields.

For the HSR the following simplifications apply. The beam is ultrarelativistic ($\gamma \gg 1$), so $1/\gamma \approx 0$ and $\gamma/(\gamma+1) \approx 1$. On the design orbit there are no electric fields ($\vec{E} = 0$) and the bending field is purely transverse ($\vec{\beta} \cdot \vec{B} = 0$). This $\vec{\beta} \cdot \vec{B} = 0$ simplification is exact only on the design closed orbit; for particles with finite longitudinal momentum spread or transverse motion, and in regions where the equilibrium $\vec{\beta}$ acquires a transverse component (snake and spin-rotator interiors, low- β^* insertions near the interaction points), $\vec{\beta} \cdot \vec{B}$ is generally nonzero and contributes small additional spin rotations beyond the closed-orbit precession. These contributions are the dominant source of the spread of the stable spin direction across phase space, and hence of the residual spin-tune spread σ_{ν_s} that drives the coherence-time analysis of Sec. 6.8 and the working value adopted in Table 3. Transforming from t to arc length s via $dt = ds/v$ and using $p = qB\rho$ (magnetic rigidity), Eq. (9) reduces to

$$\vec{\Omega}(s) = -\frac{q}{p} (G\gamma + 1) B_{\perp}(s) \vec{e}_{\perp}, \quad (10)$$

where $B_{\perp}$ is the field component perpendicular to the design orbit. In a vertical dipole field $\vec{B} = B_y \vec{e}_y$ the precession is purely around $\vec{e}_y$ at rate $G\gamma/\rho$ per unit arc length, so the spin accumulates a total precession angle $G\gamma\theta$ over a sector of orbital bending angle θ , giving the arc rotation matrix of Eq. (12). The quantity $G = 1.7928$ for the proton. The T-BMT equation in arc-length form is

$$\frac{d\vec{S}}{ds} = \vec{\Omega}(s) \times \vec{S}. \quad (11)$$

3.4.2 Arc sector and snake rotation matrices

In a pure vertical dipole field the spin precesses around $\vec{e}_y$ by $G\gamma\theta$ for an arc sector of orbital bending angle θ , leading to

$$\mathbf{R}_{\text{arc}}(\theta) = \mathbf{R}_y(G\gamma\theta), \quad \mathbf{R}_y(\varphi) = \begin{pmatrix} \cos \varphi & 0 & \sin \varphi \\ 0 & 1 & 0 \\ -\sin \varphi & 0 & \cos \varphi \end{pmatrix}. \quad (12)$$

Snake k rotates the spin by 180° around the horizontal axis $\vec{n}_k = \sin \phi_k \vec{e}_x + \cos \phi_k \vec{e}_z$, where ϕ_k is measured from the local beam direction $\vec{e}_z$ [30]. Applying the Householder reflection formula [31], $\vec{S}' = 2(\vec{S} \cdot \vec{n}_k) \vec{n}_k - \vec{S}$, where $\vec{S}'$ is the spin after the snake rotation, yields

$$\mathbf{R}_{\text{snake}}(\phi_k) = \begin{pmatrix} -\cos 2\phi_k & 0 & \sin 2\phi_k \\ 0 & -1 & 0 \\ \sin 2\phi_k & 0 & \cos 2\phi_k \end{pmatrix}. \quad (13)$$

The action on the spin components is most transparent for three special cases:

Axis	ϕ_k	S'_x	S'_y	S'_z
$\vec{e}_z$ (longitudinal)	0°	$-S_x$	$-S_y$	$+S_z$
diagonal	45°	$+S_z$	$-S_y$	$+S_x$
$\vec{e}_x$ (radial)	90°	$+S_x$	$-S_y$	$-S_z$

In every case $S_y \rightarrow -S_y$ because the snake axis is always horizontal. The component along $\vec{n}_k$ is preserved; the component perpendicular to $\vec{n}_k$ in the xz plane is flipped. The arc rotations act only in the xz plane and leave S_y unchanged.

3.4.3 One-turn spin transfer matrix

The one-turn matrix $\mathbf{M}(s_0)$ is the ordered product of arc and snake matrices for one full revolution. For the analytical six-snake model (six equal arcs, $\theta = \pi/3$, six point-like snakes), one obtains

$$\mathbf{M}(s_0) = \prod_{k=6}^1 \mathbf{R}_{\text{snake}}(\phi_k) \mathbf{R}_{\text{arc}}(\pi/3), \quad (14)$$

evaluated in the order of traversal (indices modulo 6).

3.4.4 Stable spin axis $\vec{n}_0$ and spin tune ν_s

The one-turn matrix $\mathbf{M}(s)$ belongs to $SO(3)$, the group of proper rotations in $\mathbb{R}^3$, defined by $\mathbf{M}^T \mathbf{M} = \mathbf{I}$ and $\det(\mathbf{M}) = +1$. Its three eigenvalues therefore take the form $\{+1, e^{+2\pi i \nu_s}, e^{-2\pi i \nu_s}\}$, where the real eigenvalue $+1$ always exists because every $SO(3)$ rotation has a fixed axis. The stable spin axis $\vec{n}_0(s)$ is the eigenvector belonging to $\lambda = +1$: it is the direction around which all other spin components precess, and it returns to itself after every full revolution:

$$\mathbf{M}(s) \vec{n}_0(s) = +\vec{n}_0(s). \quad (15)$$

The complex pair $e^{\pm 2\pi i \nu_s}$ describes the precession of spin components perpendicular to $\vec{n}_0$: they rotate by angle $2\pi \nu_s$ per turn. Since $|e^{\pm 2\pi i \nu_s}| = 1$ the precession is length-preserving, and the spin vector does not grow or decay. The spin tune ν_s is extracted from the matrix trace via

$$\text{tr}(\mathbf{M}) = 1 + 2 \cos(2\pi \nu_s). \quad (16)$$

3.4.5 Phasor derivation of $\nu_s = \frac{1}{2}$

Since every arc leaves S_y unchanged and every snake flips $S_y \rightarrow -S_y$, after six snakes $S_y \rightarrow (-1)^6 S_y = +S_y$. Hence $\vec{e}_y$ is always an eigenvector of $\mathbf{M}$ with eigenvalue $+1$: $\vec{n}_0 = \vec{e}_y$ in the arc regions, independently of energy.

To find ν_s we work in the xz plane using the complex variable $w = S_x + iS_z$, in which a rotation by angle φ becomes multiplication by $e^{i\varphi}$ and a reflection becomes complex conjugation. The spin precession angle accumulated per arc sector is

$$\alpha = G\gamma \frac{\pi}{3}, \quad (17)$$

which depends on the beam energy through $G\gamma$. The two building blocks in the xz plane are:

$$\text{Arc: } w \rightarrow e^{-i\alpha} w, \quad (18)$$

$$\text{Snake } k: w \rightarrow -e^{-2i\phi_k} \bar{w}. \quad (19)$$

The conjugation $\bar{w} = S_x - iS_z$ in the snake operation flips the sign of S_z , equivalent to a reflection in the xz plane; the prefactor $-e^{-2i\phi_k}$ rotates this reflection to the actual snake axis orientation.

The key observation is that the conjugation in each snake turns the arc phase $e^{-i\alpha}$ accumulated before it into $e^{+i\alpha}$, which then cancels exactly with the $e^{-i\alpha}$ of the following arc. Tracing through two consecutive arc-snake pairs:

$$w \xrightarrow{A_1} e^{-i\alpha} w \xrightarrow{S_1} -e^{-2i\phi_1} e^{+i\alpha} \bar{w} \xrightarrow{A_2} -e^{-2i\phi_1} \bar{w} \xrightarrow{S_2} e^{2i(\phi_1-\phi_2)} w. \quad (20)$$

The arc phase α has cancelled completely after each pair. After all three pairs of arc-snake steps the one-turn map in the xz plane is a pure rotation by the total spin phase Φ :

$$w_{\text{final}} = e^{i\Phi} w_0, \quad \Phi = -2 \sum_{k=1}^3 (\phi_{2k} - \phi_{2k-1}), \quad (21)$$

independent of $G\gamma$. Comparing with $e^{\pm 2\pi i \nu_s}$ from Eq. (16), $\nu_s = 1/2$ requires $\Phi = \pi$.

With the LC axes $\phi_k = (-1)^{k+1} \times 45^\circ$, each pair contributes $\phi_{2k} - \phi_{2k-1} = -90^\circ$, giving:

$$\Phi = -2 \times 3 \times (-90^\circ) = +540^\circ \equiv \pi \pmod{2\pi}. \quad (22)$$

Hence $e^{i\Phi} = e^{i\pi} = -1$, and the one-turn trace is $\text{tr}(\mathbf{M}) = 1 + 2\cos(\pi) = -1$, giving $\nu_s = 1/2$ at every beam energy. This is verified numerically in Section 3.4.7.

3.4.6 Lee–Courant and Doubly Lee–Courant configurations for the HSR

The HSR employs six full Siberian snakes, distributed approximately uniformly around the ring. Each snake rotates the spin by 180° around a horizontal axis with angle ϕ_k . For an even number of full snakes, the closed-orbit spin tune is fixed at

$$\nu_s = \frac{1}{2}, \quad (23)$$

independently of energy. The snake-axis pattern therefore does not primarily determine the closed-orbit spin tune. Instead, it determines the closed-orbit stable spin direction $\vec{n}_0(s)$ and the degree to which higher-order spin-orbit resonances, and therefore amplitude-dependent spin-tune spread, are suppressed during the ramp.

The Lee–Courant (LC) condition imposes local cancellation of spin motion across selected snake pairs [32]. In the simplified six-snake model with equal 60° arc sectors, this corresponds to successive snake axes differing by 90° in absolute value,

$$|\phi_{k+1} - \phi_k| = 90^\circ. \quad (24)$$

A convenient representative is

$$\phi_k = (-1)^{k+1} \times 45^\circ, \quad k = 1, \dots, 6, \quad (25)$$

i.e., $\phi = [+45^\circ, -45^\circ, +45^\circ, -45^\circ, +45^\circ, -45^\circ]$. These are the axes shown in Fig. 3. In the idealized analytical model this configuration gives $\vec{n}_0 = \vec{e}_y$ throughout the arc regions, as confirmed numerically in Section 3.4.7.

A stronger symmetry is obtained in the Doubly Lee–Courant (DLC) scheme. The DLC condition enforces the local π spin-phase cancellation across every consecutive pair of snakes, so that each snake participates in two overlapping cancellation segments. For six snakes this reduces the allowed axis pattern to one continuous overall rotation angle ψ ,

$$\phi_k = \psi + (k-1) \times 45^\circ, \quad k = 1, \dots, 6. \quad (26)$$

Equivalently, neighboring axes step by 45° around the ring. Tracking studies of the realistic HSR lattice [23] find an optimized DLC pattern with $\psi \simeq 22.5^\circ$, giving

$$\phi \simeq [22.5^\circ, 67.5^\circ, 112.5^\circ, 157.5^\circ, 202.5^\circ, 247.5^\circ]. \quad (27)$$

The important point for the present SQUID estimate is that the DLC condition suppresses spin-orbit resonance driving terms over the full energy ramp without requiring dynamic adjustment of the snake axes.

Recent HSR spin-tracking studies by Hamwi provide useful context for the expected scale of spin-tune spread in different six-snake configurations, especially the snake-axis comparison in Sec. 4.4.4 and Figs. 4.12–4.13 of Ref. [33]. In that work, Lee–Courant schemes are shown to improve the polarization-preserving properties of the HSR snake lattice, while the Doubly Lee–Courant (DLC) scheme imposes a stronger overlapping local-cancellation condition in which each snake participates in two cancellation segments.

Hamwi quantifies the residual spin spread through the amplitude-dependent spin tune (ADST). The ADST is the dependence of the spin tune on the orbital actions of a particle, caused by spin-orbit coupling in the real lattice. Particles at different transverse amplitudes therefore do not all precess with exactly the closed-orbit spin tune $\nu_s = 1/2$, but acquire slightly different spin tunes. This spread is precisely the quantity that limits the coherent transverse-spin signal. Hamwi reports the maximum ADST spread over the tracked ensemble and energy range. In the present discussion we denote the corresponding conservative spin-tune-spread scale by σ_{ν_s} , to keep the notation consistent with the coherence-time estimate. The resulting hierarchy is

$$\begin{aligned}\sigma_{\nu_s}^{\text{DLC}} &\lesssim 10^{-2}, \\ \sigma_{\nu_s}^{\text{LC}} &\sim 0.02\text{--}0.04, \\ \sigma_{\nu_s}^{\text{gen}} &\sim 0.04\text{--}0.08.\end{aligned}\tag{28}$$

These values should be understood as configuration-dependent spread scales inferred from maximum ADST measures, not as final rms values for the stored beam. Nevertheless, they show that the snake-axis configuration is a central control knob for σ_{ν_s} and provide a scale against which the coherent integration assumption should be tested.

Table 3: Input parameters and spin-tune-spread scales used for the analytical six-snake model and SQUID coherence estimates. The working value $\sigma_{\nu_s}^{\text{eff}} = 1.0 \times 10^{-3}$ is used for the baseline FID coherence and matched-filter estimates. The LC, DLC, and general-axis entries are conservative spin-tune-spread scales inferred from Hamwi’s maximum ADST measures [33] and are used as test scales rather than fixed rms beam parameters.

Description	Parameter	Unit	Value
Proton anomalous magnetic moment	G	—	1.7928
Lorentz factor (injection)	γ_{inj}	—	25.05
Lorentz factor (flatop)	γ_{flat}	—	293.1
Number of snakes	N_s	—	6
Arc bending angle per sector	θ	$^\circ$	60
LC axis angles	ϕ_k	$^\circ$	alternating ± 45
DLC axis offset	ψ	$^\circ$	$\simeq 22.5$
FID tipping angle	α	mrad	30
Working effective RMS spin-tune spread	$\sigma_{\nu_s}^{\text{eff}}$	—	1.0×10^{-3}
DLC ADST spin-tune-spread scale	$\sigma_{\nu_s}^{\text{DLC}}$	—	$\lesssim 1.0 \times 10^{-2}$
LC ADST spin-tune-spread scale	$\sigma_{\nu_s}^{\text{LC}}$	—	$2.0 \times 10^{-2}\text{--}4.0 \times 10^{-2}$
General-axis ADST spin-tune-spread scale	$\sigma_{\nu_s}^{\text{gen}}$	—	$4.0 \times 10^{-2}\text{--}8.0 \times 10^{-2}$
Revolution frequency (injection)	$f_{\text{rev},\text{inj}}$	kHz	78.1339
Revolution frequency (flatop)	$f_{\text{rev},\text{flat}}$	kHz	78.1957

3.4.7 Numerical results

To illustrate the spin motion in the LC configuration introduced in Sec. 3.4.6, we evaluate the transport for the hexagonal six-snake model shown in Fig. 3. The calculation uses the one-turn map of Eq. (14) together with the sector-wise rotations of Eqs. (13) and (12). In this analytical model, the lattice is represented by six identical arc sectors separated by localized snake rotations at the positions S_k , $k = 1, \dots, 6$. The orbit is discretized into straight sections and curved arc segments so that the continuous evolution in the arcs and the localized snake rotations can be resolved explicitly.

The invariant spin axis $\vec{n}_0$ is obtained from the eigenvalue condition of Eq. (15) and transported along the ring using the same sequence of spin maps, thereby defining the local stable spin direction $\vec{n}_0(s)$ along the closed orbit. Due to the localized action of the snakes, $\vec{n}_0(s)$ undergoes finite rotations at the snake locations and evolves continuously in the arcs. A physical spin vector $\vec{S}(s)$ is initialized at the reference point with a fixed opening angle ψ relative to $\vec{n}_0$ and propagated with the same spin-transfer sequence, ensuring that the angle between $\vec{S}(s)$ and $\vec{n}_0(s)$ remains constant while the phase around $\vec{n}_0(s)$ evolves.

The resulting spin lattice is shown in Fig. 4. It corresponds to the same hexagonal layout and snake labeling shown in Fig. 3. The piecewise behavior of $\vec{n}_0(s)$ reflects the symmetry of the lattice, with finite rotations at the localized snakes and smooth evolution in the intervening arc sectors. The spin vectors $\vec{S}(s)$ populate the corresponding precession cone with fixed opening angle ψ . The different blue vectors at each selected location represent distinct phase angles around $\vec{n}_0(s)$ and are included to visualize the local precession cone; they do not represent subsequent turns of a single particle. All parameters used in the calculation correspond to the LC configuration defined in Eq. (25) and are summarized in Table 3.

3.5 Net magnetic dipole moment of a polarized bunch

A proton has magnetic moment

$$\mu_p = (1 + G)\mu_N = 1.410\,583 \times 10^{-26} \text{ J T}^{-1}, \quad (29)$$

where $\mu_N = e\hbar/(2m_p)$ is the nuclear magneton and $G = 1.7928$ is the proton anomalous magnetic moment [27]. For a fully polarized bunch of N_p protons with spins aligned along the stable spin axis $\vec{n}_0 \approx \vec{e}_y$, the net magnetic dipole moment is

$$\vec{m} = N_p \mu_p \vec{e}_y. \quad (30)$$

For a partially polarized beam with polarization P , this becomes

$$\vec{m} = P N_p \mu_p \vec{e}_y. \quad (31)$$

In steady state the moment is static along $\vec{e}_y$; the cosine- θ transverse SQUID pickup, which is sensitive to the local m_x component, therefore sees no FID signal before a tipping pulse is applied. How an oscillating in-plane signal is generated via a tipping pulse is discussed in Sec. 3.6.

Since $\nu_s = 0.5$, once tipped, the in-plane component reverses sign every turn, as shown in Fig. 1. The moment sampled by a fixed cosine- θ pickup after n turns is

$$m_x(n) = P N_p \mu_p \sin \alpha (-1)^n, \quad (32)$$

where α is the tipping angle. This is the same spin-precession signal introduced in Eq. (3), with

$$f_s = \nu_s f_{\text{rev}} = \frac{1}{2} f_{\text{rev}}. \quad (33)$$

Numerically, this gives $f_s = 39.0669 \text{ kHz}$ at injection and $f_s = 39.0979 \text{ kHz}$ at flattop.




Figure 4: Numerical spin lattice for the hexagonal six-snake model introduced in Fig. 3. The black squares mark the localized Siberian snakes S_k , $k = 1, \dots, 6$, and the gray curve represents the simplified closed orbit. The red vectors show the transported stable spin axis $\vec{n}_0(s)$ at selected locations, while the blue vectors show spin vectors $\vec{S}(s)$ with a fixed opening angle $\psi = 30^\circ$ around $\vec{n}_0(s)$. The multiple blue vectors at a given location represent different phase angles on the local precession cone; they do not represent successive turns of one particle.

For a single fully polarized bunch at injection and flattop (Table 1, Eq. (29)), the magnetic-moment scale is

$$|m|^{\text{inj}} = N_p^{\text{inj}} \mu_p \approx 3.893 \times 10^{-15} \text{ J T}^{-1}, \quad (34)$$

$$|m|^{\text{flat}} = N_p^{\text{flat}} \mu_p \approx 9.733 \times 10^{-16} \text{ J T}^{-1}. \quad (35)$$

The fourfold reduction at flattop reflects RF bunch splitting. Signal estimates in subsequent sections use $P = 0.7$ and include the $\sin \alpha$ factor from Eq. (32). For the operating tipping angle $\alpha = 30 \text{ mrad}$ (Table 1), the corresponding FID magnetic moments are

$$|m_x|_{\text{FID}}^{\text{inj}} = P N_p^{\text{inj}} \mu_p \sin \alpha \approx 8.174 \times 10^{-17} \text{ J T}^{-1}, \quad (36)$$

$$|m_x|_{\text{FID}}^{\text{flat}} = P N_p^{\text{flat}} \mu_p \sin \alpha \approx 2.044 \times 10^{-17} \text{ J T}^{-1}. \quad (37)$$

3.6 Tipping the spin ensemble: the FID protocol

In steady state the spins are aligned with the local stable spin axis, taken here as $\vec{n}_0 \simeq \vec{e}_y$. The bunch magnetic moment is then static and points approximately along $\vec{e}_y$, as in Eq. (31). The cosine- θ pickup, which measures the local m_x component, therefore sees no FID signal before the spin ensemble is tipped.

To initiate a measurement, a dedicated spin-kick element tips the spin ensemble away from $\vec{e}_y$ by a small angle α into the ring plane. For the nominal arc geometry with $\vec{n}_0 \simeq \vec{e}_y$, a longitudinal magnetic field B_z provides a spin-kick axis perpendicular to $\vec{n}_0$. The tipping angle

is related to the integrated longitudinal field by the T-BMT equation [28],

$$\alpha = \frac{(1+G)}{B\rho} \int B_z dl, \quad (38)$$

where $B\rho = p/e$ is the magnetic rigidity and G is the proton anomalous magnetic moment (Table 3). The factor $(1+G)$ appears because the spin rotation in a longitudinal field combines the anomalous-moment contribution with the rotation along the velocity direction; see, e.g., [28, 34]. A longitudinal field is attractive because $B_z \parallel \vec{v}$ gives $\vec{v} \times \vec{B} = 0$ for the reference particle. It therefore produces a magnetic spin kick without a transverse Lorentz-force kick on the closed orbit.

After the pulse, the tipped in-plane component precesses freely at the spin-precession frequency introduced in Eq. (3). For $\nu_s = 0.5$, this component changes sign on every turn. The moment sampled by a fixed cosine- θ pickup after n turns is therefore

$$m_x^{\text{FID}}(n) = P N_p \mu_p \sin \alpha (-1)^n. \quad (39)$$

This freely precessing transverse signal is the free-induction decay (FID), in direct analogy with nuclear magnetic resonance (NMR).

The kick element must create the transverse spin component without significantly perturbing the orbit. A purely longitudinal magnetic kick satisfies this condition geometrically, but a conventional high-inductance solenoid may not be the most practical implementation for a fast, well-defined pulse. An alternative is a transverse RF spin-kick element operated in Wien-filter mode, where the electric and magnetic fields satisfy

$$\vec{E} + \vec{v} \times \vec{B} = 0 \quad (40)$$

for the reference orbit. This cancels the orbital Lorentz force while retaining the spin rotation. RF Wien filters for spin manipulation have been developed and operated in the COSY/JEDI program [19, 35]. Phase-locking such a transverse Wien-filter field to a longitudinal RF field, 90° out of phase, selects a single rotating field component and drives a coherent reversal of a bunch's stored polarization at $\nu_s = 1/2$; gated bunch by bunch, this provides per-bunch polarization reversal, which is developed in Appendix A.

For the kick to be treated as impulsive in the spin dynamics, its duration should be short compared with the spin-precession period,

$$T_s = \frac{1}{f_s} \simeq 25.6 \mu\text{s}. \quad (41)$$

A sub-microsecond to few-microsecond spin-kick element is therefore the natural target. Slower or resonant implementations are not excluded, but then the finite time profile of the kick must be included explicitly in the spin map rather than approximated as an instantaneous rotation.

The tipping angle α sets the size of the observed FID signal. If the spin ensemble were tipped and then simply left unrestored, the stored projection would be reduced by $1 - \cos \alpha \simeq \alpha^2/2$. In the intended tip-read-restore protocol, however, a restore pulse is applied after the readout. The residual operational polarization loss is therefore determined by the accuracy and reproducibility of the restore operation, not by the tipping angle alone.

If the local stable spin axis differs appreciably from $\vec{e}_y$, the spin-kick field must be oriented perpendicular to the actual local $\vec{n}_0(s)$. This is a placement and hardware-orientation constraint, not a change to the basic FID concept.

4 Spin-tune measurement: spectral search and phase lock

The first operational task of the SQUID system is not absolute polarimetry but the direct, on-line measurement of the ensemble spin tune ν_s and of the spin-tune spread σ_{ν_s} . Both must be

determined from the SQUID signal itself, before any phase-tracking or coherent-accumulation strategy of Sec. 6.8 can be invoked. The FID protocol of Sec. 3.6, which produces a freely precessing transverse magnetic moment after a controlled spin tip, supplies the required time-domain signal; what is needed in addition is a data-analysis strategy that can extract the precession frequency from this signal without prior knowledge of either the frequency itself or the coherence window.

The principal difficulty is that the SQUID polarimeter must commission itself in a machine state whose spin-coherence time is essentially unknown and likely short. A first measurement concept that assumed an observation time longer than the untuned coherence time would be circular: it would require the machine to be tuned for long coherence before the diagnostic capable of guiding that tuning is operational. The initial measurement strategy must therefore work in the short-coherence-time regime expected from conservative start-up estimates ($\sigma_{\nu_s} \sim 10^{-2}$ – 10^{-3} , corresponding to $\tau_{\text{coh}} \sim 0.2$ – 2 ms) and must extract both f_s and σ_{ν_s} from the SQUID signal alone. The technique introduced in this section meets this constraint with a noninvasive spectral search performed entirely with the SQUID polarimeter.

4.1 Spectral-search strategy

One measurement cycle consists of (i) a longitudinal kicker pulse, defined in Sec. 3.6, that tips the stored polarization by the angle $\alpha = 30$ mrad; (ii) a SQUID recording of the resulting FID for a duration $T_{\text{meas}} \leq \tau_{\text{coh}}^{\text{untuned}}$; (iii) a fast Fourier transform of the recorded record to produce a single power spectrum; and (iv) a closing restore pulse that returns the polarization to vertical. This cycle is repeated N times. The N single-record power spectra are then averaged incoherently.

The qualitative behaviour expected from this averaging is the content of Fig. 5. The underlying physical signal is a damped cosine at the precession frequency f_s , with envelope $\exp(-t/\tau_{\text{coh}})$ set by the spin-tune spread (panel a). In any single record this signal is buried in SQUID noise, the per-record signal-to-noise ratio at $\sigma_{\nu_s} = 10^{-3}$ is only $\text{SNR}_1 \approx 0.75$ (panel b), so the single-record power spectrum carries no unambiguous peak (panel c). The key feature of incoherent power-spectrum averaging is that the deterministic signal contribution at f_s adds in phase from record to record (because power is a positive quantity), while the random noise contributions reduce as $1/\sqrt{N}$ in the averaged spectrum. After enough records the peak at f_s therefore emerges above the noise floor at a significance that grows as $\sqrt{N} = \sqrt{T_{\text{tot}}/\tau_{\text{coh}}}$ (panel d). Importantly, this strategy requires no a priori knowledge of the precession phase between records; what is preserved across the average is the precession frequency, which is the quantity to be measured.

Two pieces of information are extracted from the averaged spectrum. The peak position determines the spin tune, $\nu_s = f_s/f_{\text{rev}}$, with statistical uncertainty $\delta\nu_s$ that decreases as $1/\sqrt{T_{\text{tot}}}$ once the peak is resolved. The peak width determines the spin-tune spread σ_{ν_s} through the relation $\tau_{\text{coh}} = 1/(2\pi f_{\text{rev}}\sigma_{\nu_s})$, with narrower peaks corresponding to better-tuned machine states. A single spectral search therefore returns simultaneously the spin tune that will be used to set up phase-locked operation (Sec. 2.4) and the coherence-time diagnostic that will guide subsequent lattice tuning.

4.2 Detection feasibility and operating-point numbers

Figure 6 extends the single-energy demonstration of Fig. 5 to both storage energies and quantifies the detection and precision scaling. Panels (a) and (b) show averaged spectra at injection (1.7 min of integration) and at flattop (27.1 min). The peak at f_s is resolved at both energies. The slower emergence at flattop is not a consequence of fewer total stored protons (RF bunch-splitting quadruples the bunch count and conserves the total integrated proton current) but of the SQUID pickup responding to the magnetic dipole moment of each *individual* bunch as it passes. Each flattop bunch carries one quarter of the protons of an injection bunch, hence one quarter of the per-bunch FID amplitude. Since detection significance scales linearly in this




Figure 5: Spectral search with the SQUID polarimeter at injection energy: the precession signal emerges from SQUID noise through incoherent power-spectrum averaging. All panels use the locked operating point $\alpha = 30$ mrad, $P = 0.5$, $\sigma_{\nu_s} = 10^{-3}$, yielding a coherence time $\tau_{\text{coh}} = 2.04$ ms and an expected precession frequency $f_s = 39.07$ kHz. (a) Clean reference FID over one record of duration $T_{\text{meas}} = \tau_{\text{coh}}$; the red trace is the underlying physical signal, the dashed gray curves mark its exponential envelope. (b) Same record plus realistic SQUID noise; the per-record signal-to-noise ratio $\text{SNR}_1 \approx 0.75$ leaves the signal entirely buried in noise. (c) Power spectrum of one such record; the underlying peak at f_s is indistinguishable from random spectral spikes. (d) Power spectrum averaged over N records; the peak at f_s emerges with N and reaches $\sim 14\sigma$ above the noise floor in $T_{\text{tot}} = N T_{\text{meas}} = 1.7$ min of integration, well inside a routine machine-development session. The detection significance grows as $\sqrt{N} = \sqrt{T_{\text{tot}}/\tau_{\text{coh}}}$.

amplitude and only as $\sqrt{T_{\text{tot}}}$ in time, the integration required for matched significance grows by a factor of about sixteen relative to injection. Panel (c) plots the universal $\sqrt{T_{\text{tot}}}$ scaling of the detection significance against both energies, with the 5σ and 10σ detection thresholds shown for orientation. At injection, the spectral peak is detected at 10σ after $T_{\text{tot}} \lesssim 1$ s; at flat-top, after $T_{\text{tot}} \sim 6$ s. Table 4 collects the corresponding numerical timing scenarios for three machine-tuning states characterized by σ_{ν_s} . The canonical value $\sigma_{\nu_s} = 10^{-3}$ reflects conservative start-up expectations from COSY-class lattices [13, 36]; $\sigma_{\nu_s} = 10^{-2}$ represents a poorly tuned commissioning state, while $\sigma_{\nu_s} = 10^{-4}$ represents the level projected after lattice optimization informed by an initial SQUID measurement (still two orders of magnitude above the best COSY deuteron benchmark of $\sigma_{\nu_s} \sim 2 \times 10^{-10}$ achieved with deliberate spin-coherence tuning [36]). The required integration times to reach $\delta\nu_s = 10^{-5}$ at injection range from 1.0 h in the poorly-tuned case down to 0.36 s in the ML/AI-tuned case. The corresponding integrated longitudinal kicker field, fixed by α and the magnetic rigidity, is $\int B_z dl = 0.84$ T m at injection and $\int B_z dl =$

Table 4: Spin-tune-search timing scenarios as a function of the assumed spin-tune spread σ_{ν_s} . All entries assume the locked operating point $\alpha = 30$ mrad, $P = 0.5$ (a conservative start-up polarization). The per-record duration is matched to the coherence time $T_{\text{meas}} = \tau_{\text{coh}} = 1/(2\pi f_{\text{rev}}\sigma_{\nu_s})$ set by the assumed σ_{ν_s} . SNR_1 is the spectral signal-to-noise ratio of the peak at f_s after one such record, computed from the SQUID polarimeter and beam parameters in Table 1. T_{tot} is the total integration time required to reach $\delta\nu_s = 10^{-5}$, corresponding to a 1% relative measurement of σ_{ν_s} in the canonical scenario ($\sigma_{\nu_s} = 10^{-3}$). The required integrated longitudinal kicker field is $\int B_z dl = 0.84$ T m at injection and $\int B_z dl = 9.85$ T m at flattop, independent of σ_{ν_s} .

σ_{ν_s}	τ_{coh}	SNR_1 (inj)	SNR_1 (flat)	T_{tot} (inj)	T_{tot} (flat)
10^{-2}	0.20 ms	0.237	0.059	1.0 h	16.1 h
10^{-3}	2.04 ms	0.750	0.187	36.2 s	9.7 min
10^{-4}	20.37 ms	2.372	0.593	0.36 s	5.8 s

9.85 T m at flattop.

4.3 Iterative improvement loop

Panel (d) of Fig. 6 makes explicit the operational consequence of measuring f_s and σ_{ν_s} simultaneously: for a fixed target precision $\delta\nu_s$, the required integration time scales as $\sigma_{\nu_s}^2$. A factor-of-ten reduction in the spin-tune spread therefore compresses the measurement time by two orders of magnitude. This is the qualitative advantage of noninvasive SQUID polarimetry over destructive carbon-target spin-tune measurement [13]: each measurement returns not only ν_s but also a fresh, non-destructive estimate of σ_{ν_s} from the peak linewidth, providing the diagnostic feedback needed to guide sextupole, chromaticity, and bunch-cooling adjustments without heating the beam. The next measurement is then both more precise at fixed time and faster at fixed precision. A realistic factor-of-ten reduction in σ_{ν_s} , obtainable from, e.g., ML/AI-guided lattice tuning, compresses an initial $\sim$ twenty-minute spectral search at $\sigma_{\nu_s} \sim 10^{-2}$ to a few-second update at $\sigma_{\nu_s} \sim 10^{-4}$ (Table 4), once the lattice has been optimized for long spin coherence.

4.4 Transition to phase-locked operation

Once the spectral search has located f_s to a small fraction of the coherence-broadened peak width, the SQUID reference oscillator is phase-locked to f_s [16] and the system transitions to a precision operating mode that reaches $\delta P/P = 1\%$ within seconds at injection and minutes at flattop. Two ingredients enable this transition; both are developed in the next two sections, and the present subsection sketches them in enough detail to explain the operational picture before the reader meets them in full.

The first is the spin echo, developed in Sec. 5. A phase-coherent π rephasing pulse applied at time $t = \tau$ after the initial tipping pulse reflects the phase fan in the transverse precession plane, so that the dephasing accumulated during $0 \leq t \leq \tau$ closes again at $t = 2\tau$. The echo extends the usable in-plane coherence window from the reversible $T_2^* = \tau_{\text{coh}}$ set by the spin-tune spread (Eq. 42) to the much longer irreversible T_2 (Eq. 43), so a single measurement record is no longer bounded by τ_{coh} . The full four-pulse cycle (tip, π rephase, echo readout, restore) is closed by a phase-coherent restore pulse that returns the tipped in-plane component to the stable spin axis with a per-cycle polarization cost of $\mathcal{O}(\alpha^2/\pi^2) \sim 10^{-4}$, so many cycles can be accumulated within a single fill with negligible cumulative depletion.

The second is the matched-filter coherent summation across all bunches, developed in Sec. 6.6. The injected spin-sign pattern and the deterministic lattice phase advance from bunch to bunch are both known, so each per-bunch contribution can be rotated into a common phase before




Figure 6: Operational analysis of the spin-tune spectral search. (a) Averaged SQUID power spectrum at injection ($N = 5 \times 10^4$ records, $T_{\text{tot}} = 1.7$ min): the peak at $f_s = 39.07$ kHz stands $\sim 15\sigma$ above the noise floor. (b) Same at flattop ($N = 8 \times 10^5$ records, $T_{\text{tot}} = 27.1$ min); RF bunch-splitting reduces the proton number per bunch by a factor of four, so the same detection significance requires a factor of sixteen more integration time than at injection, but the peak is still resolved within one hour of machine-development time. (c) Detection significance vs total integration time T_{tot} at $\alpha = 30$ mrad, $P = 0.5$, for both energies; the universal $\sqrt{T_{\text{tot}}}$ scaling shows that 10σ detection is reached in $T_{\text{tot}} \lesssim 1$ s at injection and $T_{\text{tot}} \sim 6$ s at flattop. Horizontal dotted lines mark 5σ and 10σ thresholds; vertical dotted lines indicate 1 min and 1 h time budgets. d) Spin-tune precision $\delta\nu_s$ at injection vs T_{tot} for three machine-tuning states characterized by the spin-tune spread σ_{ν_s} . The horizontal dotted line marks the representative target $\delta\nu_s = 10^{-5}$, corresponding to a 1% relative measurement of σ_{ν_s} in the canonical scenario ($\sigma_{\nu_s} = 10^{-3}$). Reducing σ_{ν_s} by one order of magnitude (e.g., through ML/AI-guided lattice tuning informed by an initial SQUID measurement) compresses the integration time required to reach this target by two orders of magnitude.

summation; the result is a per-record signal-to-noise ratio that scales as $\sqrt{N_{\text{fill}}}$ rather than the order-unity gain of an uncorrected sum. Combined with the phase-locked tip- π -echo-restore cycle, the relative statistical precision on the polarization magnitude P scales as $1/\sqrt{T}$ in the total integration time T , with the stored beam polarization preserved across the measurement. A slow feedback on the fitted f_s from rolling spectral windows tracks drift in the lattice, the closed orbit, and the snake matching during the fill [13], maintaining phase lock across hours of running.

A complementary precision tool for the coherent ring-magnetic-field environment, of interest to BSM observables such as charged-particle electric dipole moments, is the radio-frequency Wien filter operated at $\nu_s = 1/2$. Its physics, the corresponding waveguide implementation, and its connection to the spin-tune measurement established here are discussed in Appendix A.

5 Spin-echo rephasing and the noninvasive measurement cycle

The spectral search of Sec. 4 returns the ensemble spin tune ν_s from the location of the spectral peak, the spin-tune spread σ_{ν_s} from its width, and concludes by phase-locking the SQUID reference oscillator to the measured precession at $f_s = \nu_s f_{\text{rev}}$. With phase lock in hand, the next step is to extend the measurement window beyond the reversible-dephasing time $\tau_{\text{coh}} = (2\pi f_{\text{rev}} \sigma_{\nu_s})^{-1}$ set by the spin-tune spread. The instrument for this extension is a Hahn spin-echo sequence [37]: a phase-coherent π pulse applied at $t = \tau$ rephases the dephased ensemble at $t = 2\tau$, cancelling the reversible σ_{ν_s} -driven dephasing and exposing the irreversible decoherence due to processes that change individual spin tunes during the measurement.

Equally important, when the echo sequence is closed by a phase-coherent restore pulse the entire cycle (tip, π rephase, echo readout, restore) returns the beam polarization to the stable spin direction with negligible loss per cycle. This is the operational property that makes the SQUID polarimeter noninvasive: many measurement cycles can be repeated within the same fill without measurable depletion of the stored polarization. By contrast, the fast carbon-target polarimeters used at RHIC [38] sampled the stored beam destructively, with each measurement depleting beam intensity and degrading the emittance. Operationally this restricted RHIC pC polarimetry to a small number of measurements at the beginning, middle, and end of each store, rather than continuous monitoring throughout the fill.

The hardware requirement is no more demanding than what the spin-tune search of Sec. 4 already needs. The longitudinal kicker that delivers the tipping pulse also delivers, when synchronized to the SQUID reference oscillator, both the π rephasing pulse and the restore pulse. The same lattice element therefore performs all three roles, and the phase lock established at the end of Sec. 4 is the prerequisite that makes the timed sequence close.

5.1 Reversible and irreversible decoherence

Two physically distinct mechanisms reduce the coherent transverse magnetic moment of the stored ensemble. The first is reversible dephasing: particles with different spin tunes $\nu_{s,i} = \nu_s + \delta\nu_{s,i}$ accumulate transverse phase differences $\Delta\phi_i(t) = 2\pi f_{\text{rev}} \delta\nu_{s,i} t$ relative to the ensemble mean, so the coherent vector sum decays with the rms spin-tune spread σ_{ν_s} . The individual spin vectors are not lost; only their phase-coherent sum is reduced. This is the FID envelope of Sec. 4, characterized by the reversible coherence time

$$T_2^* = \tau_{\text{coh}} = \frac{1}{2\pi f_{\text{rev}} \sigma_{\nu_s}}. \quad (42)$$

At canonical start-up ($\sigma_{\nu_s} = 10^{-3}$), $T_2^* = 2.04$ ms (Table 4). Reverting the relative phases by an external operation can in principle recover the lost coherence, because the dephasing is deterministic in the individual spin frequencies.

The second mechanism is irreversible. Processes that change individual spin tunes during the measurement window, intra-bunch Coulomb scattering, slow drifts of the closed orbit and the snake matching, ionization of residual gas anywhere in the ring vacuum, and beam losses, add a random walk to each particle's transverse phase. These mechanisms cannot be undone by an external pulse and define the irreversible (homogeneous) coherence time

$$T_2 \gg T_2^*, \quad (43)$$

which sets the ultimate envelope of the polarimeter signal regardless of any rephasing strategy. The spin-echo sequence developed in this section is the instrument by which T_2 is measured separately from T_2^* .

5.2 Echo rephasing and the π -pulse axis




Figure 7: Functional layout of the tip- π -echo-restore sequence. The measurement station contains the longitudinal kicker (which delivers the tip, π , and restore pulses) and the SQUID pickup. The tipping pulse at $t = 0$ creates the transverse spin component; the π pulse at $t = \tau$ reflects the phase fan; the echo maximum is sampled near $t \simeq 2\tau$; and the restore pulse, fired after sampling, returns the polarization to the stored direction. All four pulses are synchronized to the SQUID reference oscillator phase-locked at f_s .

Figure 7 shows the functional layout of the four-pulse sequence around the ring: the longitudinal kicker delivers the tipping pulse at $t = 0$, the π rephasing pulse at $t = \tau$, and the closing restore pulse, while the SQUID pickup samples the echo maximum near $t \simeq 2\tau$. The remainder of this section develops the spin-dynamics content of each step in turn.

The spin-echo sequence applies a phase-coherent π rotation pulse at time $t = \tau$ after the initial tipping pulse, reflecting the phase fan in the transverse precession plane (Fig. 8). Particles that had advanced furthest in phase before the pulse are repositioned to lag the ensemble mean afterwards, while those that had lagged behind move ahead. Continued free precession then drives the ensemble toward rephasing, with the coherent transverse vector reaching an echo maximum at $t = 2\tau$.

The π pulse is applied about the longitudinal axis $\vec{e}_z$ (the beam direction). This axis lies in the transverse precession plane, perpendicular to the local stable spin axis $\vec{n}_0 \simeq \vec{e}_y$, and is therefore one of the two possible choices for an echo rephasing pulse (a rotation about $\vec{e}_x$ would also rephase). The longitudinal axis is the preferred choice because the tipping kicker of Sec. 3.6 already delivers rotations about $\vec{e}_z$: the same element provides the π pulse and the restore pulse without modification. In spin components,

$$R_z(\pi) : (S_x, S_y, S_z) \rightarrow (-S_x, -S_y, S_z), \quad (44)$$

so the transverse components (S_x, S_y) are reflected across $\vec{e}_z$ while the stored polarization along




Figure 8: Simplified rotating-frame phase-fan picture of the FID and spin-echo sequence. Each colored vector represents a representative particle or beam slice with a slightly different spin tune, and the black vector denotes the ensemble-averaged transverse component. (a) Immediately after the tipping pulse, the transverse phases are aligned. (b) During $0 < t < \tau$, the phases fan out because of the spin-tune spread. (c) The π -pulse at $t = \tau$ reflects the phase fan across the longitudinal axis $\vec{e}_z$. (d) During $\tau < t < 2\tau$, the reflected phases re-align and the echo envelope reaches its maximum near $t \simeq 2\tau$. (e) The restore pulse, applied after the echo signal has been sampled, returns the ensemble to the stored-spin direction.

$\vec{n}_0$ is reversed in sign. The S_y sign reversal is the operational price of using $\vec{e}_z$ as the rotation axis; it is undone by the restore pulse in Sec. 5.4 below.

The rephasing can be verified directly. Consider a particle with spin tune $\nu_{s,i} = \nu_s + \delta\nu_{s,i}$, angular precession rate $\Omega_i = 2\pi(\nu_s + \delta\nu_{s,i})f_{\text{rev}}$, and an initial tipping pulse $R_z(\alpha)$ applied at $t = 0$ to the stored state $(0, 1, 0)$. The post-tip state is $(-\sin \alpha, \cos \alpha, 0)$. Free precession about $\vec{n}_0 = \vec{e}_y$ to $t = \tau^-$ gives

$$(S_x, S_y, S_z)_{\tau^-} = (-\sin \alpha \cos \Omega_i \tau, \cos \alpha, \sin \alpha \sin \Omega_i \tau). \quad (45)$$

Application of $R_z(\pi)$ from Eq. (44) flips S_x and S_y :

$$(S_x, S_y, S_z)_{\tau^+} = (\sin \alpha \cos \Omega_i \tau, -\cos \alpha, \sin \alpha \sin \Omega_i \tau). \quad (46)$$

Free precession from τ to 2τ rotates the transverse components by an additional $\Omega_i \tau$. At $t = 2\tau$,

$$(S_x, S_y, S_z)_{2\tau} = (\sin \alpha, -\cos \alpha, 0), \quad (47)$$

independent of $\delta\nu_{s,i}$. The transverse phase fan that grew during $[0, \tau]$ closes exactly during $[\tau, 2\tau]$; every particle, regardless of its individual spin tune, lands in the same transverse orientation at $t = 2\tau$. This is the spin echo.

The irreversible decoherence on the time scale T_2 is not cancelled by this construction. Including it, the echo envelope decays from its tip-pulse value as

$$M_{\perp}(2\tau) = M_{\perp}(0) \exp\left(-\frac{2\tau}{T_2}\right), \quad (48)$$

with the T_2^* contribution having cancelled at $t = 2\tau$ by construction. The echo amplitude as a function of τ thus measures T_2 directly, separately from T_2^* .

5.3 Hardware reuse, phase-locked timing, and field budget

The tipping pulse of Sec. 3.6 produces a spin rotation about the longitudinal axis whose magnitude is set by the integrated longitudinal field,

$$\alpha = \frac{(1+G)}{B\rho} \int B_z dl. \quad (49)$$

Table 5: Per-pass integrated-field requirement for the synchronized π rephasing pulse as a function of the spin-tune spread σ_{ν_s} . The number of passes N is set to the safety-factor-2 maximum allowed by Eq. (51). For reference, the tipping kicker of Table 4 delivers $\int B_z dl = 0.84$ T m at injection and 9.85 T m at flattop in a single pass.

σ_{ν_s}	τ_{coh}	N (max)	$\int B_z dl _{\pi, \text{per-pass}}$ (inj)	$\int B_z dl _{\pi, \text{per-pass}}$ (flat)
10^{-2}	0.20 ms	8	11 T m	129 T m
10^{-3}	2.04 ms	80	1.1 T m	12.9 T m
10^{-4}	20.4 ms	796	0.111 T m	1.30 T m

The π rephasing pulse and the restore pulse rotate the spin about the same longitudinal axis, so the relation Eq. (49) applies to all three pulses, with only the rotation angle changing. From Table 4, the tipping rotation of $\alpha = 30$ mrad already requires $\int B_z dl = 0.84$ T m at injection. For the π rephasing pulse the rotation angle is larger by a factor $\pi/\alpha \simeq 105$, and the restore rotation $R_z(\pi - \alpha)$ demands almost the same. If delivered in a single passage of the beam through the kicker, the π pulse would therefore call for $\int B_z dl = 88$ T m at injection and 1032 T m at flattop, two orders of magnitude above the tipping requirement and well beyond what a single short pulse from any practical solenoid can deliver.

The way around the single-shot limit is to spread the π rotation over N machine turns, with the kicker fired synchronously with the SQUID reference oscillator at the precession frequency f_s . The synchronization is essential: at $\nu_s = 1/2$ the transverse spin projection in the (x, z) plane reverses sign on every machine turn, so a static field would give per-turn contributions that average to zero. A pulsed waveform at f_s instead produces per-turn contributions of the same sign, which accumulate constructively. Each pulse then delivers a per-pass rotation π/N requiring per-pass integrated field

$$\int B_z dl|_{\pi, \text{per-pass}} = \frac{1}{N} \cdot \frac{\pi B \rho}{1 + G}. \quad (50)$$

The maximum allowed N is set by the requirement that the entire π pulse complete in much less than the coherence time, so that particles with different $\delta\nu_{s,i}$ have not yet drifted significantly during the pulse itself. Taking a safety factor of two,

$$N \leq \frac{f_{\text{rev}} \tau_{\text{coh}}}{2} = \frac{1}{4\pi\sigma_{\nu_s}}. \quad (51)$$

Table 5 translates Eqs. (50)–(51) into per-pass integrated-field requirements for the three machine-tuning scenarios of Sec. 4 at both energies.

At injection, all three machine-tuning scenarios in Table 5 sit at or below the tipping-kicker requirement of Table 4; the same longitudinal kicker delivers tip, π rephase, and restore without modification. At flattop the same is true for the canonical and ML/AI-tuned cases ($\sigma_{\nu_s} \leq 10^{-3}$): the per-pass $\int B_z dl$ for the π pulse is between 1.3 T m and 13 T m, comparable to or below the flattop tipping requirement of 9.85 T m. Only the worst-case flattop scenario ($\sigma_{\nu_s} = 10^{-2}$) is a demanding condition, requiring an order of magnitude larger per-pass field than tipping (132 T m). This single design point motivates the spin-coherence improvement loop of Sec. 4.3 as a flattop commissioning priority: bringing σ_{ν_s} down to the canonical 10^{-3} level reduces the per-pass π -pulse field requirement by a factor of ten and places it comfortably within the existing hardware envelope.

Three phase-locked timings follow from the SQUID reference oscillator of Sec. 4. The tipping pulse, fired at $t = 0$, has no phase requirement: it acts on a fresh stored S_y . The π rephasing pulse must be fired at a precession phase that cleanly reflects (S_x, S_z) across $\vec{e}_z$, at a fixed offset from the SQUID reference, calibrated once per fill. The restore pulse must be timed so that the

spin is again at the precession phase it occupied immediately after the tipping pulse, modulo the S_y flip introduced by the π pulse. All three timings reduce operationally to fixed offsets from the SQUID reference oscillator.

5.4 Restore pulse and the closed measurement cycle

The state of the ensemble at the echo maximum, Eq. (47), is $(\sin \alpha, -\cos \alpha, 0)$. After the echo signal has been sampled by the SQUID, a phase-coherent restore pulse rotates the ensemble back to the stored-spin direction. The required rotation is $R_z(\pi - \alpha)$:

$$R_z(\pi - \alpha) (\sin \alpha, -\cos \alpha, 0) = (0, 1, 0), \quad (52)$$

which simultaneously undoes the S_y sign reversal introduced by the π pulse and the small tipping rotation α . The integrated-field requirement is therefore $(1 - \alpha/\pi) \simeq 99\%$ of the π -pulse field requirement already established in Eqs. (49)–(50). Figure 8(e) shows the action of the restore pulse in the rotating-frame phase-fan picture.

The full four-pulse cycle (tip, π rephase, echo readout, restore) thus returns the beam polarization to within $\mathcal{O}(\alpha^2/\pi^2) \sim 10^{-4}$ of its starting value per cycle, modulo irreversible T_2 loss accumulated during the 2τ free-evolution interval. The cycle can be repeated continuously throughout a fill at a duty cycle determined by τ . This is the operational property that makes the SQUID polarimeter noninvasive: the stored polarization is preserved across the measurement, and many cycles can be averaged to improve the statistical precision on the echo amplitude and on the spin tune extracted from the SQUID signal between pulses.

Cumulative drift from imperfect calibration of the restore angle can be eliminated entirely by a Carr–Purcell–Meiboom–Gill (CPMG) sequence [39, 40], in which a second π pulse at $t = 3\tau$ produces a second echo at $t = 4\tau$, after which the spin state coincides with the configuration immediately following the tipping pulse, so that the stored polarization sign is restored automatically by the second π flip. A single $R_z(-\alpha)$ tip then closes the cycle, with the restore-pulse field requirement reduced to that of the original tipping pulse.

The closed cycle enables a qualitatively new operating mode at the HSR. With a fill lasting several hours and the tip- π -echo-restore cycle returning the beam polarization to its starting value each time, the spin-precession frequency f_s , the spin-tune spread σ_{ν_s} , and the in-plane beam polarization can be monitored essentially continuously throughout the entire ~ 8 -hour store, with negligible cumulative impact on the beam intensity or polarization. A slow feedback on the fitted f_s from rolling spectral windows tracks the spin tune as it drifts during the store, maintaining phase lock across hours of running. The drift itself carries physics information: the spin tune is acutely sensitive to lattice imperfections, closed-orbit distortions, and snake-matching errors, as demonstrated in dedicated spin-tune mapping experiments at COSY [15], so continuous spin-tune monitoring with the SQUID polarimeter functions in parallel as a high-precision diagnostic of HSR magnetic imperfections and operational stability. The precision achievable by this class of free-precession measurement is high: at COSY the JEDI free-precession method determined the spin tune to order 10^{-8} in 2.6 s and to 1×10^{-10} over a continuous 100 s cycle [13], among the most precisely measured quantities in any accelerator, and spin-tune mapping fixed the angular orientation of the stable spin axis $\vec{n}_0$ to better than $2.8 \mu\text{rad}$ [15]. These benchmarks were reached in a regime far removed from HSR operating conditions: in-plane spin-coherence times near 1000 s obtained through dedicated bunching, electron cooling, and sextupole correction [36], with the stored beam current deliberately limited to below 10^9 deuterons per fill to suppress collective effects [14]. They indicate the ultimate reach of the free-precession method rather than the precision expected at HSR start-up, where the coherence time is initially of order milliseconds and the stored proton intensity is far higher. Continuous polarimetry of this kind is beyond the reach of the polarimetry methods deployed at RHIC and required at the EIC. The fast carbon-target (pC) polarimeter [38] samples the stored beam destructively, with each measurement depleting beam intensity and degrading the




Figure 9: Continuous SQUID readout for the full tip- π -echo-restore sequence. The upper panel shows the coherent transverse-spin envelope: after the initial FID decay on the reversible time scale $T_2^* = \tau_{\text{coh}}$, the π pulse at $t = \tau$ reverses the phase fan and produces an echo maximum near $t \simeq 2\tau$, reduced from the tip-pulse value by $\exp(-2\tau/T_2)$ from irreversible decoherence on the time scale T_2 . The lower panel zooms on the SQUID signal near the echo maximum: the precession at $f_s = \nu_s f_{\text{rev}} \simeq 39.1 \text{ kHz}$ continues uninterrupted throughout the sequence, and the echo restores the slowly varying coherent envelope, not the underlying precession.

emittance, restricting pC polarimetry operationally to a small number of measurements at the beginning, middle, and end of each store rather than continuous monitoring. The noninvasive atomic-hydrogen-jet (HJET) polarimeter, also operated at RHIC and required at the EIC for absolute proton beam polarization calibration, removes the destructive aspect but provides only marginal statistical precision per unit time relative to what the SQUID concept offers at HSR conditions [1].

5.5 The echo amplitude τ -scan and what is measured

The deliverable of the spin-echo sequence is the echo amplitude as a function of the operator-chosen delay τ . From Eq. (48),

$$\frac{M_{\perp}(2\tau)}{M_{\perp}(0)} = \exp\left(-\frac{2\tau}{T_2}\right), \quad (53)$$

so a logarithmic-slope fit of the SQUID echo amplitude versus τ extracts T_2 directly. The FID linewidth measured from the same SQUID data within each cycle continues to yield T_2^* via Eq. (42) as in Sec. 4. The two coherence times can therefore be measured simultaneously in the same fill, and the ratio T_2/T_2^* quantifies how much of the apparent τ_{coh} is recoverable dephasing versus genuine irreversible loss. Figure 9 illustrates the appearance of the continuous SQUID signal across one such cycle.

The HSR-specific contributions to T_2 include intra-bunch Coulomb scattering, slow drifts of the closed orbit and the snake matching, residual-gas scattering throughout the ring, and beam losses. Quantitative estimates of these contributions require dedicated HSR spin tracking and are beyond the scope of this section. The operational point is that T_2 is a measurable, noninvasive diagnostic of these processes via Eq. (53), complementing the σ_{ν_s} feedback used by the lattice-tuning loop of Sec. 4.3.

6 Magnetic flux signal and matched-filter reconstruction




Figure 10: Magnetic-field pattern of a transverse magnetic dipole before and after a Lorentz boost. Both panels are shown in the y - z plane, with y vertical and z horizontal. Left: magnetic field lines of a dipole moment $\vec{m} = m_y \vec{e}_y$, shown in red, at rest. Right: the corresponding Lorentz-boosted field pattern for the same dipole boosted along $+\vec{e}_z$, shown for $\beta = 0.98$. The boost compresses the field pattern in the longitudinal z direction, so that a fixed pickup samples the field over a short time interval set by the bunch length.

6.1 Relativistic field of a moving transverse magnetic dipole

The SQUID pickup coil is placed at position $r_{\text{coil}} \vec{e}_x$ (distance r_{coil} from the beam axis along $\vec{e}_x$), with its effective area normal along $\vec{e}_x$, outside the beam pipe. For a magnetic dipole $\vec{m} = m_x \vec{e}_x$ moving along $\vec{e}_z$ at $v \approx c$, the Lorentz transformation of the rest-frame magnetic field gives, at the moment of closest approach of the bunch centre to the coil (i.e. when the bunch centre

passes the transverse plane $z = 0$ containing the coil)

$$B_x^{\text{lab}} = + \frac{\mu_0}{4\pi} \frac{2m_x}{r_{\text{coil}}^3}. \quad (54)$$

At the peak this transverse field is in fact enhanced by γ : with $\vec{E}' = 0$ in the bunch rest frame, the Lorentz transformation of the field gives $\vec{B}_{\perp}^{\text{lab}} = \gamma \vec{B}'_{\perp}$ pointwise, while the transverse separation r_{coil} is invariant. What is boost-invariant is not the peak amplitude but the longitudinal field integral: the amplitude factor γ is compensated by the $1/\gamma$ longitudinal compression of the field pattern (Fig. 10). Because the bunch is much longer than this Lorentz-contracted profile ($\sigma_L \gg r_{\text{coil}}/\gamma$), and because the matched filter of Sec. 6.6 responds to the single-pass flux integral $\Phi_1(\omega_s)$ with $\omega_s \sigma_t \ll 1$ rather than to the instantaneous peak, the relevant flux template is this boost-invariant integral. Equation (54) may therefore be used with $\gamma = 1$, leaving the per-bunch flux of Sec. 6.2 and all downstream sensitivity estimates unchanged.

The bunch has a lab-frame rms length σ_L (listed in Table 1), related to the rms temporal width by $\sigma_t = \sigma_L/(\beta c)$; both are already lab-frame quantities and require no Lorentz correction. The time for the bunch centre to traverse the coil is negligible; what matters is the duration of the magnetic pulse at the coil, which is set by the bunch length divided by c :

$$\Delta t \approx 2\sigma_t, \quad (55)$$

where the factor of 2 accounts for the full $\pm 1\sigma$ passage time and σ_t is the lab-frame rms temporal width from Table 1. At EIC flattop ($\sigma_t = 0.200$ ns): $\Delta t \approx 0.40$ ns. The magnetic pulse is far too brief for time-domain detection; the signal is therefore extracted in the frequency domain via the matched-filter coherent summation described in Sec. 6.6.

6.2 Flux per bunch pass

The single-turn flux at the pickup for a bunch carrying transverse magnetic moment m_x is

$$\Phi_{\text{bunch}} = \frac{\mu_0}{4\pi} \frac{|m_x|}{r_{\text{coil}}^3} A f_{\text{geom}}, \quad (56)$$

with r_{coil} , A , and f_{geom} from Table 8. The longitudinal form factor $f_{\text{geom}} \sim 1$ –3 accounts for the finite saddle extent along $\vec{e}_z$ relative to the dipole field profile; it is the radial dipole field integrated over the saddle winding, normalized to the on-axis $1/r_{\text{coil}}^3$ scale, so the point value of Eq. (54) sets only this scale and the flux actually linked by the coil is carried by f_{geom} .

For unit polarization ($P = 1$) and full transverse projection ($\sin \alpha = 1$), the bunch magnetic moment is

$$|m_x| = N_p \mu_p, \quad (57)$$

with N_p from Table 1 and $\mu_p = 1.4106 \times 10^{-26}$ J T $^{-1}$ the proton magnetic moment. Evaluating Eq. (56),

$$\Phi_{\text{bunch}}^{\text{inj}} \approx 17.6 \mu\Phi_0, \quad \Phi_{\text{bunch}}^{\text{flat}} \approx 4.41 \mu\Phi_0, \quad (58)$$

where $\Phi_0 = h/(2e) = 2.07 \times 10^{-15}$ Wb is the flux quantum. These values represent the flux per single saddle turn that would be delivered by a fully polarized ensemble with the spin vector lying entirely in the transverse plane, and serve as the geometric flux template for the matched-filter sensitivity that follows.

The pickup is wound as a multi-turn cosine- θ saddle connected to the SQUID input through a superconducting flux transformer. With winding multiplicity N_{turns} from Table 8 and flux-transformer coupling efficiency η from Table 6, the flux delivered to the SQUID input per bunch passage is

$$\Phi_{\text{pickup}} = N_{\text{turns}} \eta \Phi_{\text{bunch}}. \quad (59)$$

For unit transverse polarization,

$$\Phi_{\text{pickup}}^{\text{inj}} \approx 1236 \mu\Phi_0, \quad \Phi_{\text{pickup}}^{\text{flat}} \approx 309 \mu\Phi_0. \quad (60)$$

The operating polarization P and tipping angle α (Table 1) enter as a single transverse-polarization factor $P_{\perp} = P \sin \alpha$ at the matched-filter sensitivity stage in Sec. 6.8, not in the per-bunch flux template here. The SQUID array architecture with N_{squids} independent channels (Table 6) contributes an additional factor $\sqrt{N_{\text{squids}}}$ to the matched-filter signal-to-noise ratio, also derived in Sec. 6.8.

Table 6: SQUID readout parameters. Flux noise is the conservative published value for the Magnicon-class two-stage current sensor at 4.2 K read out by the XXF-1 electronics [41], anchoring the canonical 4 K operating temperature assumed throughout this paper.

Parameter	Symbol	Unit	Value
Flux-transformer coupling efficiency	η	1	0.7
SQUID array channels	N_{squids}	1	4
SQUID white flux noise	$S_{\Phi}^{1/2}$	$\mu\Phi_0/\sqrt{\text{Hz}}$	0.4
SQUID operating temperature	T	K	4

The SQUID sensors and pickup coils are operated at a base temperature of $T = 4$ K, anchored to the published noise specification of the Magnicon-class two-stage current sensor at that temperature [41]. The cold mass is contained in a vacuum envelope with two beam-tube apertures of 100 mm diameter (combined aperture area $\approx 157 \text{ cm}^2$) for the upstream and downstream beam pipes. On the beam axis, the transverse dipole field of a single polarized bunch falls as $1/z^3$ outside the active pickup region, dropping to $\sim 10^{-19} \text{ T}$ at $z = 10 \text{ cm}$. For comparison, the single-bunch SQUID noise floor on a 40-mm-radius single-turn pickup is $\sim S_{\Phi}^{1/2} \sqrt{\text{BW}} / (\pi r_{\text{coil}}^2) \sim 10^{-14} \text{ T}$ for the per-bunch bandwidth $\text{BW} \sim 1/\sigma_t \sim 1 \text{ GHz}$ at injection, i.e. five orders of magnitude above the on-axis dipole field at $z = 10 \text{ cm}$; the ceramic former that carries the pickup coils can therefore be extended axially upstream and downstream of the active region at no signal cost. We assume such an extension is configured as a graded thermal shield, with the warm end at $\sim 300 \text{ K}$ at the beam-tube interface, an intermediate stage at $\sim 40 \text{ K}$, and the active region at 4 K , so that the room-temperature radiative load through the apertures is intercepted before reaching the active pickup. With this geometry the residual load at the 4 K stage is well within the $\sim 1 \text{ W}$ capacity of a standard pulse-tube cryocooler. Operation at sub-K temperatures with a multi-stage refrigerator would further reduce the flux noise but is not assumed for the baseline performance estimates of this paper.

6.3 Bunch spin pattern at injection

The bunch spin pattern at the EIC HSR has not yet been settled. Two operationally realistic options are the RHIC-style alternating pattern $s_j = (-1)^j$ for $j = 0, \dots, 289$, which is the canonical fill at RHIC for systematic cancellation of spin asymmetry errors, and a random pattern, which may be preferred at the EIC for added systematic robustness in spin-correlation measurements. Whatever pattern is loaded at injection is not modified during the store. At flattop, each parent bunch j produces four daughters at positions $4j, \dots, 4j + 3$, all inheriting the parent sign s_j . The matched-filter formulation developed below in Sec. 6.6 accommodates any known spin-sign pattern $\{s_j\}$ (Eq. (70)); the alternating pattern is used throughout this paper as the worked example because it is concrete, operationally familiar from RHIC, and is the most demanding case for an uncorrected sum (Sec. 6.5), providing a worst-case anchor for the matched-filter gain.

The important point is not the visual pattern itself, but the cancellation it produces in an unweighted sum. For the alternating injection pattern, a naive sum over bunches gives only

$$C_{\text{naive}} = \sum_{j=0}^{N_{\text{fill}}-1} s_j \cos \psi_j \simeq 1 \quad (\text{injection}), \quad (61)$$

and, after fourfold RF splitting,

$$C_{\text{naive}} \simeq 4 \quad (\text{flattop}). \quad (62)$$

Thus an every-other-turn or uncorrected bunch sum discards essentially the entire multi-bunch gain. The matched filter avoids this cancellation by applying the known spin sign and lattice phase correction to each bunch before summation.

6.4 Lattice spin phase offsets

The bunch index also carries a spin-phase offset. Bunch j is displaced from bunch 0 by the fraction j/N_{fill} of one ring circumference. Since the spin advances by $2\pi\nu_s$ per full revolution, the same longitudinal displacement corresponds to the spin phase advance

$$\psi_j = 2\pi\nu_s \frac{j}{N_{\text{fill}}}. \quad (63)$$

For the HSR operating point $\nu_s = 1/2$, this becomes

$$\psi_j = \frac{\pi j}{N_{\text{fill}}}, \quad (64)$$

where $N_{\text{fill}} = 290$ at injection and $N_{\text{fill}} = 1160$ at flattop.

This phase is not an additional unknown spin rotation. It is the deterministic phase offset that follows from the position of bunch j around the ring. Even if all bunches had the same polarization magnitude, their signals at a fixed pickup would not arrive with identical spin phase, because they sample different azimuthal positions of the same spin-precession wave around the ring.

The signal from bunch j on turn n can therefore be written as

$$x_{j,n} = \Phi_{\text{pickup}} s_j P_j \sin \alpha \cos(\pi n + \psi_j) + \text{noise}_{j,n}. \quad (65)$$

Here Φ_{pickup} is the per-bunch flux at the SQUID input for unit polarization and full transverse projection [Eq. (60)], $s_j = \pm 1$ is the injected spin-sign pattern, P_j is the polarization magnitude of bunch j , $\sin \alpha$ is the transverse projection set by the operating tipping angle α (Table 1), and the factor $\cos(\pi n + \psi_j)$ describes the spin-precession phase at the pickup. The term πn gives the turn-by-turn sign reversal from $\nu_s = 1/2$, while ψ_j is the fixed bunch-dependent phase offset.

The matched filter uses both known pieces of information, the injected sign s_j and the deterministic phase ψ_j , to rotate each bunch contribution into a common phase before summation. This correction chain is illustrated schematically in Fig. 11.

6.5 Naive every-other-turn summation

Summing all bunch signals on even turns yields an effective pattern-weighted sum $C_{\text{naive}} = \sum_j s_j \cos \psi_j$. For $s_j = (-1)^j$ with $\psi_j = \pi j/N_{\text{fill}}$ this is a geometric series evaluating to $C_{\text{naive}} = 1$: the alternating pattern nearly cancels, leaving only one equivalent bunch at injection and four at flattop.

6.6 Matched filter: phase-corrected coherent summation

The purpose of the matched filter is to remove the two known bunch-dependent factors before the bunches are summed: the injected spin sign s_j and the lattice spin-phase offset ψ_j . We therefore define the complex correction weight

$$w_j = s_j e^{-i\psi_j} = s_j (\cos \psi_j - i \sin \psi_j). \quad (66)$$

The real and imaginary parts of this weight are

$$\text{Re } w_j = s_j \cos \psi_j, \quad \text{Im } w_j = -s_j \sin \psi_j. \quad (67)$$

These are the two correction components shown in the middle panel of Fig. 11.

The phase-corrected bunch sum on turn n is then

$$\tilde{S}_n = \sum_{j=0}^{N_{\text{fill}}-1} w_j x_{j,n} = \sum_{j=0}^{N_{\text{fill}}-1} s_j e^{-i\psi_j} x_{j,n}. \quad (68)$$

Substituting Eq. (65), and taking $P_j = P$ for the coherent-sum estimate, gives

$$\tilde{S}_n = \Phi_{\text{pickup}} P_{\perp} \sum_{j=0}^{N_{\text{fill}}-1} s_j e^{-i\psi_j} s_j \cos(\pi n + \psi_j) + \tilde{n}_n, \quad (69)$$

where $\tilde{n}_n$ denotes the filtered noise contribution. Using $s_j^2 = 1$ and writing the cosine as the real part of a complex exponential, the coherent contribution becomes

$$\text{Re } \tilde{S}_n = \Phi_{\text{pickup}} P_{\perp} (-1)^n \sum_{j=0}^{N_{\text{fill}}-1} 1 = \Phi_{\text{pickup}} P_{\perp} (-1)^n N_{\text{fill}}. \quad (70)$$

Thus the matched filter rotates all bunch contributions into the same phase before summation. The uncorrected partial sum cancels for the alternating pattern, while the corrected sum grows coherently, as illustrated in the bottom panel of Fig. 11.

The optimal pattern sum is therefore

$$C_{\text{opt}} = N_{\text{fill}} = 290 \text{ (injection)}, \quad 1160 \text{ (flattop)}. \quad (71)$$

This result is independent of the particular spin-sign pattern $\{s_j\}$, provided that the pattern is known and included in w_j .

The matched-filter output of Eq. (70) is a turn-by-turn signal at the spin-precession frequency $f_s = \nu_s f_{\text{rev}}$ defined in Eq. (3), which evaluates to 39.07 kHz at injection and 39.10 kHz at flattop for $\nu_s = 1/2$.

A real-time implementation of the matched-filter sum, Eq. (68), does not require sampling each bunch individually and applying the weights in software. The general firmware implementation mixes the SQUID time-stream with a digital per-bunch reference encoding the known pattern $\{s_j\}$ and the lattice phase offsets ψ_j , generated by a digital signal processor locked to the machine RF clock. For the alternating-pattern worked example $s_j = (-1)^j$ this digital reference reduces to a single sinusoidal oscillator at the bunch-pattern alternation rate $f_{\text{ref}} = N_{\text{fill}} f_s$, where f_s is the spin-precession frequency, defined in Eq. (3)

$$f_{\text{ref}}^{\text{inj}} = N_{\text{fill}}^{\text{inj}} f_s^{\text{inj}} = 290 \times 39.07 \text{ kHz} = 11.33 \text{ MHz}, \quad (72)$$

$$f_{\text{ref}}^{\text{flat}} = N_{\text{fill}}^{\text{flat}} f_s^{\text{flat}} = 1160 \times 39.10 \text{ kHz} = 45.35 \text{ MHz}. \quad (73)$$

The relation $f_{\text{ref}} = N_{\text{fill}} f_s$ makes the physical content of the alternating-pattern case explicit: f_{ref} is N_{fill} times the spin-precession frequency, equivalently one half of the bunch-passage rate




Figure 11: Schematic illustration of the matched-filter correction for a representative subset of bunches. Top: raw bunch samples $x_{j,n}$ contain the injected spin sign s_j , the lattice spin-phase offset ψ_j , and the turn-by-turn factor $(-1)^n$ from $\nu_s = 1/2$ [Eq. (65)]. Middle: real and imaginary parts of the known per-bunch correction factor, as defined in Eqs. (66) and (67). Bottom: running sum over the displayed bunch subset. The uncorrected partial sum corresponds to the naive pattern-weighted sum C_{naive} , which remains small because adjacent bunch contributions cancel [Sec. 6.5]. The phase-corrected sum corresponds to $\text{Re}[\tilde{S}_n]$ after applying the matched-filter weights [Eq. (68)], and grows coherently as in Eq. (70). The full analysis applies the same correction to all N_{fill} bunches and then accumulates over turns.

$N_{\text{fill}} f_{\text{rev}}$ at $\nu_s = 1/2$. It is the frequency at which the alternating pattern $s_j = (-1)^j$ deposits a spectral line in the SQUID time-stream. The polarization signal appears not as a line at f_{ref} itself but as a pair of sidebands at $f_{\text{ref}} \pm f_s$ around it. Mixing the SQUID output with the sinusoidal reference at f_{ref} and applying a low-pass filter at a bandwidth slightly above f_s demodulates these sidebands down to baseband, leaving a complex oscillator at $\pm f_s$ that carries the bunch-pattern-corrected polarization signal. For a non-alternating pattern (e.g. a random pattern) the spin signal is distributed over multiple spectral lines determined by the Fourier content of $\{s_j\}$, and the firmware reference is the corresponding digital per-bunch sequence rather than a single sinusoid; the matched-filter principle and the per-bunch phase rotation by $e^{-i\psi_j}$ are unchanged.

The lattice spin-phase offsets ψ_j enter as a fixed per-bunch phase that the firmware-side matched filter applies as a static phase rotation to the demodulated samples; this is the hardware analog of the complex weight w_j in Eq. (66). Phase-locking the reference oscillator to the RF clock ensures that the bunch timing, the spin-sign pattern, and the phase correction share the same temporal reference, so that the matched filter remains coherent over the full integration time.

6.7 In-plane coherence time and spin tune spread

The matched-filter sum assumes that the transverse spin phase remains coherent over the integration time. The synchronous particle has $\nu_s = 0.5$, but off-momentum particles and particles at different betatron amplitudes acquire spin-tune offsets,

$$\nu_{s,i} = \nu_{s,0} + \delta\nu_{s,i}, \quad (74)$$

with $\nu_{s,0} = 1/2$ for the ideal HSR spin tune. The rms spread σ_{ν_s} of these offsets sets the FID envelope decay rate, with characteristic coherence time given by Eq. (42).

The working value adopted throughout this analysis is the effective rms spin-tune spread $\sigma_{\nu_s}^{\text{eff}} = 10^{-3}$ (Table 3), which corresponds to

$$\tau_{\text{coh}}^{\text{eff}} \simeq 2.04 \text{ ms}. \quad (75)$$

This is the value used for the matched-filter sensitivity estimate in Sec. 6.8 and for the multi-cycle accumulation that follows. Two larger reference scales are tabulated for context: the DLC ADST scale $\sigma_{\nu_s}^{\text{DLC}} \lesssim 10^{-2}$ (Table 3) gives $\tau_{\text{coh}} \simeq 0.20 \text{ ms}$, and the LC ADST scale $\sigma_{\nu_s}^{\text{LC}} \simeq 2 \times 10^{-2}$ gives $\tau_{\text{coh}} \simeq 0.10 \text{ ms}$. These scales are used in Sec. 4 for sizing the initial longitudinal kicker, not for the matched-filter sensitivity estimate.

At the operating tipping angle α (Table 1), the matched-filter integration time required to reach $\delta P/P = 1\%$ exceeds $\tau_{\text{coh}}^{\text{eff}}$. The measurement therefore proceeds as a sequence of N tip-read-restore cycles, each within one coherence window, and is averaged statistically; the corresponding cycle counts and integration times are given in Sec. 6.8. For the first spin-tune determination, before the spin-coherence environment has been characterized, a larger tipping angle and a single-record (short-window) detection are used, as discussed in Sec. 4.

The corresponding sideband linewidth around the spin-precession frequency f_s [Eq. (3)] is set by the spin-tune spread itself,

$$\sigma_f \simeq f_{\text{rev}} \sigma_{\nu_s}, \quad (76)$$

which evaluates to $\sigma_f^{\text{eff}} \simeq 78 \text{ Hz}$ for the working value $\sigma_{\nu_s}^{\text{eff}} = 10^{-3}$, and to 0.78 kHz and 1.56 kHz for the DLC and LC reference scales, respectively. Measuring this linewidth provides a direct noninvasive diagnostic of the spin-tune spread and of subsequent machine tuning.

6.8 Coherent integration

The SQUID readout measures a beam-synchronous sequence of single-pass flux waveforms. Let $\Phi_1(t)$ denote the flux waveform produced at the SQUID input by one bunch passage for unit transverse polarization. This waveform includes the magnetic coupling to the pickup, the finite bunch length, the pickup and flux-transformer response, and the SQUID/electronics transfer function. Its Fourier component at angular frequency ω is

$$\tilde{\Phi}_1(\omega) = \int_{-\infty}^{+\infty} \Phi_1(t) e^{-i\omega t} dt. \quad (77)$$

The quantity $\tilde{\Phi}_1(\omega)$ has units of flux times time.

For the transverse FID measurement, the relevant frequency is the spin-precession angular frequency

$$\omega_s = 2\pi\nu_s f_{\text{rev}}. \quad (78)$$

The bunch train converts the single-pass response into a coherent spectral line at this spin sideband. After applying the known injected spin sign and the lattice phase correction described in Sec. 6.6, the coherent line amplitude is

$$\Phi_{\text{line}} = f_{\text{rev}} C_{\text{opt}} P_{\perp} \left| \tilde{\Phi}_1(\omega_s) \right|, \quad (79)$$

where $P_{\perp} = P \sin \alpha$ is the transverse polarization component, with P the beam polarization and α the operating tipping angle (Table 1), and $C_{\text{opt}} = N_{\text{fill}}$ for the matched-filter sum.

For white SQUID flux noise with amplitude spectral density $S_{\Phi}^{1/2}$, coherent integration for a time T gives

$$\text{SNR}(T) = \frac{\Phi_{\text{line}}}{S_{\Phi}^{1/2}} \sqrt{T} = \frac{f_{\text{rev}} C_{\text{opt}} P_{\perp} |\tilde{\Phi}_1(\omega_s)|}{S_{\Phi}^{1/2}} \sqrt{T}. \quad (80)$$

This expression is the normalization used for the numerical sensitivity estimates. The peak single-pass flux scale in Sec. 6.2 fixes the magnetic coupling strength, while the matched-filter sensitivity is set by the spectral response $|\tilde{\Phi}_1(\omega_s)|$ of the full pickup and readout chain.

An alternative time-domain template form of the matched-filter signal-to-noise ratio exposes the hardware and ensemble inputs separately and is the form used to evaluate the numerical sensitivity in Sec. 6.9. For a Gaussian bunch profile $\Phi_1(t) = \Phi_{\text{pickup}} \exp(-t^2/2\sigma_t^2)$ with peak Φ_{pickup} at the SQUID input [Eq. (60)] and rms width σ_t (Table 1), define the matched-filter effective pulse width

$$\tau_h = \sqrt{\pi} \sigma_t, \quad (81)$$

chosen so that $\int |\Phi_1(t)|^2 dt = \Phi_{\text{pickup}}^2 \tau_h$. Accounting for an array of N_{squids} independent SQUID channels reading the same pickup (Table 6), the matched-filter signal-to-noise ratio after coherent integration over time T is

$$\text{SNR}(T) = K_{\text{template}} P_{\perp} \sqrt{T}, \quad (82)$$

with

$$K_{\text{template}} = \frac{\Phi_{\text{pickup}}}{S_{\Phi}^{1/2}} \sqrt{N_{\text{fill}} f_{\text{rev}} \tau_h N_{\text{squids}}}, \quad (83)$$

where N_{fill} and f_{rev} are the fill bunch count and revolution frequency (Tables 1 and 3), and $S_{\Phi}^{1/2}$ is the SQUID white flux noise (Table 6).

The template form factorizes the matched-filter sensitivity K_{template} into four physical contributions: $\Phi_{\text{pickup}}/S_{\Phi}^{1/2}$ is the single-channel flux-to-noise ratio set by the pickup geometry and the SQUID hardware; $\sqrt{N_{\text{fill}} f_{\text{rev}}}$ is the statistical gain from $N_{\text{fill}} f_{\text{rev}}$ bunch passages per unit time; $\sqrt{\tau_h}$ is the matched-filter window per pulse, set by the bunch length; and $\sqrt{N_{\text{squids}}}$ is the parallel-channel gain from the SQUID array. The transverse polarization factor $P_{\perp} = P \sin \alpha$ multiplies this template linearly, separating the operating-point choice of polarization and tipping angle from the hardware-fixed sensitivity scale K_{template} .

6.9 Numerical sensitivity estimates

Inserting the values from Tables 1, 3, 8, and 6 into Eq. (83) gives the matched-filter sensitivity at the two operating energies,

$$K_{\text{template}}^{\text{inj}} \approx 1108 \sqrt{\text{s}}, \quad (84)$$

$$K_{\text{template}}^{\text{flat}} \approx 277 \sqrt{\text{s}}. \quad (85)$$

The injection sensitivity exceeds the flattop sensitivity by a factor of four: the four times larger protons-per-bunch $N_p^{\text{inj}}/N_p^{\text{flat}} = 4$ enters Φ_{pickup} linearly, while the four times larger fill count $N_{\text{fill}}^{\text{flat}}/N_{\text{fill}}^{\text{inj}} = 4$ enters K_{template} only as $\sqrt{N_{\text{fill}}}$. The remaining factor of two comes from the matched-filter window $\sqrt{\tau_h} \propto \sqrt{\sigma_t}$: the injection bunch is four times longer ($\sigma_t = 0.801 \text{ ns}$ versus 0.200 ns), so $\sqrt{\tau_h^{\text{inj}}/\tau_h^{\text{flat}}} = 2$.

At the operating tipping angle $\alpha = 30 \text{ mrad}$ and beam polarization $P = 0.7$ (Table 1), the transverse polarization factor is $P_{\perp} = P \sin \alpha \approx 0.021$. The integration time required to reach

$\delta P/P = 1\%$ (SNR = 100) follows from Eq. (82) as

$$T_{1\%} = \left(\frac{100}{K_{\text{template}} P_{\perp}} \right)^2, \quad (86)$$

giving

$$T_{1\%}^{\text{inj}} \approx 18.5 \text{ s}, \quad (87)$$

$$T_{1\%}^{\text{flat}} \approx 4.9 \text{ min}. \quad (88)$$

Both times exceed the coherence time $\tau_{\text{coh}}^{\text{eff}} \simeq 2.04 \text{ ms}$ [Eq. (75)], so the measurement proceeds as a sequence of tip-read-restore cycles, each within one coherence window, accumulated statistically over $N \sim T_{1\%}/T_{\text{cycle}}$.

For reference, the integration time at full transverse projection ($P_{\perp} = 0.7$, corresponding to $\sin \alpha = 1$) would be $T_{1\%}^{\text{inj}} \approx 16.6 \text{ ms}$ at injection and $T_{1\%}^{\text{flat}} \approx 266 \text{ ms}$ at flattop. The $1/\sin^2 \alpha$ scaling of $T_{1\%}$ at fixed precision accounts for the three-orders-of-magnitude increase between full projection and the operating tipping angle.

Table 7 summarizes the three analysis strategies discussed in Sec. 6.5 and Sec. 6.6. The near-cancellation of the alternating injection pattern in the naive every-other-turn sum leaves an effective sample of one bunch at injection and four at flattop, and the corresponding $1/N$ degradation makes the integration time impractically long. Selecting only the same-sign bunches recovers approximately half the fill at injection, but with a factor-of-two reduction in usable samples. The matched filter recovers the full N_{fill} samples and is the strategy used throughout this analysis.

Table 7: Effective pattern sum C and integration time $T_{1\%}$ for $\delta P/P = 1\%$ at the operating tipping angle $\alpha = 30 \text{ mrad}$ and beam polarization $P = 0.7$ (Table 1), with K_{template} from Eq. (83) evaluated at injection and flattop and the alternative pattern sums discussed in Sec. 6.5.

Condition	Analysis	C	$T_{1\%}$
Injection, 290 bunches	Naive every-other-turn	1	17.98 days
Injection, 290 bunches	Select same-sign only	145	73.9 s
Injection, 290 bunches	Matched filter	290	18.47 s
Flattop, 1160 bunches	Naive every-other-turn	4	287.9 days
Flattop, 1160 bunches	Matched filter	1160	4.929 min

7 Design overview: three-channel SQUID polarimeter

The single-channel transverse-pickup operation developed in Secs. 4–5 for spin-tune determination and the matched-filter precision-mode analysis of Sec. 6.8 measure one in-plane component of the polarization vector at a time. The full proton polarimetry programme at the HSR, including absolute determination of the stored vertical polarization P_y and on-line monitoring of the longitudinal component P_z at the interaction region, requires resolving all three orthogonal components of the polarization vector $\vec{P} = (P_x, P_y, P_z)$. This is achieved by a three-channel SQUID pickup that shares a single cylindrical former around the beam pipe and a single 4 K cryogenic envelope.

Figure 13 shows the three pickup channels schematically. A cosine- θ saddle-loop pair, with winding density proportional to $\cos \phi$ around the cylinder azimuth, picks up the radial in-plane magnetic-moment component and is therefore sensitive to P_x . A sine- θ saddle-loop pair, identical in geometry but rotated by 90° about the cylinder axis, picks up the vertical component and




Figure 12: SNR as a function of integration time T for the matched filter at injection and flattop, for the operating-point transverse polarization $P_{\perp} = P \sin \alpha = 0.021$ (Table 1) and the beam parameters of Tables 1, 3, and 8. Horizontal dotted lines indicate SNR = 10, 100, and 1000, corresponding to relative statistical uncertainties of 10%, 1%, and 0.1%, respectively. Vertical dotted lines mark the corresponding integration times. The grey shaded region indicates $T > \tau_{\text{coh}}^{\text{eff}}$, using the canonical spin-tune spread $\sigma_{\nu_s}^{\text{eff}} = 10^{-3}$ and $\tau_{\text{coh}}^{\text{eff}} \simeq 2.04 \text{ ms}$ from Eq. (75).

is therefore sensitive to P_y . A coaxial pair of azimuthally symmetric loops separated along the beam direction forms an axial gradiometer sensitive to the longitudinal component P_z . The three winding patterns are orthogonal in the sense that each one projects onto a single component of the bunch magnetic moment as the bunch passes through, with negligible cross-talk for an ideally-wound common-cylinder geometry.

The three channels couple to the three orthogonal polarization components in distinct temporal patterns set by the spin dynamics. The cosine- θ channel sees an oscillating signal at $f_s = \nu_s f_{\text{rev}} = f_{\text{rev}}/2$ during free-induction decay, because the P_x component reverses sign every machine turn as the spin precesses about the stable spin axis $\vec{n}_0 \simeq \vec{e}_y$; this is the channel implicitly used in Secs. 4–5. The axial gradiometer sees a similar oscillating signal at f_s for any non-zero P_z component, whether deliberately excited by a tipping pulse applied about $\vec{e}_x$ or already present in the stored beam as a residual longitudinal polarization. In a non-ideal lattice such residual P_z , and equivalently residual P_x on the cosine- θ channel, arises from imperfect injection matching, snake mismatch, rotator transients, or other lattice imperfections, and produces a continuous signal at f_s in the absence of any tipping pulse. The two transverse channels therefore serve a dual role: they support deliberate FID and spin-echo measurements after a tipping pulse, and they monitor the residual non-vertical polarization components of the stored beam in steady state. The sine- θ channel, in contrast, picks up the stored vertical polarization P_y that lies along $\vec{n}_0$ and is therefore static on the spin-precession time scale, producing a signal at the bunch repetition frequency f_{rev} that requires no tipping pulse to read out. All three channels share the same beam pass, the same mechanical support, and the same cryostat, so their relative geometric calibration is fixed by the mechanical reference frame and does not drift between channels.

The remainder of the paper develops the polarimetric reconstruction for each channel class in




Figure 13: Schematic pickup-basis geometries used to illustrate the three magnetic-field channels considered in this work. The cosine- θ and sine- θ pickups form an orthogonal transverse basis and are shown as saddle-coil patterns rotated by 90° with respect to each other. The axial gradiometer consists of two longitudinally separated loops with differential readout for common-mode rejection. The former is drawn as a transparent guide to show the winding geometry. For visual clarity, the transverse saddle-coil panels are shown with the reduced value $n_{\text{turns}} = 10$; this value is used only for the schematic visualization. The geometric input parameters used for the drawings are summarized in Table 8.

Table 8: Pickup-coil parameters for the three-channel SQUID polarimeter. Top block: saddle parameters for the cosine- θ and sine- θ transverse channels, following the construction shown schematically in Fig. 13; the schematic in that figure is drawn with reduced turn count $n_{\text{turns}} = 10$ for visual clarity, the operational value is $n_{\text{turns}} = N_{\text{turns}} = 100$. Bottom block: axial-gradiometer parameters for the longitudinal channel (Sec. 9).

Parameter	Symbol	Unit	Value
<i>Cosine-θ and sine-θ transverse saddle pickup</i>			
Saddle former radius	r_{coil}	mm	40
Base half-length of saddle loop	L_{base}	mm	7.5
Length increment between turns	ΔL	mm	1.0
End-loop height parameter	H_{end}	mm	4.5
Maximum saddle angle	θ_{max}	$^\circ$	75
Winding multiplicity	N_{turns}	1	100
Total length along z	$L_z = 2[L_{\text{base}} + (N_{\text{turns}} - 1)\Delta L + H_{\text{end}}]$	mm	222
Single-turn effective coupling area	A	cm^2	40
Longitudinal form factor	f_{geom}	1	1.5
<i>Axial gradiometer (longitudinal P_z pickup)</i>			
Axial-loop radius	r_{coil}	mm	40
Gradiometer baseline (centre-to-centre)	Δz	mm	300

turn. Section 8 treats the transverse polarimetry of P_x and P_y via the two saddle-loop channels, including the static bunch magnetic moment, the three-channel single-passage response, the finite-bunch-length convolution, the response matrix, and the cross-talk hierarchy. Section 9 treats the longitudinal polarimetry of P_z via the axial gradiometer, with the charge-current rejection geometry that distinguishes it from the saddle-loop pickups, the gradiometer-specific signal-to-noise treatment, and the operational regime in which the longitudinal channel adds information beyond what the transverse channels already provide.

8 Full spin-vector extension and static vector polarimetry

The design overview of Sec. 7 introduced the three-channel pickup basis on a common cylindrical former. This section develops the polarimetric reconstruction of the two transverse polarization components P_x and P_y from the saddle-loop channels (cosine- θ and sine- θ). The longitudinal component P_z , picked up by the axial gradiometer, is treated separately in Sec. 9 because its winding geometry and charge-current rejection mechanism differ qualitatively from the saddle-loop case. The cross-talk between channels in a realistic apparatus, and the broader systematic-effect program required to convert per-channel sensitivities into absolute polarimetric precision, are framed in Sec. 8.7 but their quantitative evaluation is deferred to a dedicated engineering paper.

The starting point is the local polarization vector

$$\vec{P} = (P_x, P_y, P_z) \quad (89)$$

at the pickup location. The FID, spin-echo, and matched-filter modes of Secs. 4–6.8 are special cases of this framework in which a controlled tipping pulse generates a deliberately oscillating in-plane component and the cosine- θ channel reads it out. The static-vector framework developed in this section is the more general case: in steady state, each saddle-loop channel reads out a definite component of the bunch magnetic moment $\vec{m} = N_p \mu_p \vec{P}$, with the time structure imposed by the bunch passage rather than by any deliberate spin manipulation.

At a pickup location in an ordinary HSR arc or straight section where the local stable spin axis is close to vertical, $\vec{n}_0 \simeq \vec{e}_y$, a stored vertical polarization component P_y produces a magnetic moment component m_y that is constant on the spin-precession time scale and pulses at f_{rev} as each bunch passes the pickup. Residual in-plane and longitudinal components P_x and P_z , arising from imperfect injection matching, snake mismatch, rotator transients, or other lattice imperfections, produce moment components that oscillate at $f_s = f_{\text{rev}}/2$ in the spin-precession frame and therefore appear as alternating bunch-to-bunch signals at the pickup. The three symmetry classes of the pickup basis separate these contributions cleanly.

8.1 Static bunch magnetic moment

For bunch b , the total magnetic moment is written as

$$\vec{m}_b = N_{p,b} \mu_p \vec{P}_b, \quad (90)$$

where $N_{p,b}$ is the bunch population, μ_p is the proton magnetic moment, and $\vec{P}_b = (P_{x,b}, P_{y,b}, P_{z,b})$ is the local bunch polarization vector. Equation (90) is the common starting point for both the FID-based transverse signal and the static vector measurement. The distinction is not the magnetic moment itself, but the time dependence imposed on it: in the FID mode the transverse component is deliberately generated and allowed to precess, whereas in the static-vector mode the pickup basis is used to reconstruct the instantaneous local vector components from the known bunch-passage pattern.




Figure 14: Radial field component $B_r(\phi, z)$ on the cylindrical pickup surface, shown in unwrapped-cylinder coordinates for the three static dipole orientations. These three symmetry classes form the natural field basis for the cosine- θ , sine- θ , and axial pickup channels.

8.2 Natural pickup basis on a cylindrical former

On a cylindrical pickup surface surrounding the beam, the three dipole components separate into different angular symmetry classes. For a point dipole on the beam axis with moment $\vec{m} = (m_x, m_y, m_z)$, evaluated on a cylindrical pickup surface of radius r_{coil} , the radial field components are

$$B_r^{(x)}(\phi, z) = \frac{\mu_0 m_x}{4\pi} \frac{2r_{\text{coil}}^2 - z^2}{(r_{\text{coil}}^2 + z^2)^{5/2}} \cos \phi, \quad (91)$$

$$B_r^{(y)}(\phi, z) = \frac{\mu_0 m_y}{4\pi} \frac{2r_{\text{coil}}^2 - z^2}{(r_{\text{coil}}^2 + z^2)^{5/2}} \sin \phi, \quad (92)$$

$$B_r^{(z)}(\phi, z) = \frac{\mu_0 m_z}{4\pi} \frac{3r_{\text{coil}} z}{(r_{\text{coil}}^2 + z^2)^{5/2}}. \quad (93)$$

The transverse components are first-harmonic angular modes on the pickup cylinder, while the longitudinal component is independent of ϕ . The natural correspondence is therefore

$$m_x \leftrightarrow \cos \phi, \quad m_y \leftrightarrow \sin \phi, \quad m_z \leftrightarrow \text{azimuthally symmetric axial response}. \quad (94)$$

The three pickup channels were introduced in Sec. 7; here we verify their mathematical orthogonality. The ideal cosine- θ and sine- θ windings are orthogonal by the angular integrals

$$\int_0^{2\pi} \cos \phi \sin \phi d\phi = 0, \quad \int_0^{2\pi} \cos \phi d\phi = \int_0^{2\pi} \sin \phi d\phi = 0. \quad (95)$$

Thus, in the ideal continuous-winding limit, the transverse channels are mutually orthogonal and are also blind to azimuthally symmetric fields. The axial channel belongs to a different symmetry class and is correspondingly orthogonal to the transverse first-harmonic channels in the ideal geometry. The corresponding symmetry patterns on the unwrapped cylindrical surface are shown in Fig. 14.

8.3 Three-channel single-passage response

The ideal single-passage responses can be written in the form

$$S_x^{\text{bunch}}(z_b) = K_x m_x g_T(z_b; \sigma_L), \quad (96)$$

$$S_y^{\text{bunch}}(z_b) = K_y m_y g_T(z_b; \sigma_L), \quad (97)$$

$$S_z^{\text{bunch}}(z_b) = K_z m_z g_z(z_b; \sigma_L), \quad (98)$$

where z_b is the bunch-centre coordinate relative to the pickup, K_x , K_y , and K_z are geometry and flux-transformer constants, and g_T and g_z are the finite-bunch-length waveform templates for the transverse and axial channels. For a mechanically symmetric implementation one expects $K_x \simeq K_y$; deviations from equality provide a direct calibration observable.

The transverse template g_T is common to the cosine- θ and sine- θ channels because they differ only by a 90° azimuthal rotation of the winding pattern. The axial template g_z differs qualitatively because the first-order axial gradiometer produces a bipolar waveform as the bunch passes the two oppositely wound axial loops.

8.4 Finite bunch-length convolution

The bunch is modeled by a Gaussian longitudinal profile of rms length σ_L ,

$$\lambda(z') = \frac{1}{\sqrt{2\pi}\sigma_L} \exp\left(-\frac{z'^2}{2\sigma_L^2}\right). \quad (99)$$

For each pickup channel, the point-dipole response is convolved with this longitudinal distribution. The transverse template can be written schematically as

$$g_T(z_b; \sigma_L) = \int_{-\infty}^{+\infty} \lambda(z') h_T(z_b - z') dz', \quad (100)$$

where h_T is the point-bunch transverse response of the cosine- θ or sine- θ winding. The axial template is similarly

$$g_z(z_b; \sigma_L) = \int_{-\infty}^{+\infty} \lambda(z') h_z(z_b - z') dz', \quad (101)$$

where h_z is the point-bunch response of the first-order axial gradiometer. The explicit axial-loop expression is developed in Sec. 9; in the present section we treat it abstractly as the P_z component of the vector-pickup basis.

At injection, σ_L is long compared with a pickup radius of order 4 cm, and finite-bunch-length effects strongly modify the peak amplitude and waveform shape. At flattop the bunch is shorter, and the axial gradiometer becomes more favorable. This is why the axial-channel benchmark cannot be inferred from the point-dipole peak field alone.

8.5 Beam-synchronous static signal formation

The measured three-channel signal can be written as

$$\vec{\Phi}^m(t) = \sum_{n,b} \mathbf{H}(t - t_{n,b}) c_b \vec{m}_b + \vec{B}^{\text{bg}}(t), \quad (102)$$

where $t_{n,b}$ are the known bunch-passage times, c_b is the known bunch-pattern or spin-sign factor, $\mathbf{H}$ is the three-channel response kernel, and $\vec{B}^{\text{bg}}(t)$ denotes spin-independent and spin-correlated backgrounds. In the ideal basis the kernel is diagonal in the channel index, with the transverse channels sharing the same waveform template and the axial channel using the gradiometer template.

For a known bunch pattern, the same phase-corrected coherent summation used in the FID analysis can be generalized to three simultaneous matched filters. The estimated vector amplitude is obtained by applying the channel-specific template and the bunch-pattern phase factor to each channel separately,

$$\hat{m}_i \propto \sum_{n,b} \int dt \Phi_i^m(t) q_i(t - t_{n,b}) c_b, \quad i = x, y, z, \quad (103)$$




Figure 15: Single-passage waveform classes after finite bunch-length convolution: (a) unipolar transverse template g_T from a cosine- θ or sine- θ saddle loop, and (b) bipolar axial template g_z from the axial gradiometer. The two curves in each panel correspond to the injection and flattop rms bunch lengths σ_L listed in Tab. 1 and the pickup radius r_{coil} in Tab. 8. Both panels are normalized to the flattop peak, so the injection curves directly show the amplitude reduction from bunch washout when $\sigma_L \gtrsim r_{\text{coil}}$.

where q_i is the matched-filter template for channel i . Calibration then converts $\hat{m}_i$ into $\hat{P}_i$ using Eq. (90). This formulation also allows separate accumulation of bunch subsets, for example up/down spin classes or individual tagged bunches.

The three channels need not be operated in only one mode. They may be read out simultaneously if cross-talk is sufficiently small, or interleaved in time to reduce live coupling and facilitate systematic checks. Interleaved operation would still provide quasi-continuous vector reconstruction because the bunch-passage pattern and machine RF clock are known.

8.6 Background rejection and cross-talk hierarchy

The three channels have different dominant background and rejection mechanisms. For the axial gradiometer, the ideal charge-current rejection is geometric: the beam current generates an azimuthal magnetic field, while the axial loop measures axial flux. A single coaxial axial loop is therefore intrinsically blind to the ideal beam-current field, and the first-order gradiometer further rejects uniform environmental B_z backgrounds.

For the transverse cosine- θ and sine- θ channels, the azimuthal symmetry of the beam-current field is rejected by the first-harmonic winding pattern in the ideal continuous limit. The rejection argument is sensitive to the azimuthal symmetry of the beam current itself. The EIC HSR beam, however, is deliberately flat by design: peak-luminosity operation requires a transverse emittance ratio of order $\varepsilon_x/\varepsilon_y \simeq 10$, produced by electron cooling at the IR2 cooling section at injection and maintained against intrabeam scattering by hadron cooling at top energy [42]. Combined with the matched-optics β -function asymmetry at the pickup location, this yields rms beam sizes that are anisotropic by a factor of several (for example $\sigma_x \simeq 3.5$ mm versus $\sigma_y \simeq 0.7$ mm at injection in the matched optics at the HJET in IR4 [1]). The wall image current of such a flat bunch therefore carries substantial higher-multipole ($m \geq 2$) azimuthal structure, while the $m = 1$ dipole component to which the saddle pickups respond is suppressed by parity for a centered beam. Residual coupling into the $m = 1$ channel comes from orbit-offset, dispersion, and small mismatches that break this parity, and its magnitude at the pickup radius is bounded by Ref. [1], whose tabulated radial field profiles for symmetric and asymmetric beam distributions are visibly distinct near the beam axis but converge for radii of order the pickup radius $r_{\text{coil}} = 40$ mm. In addition, finite winding accuracy, mechanical asymmetry, image currents, feedthroughs, nearby

conductive structures, and flux-transformer routing produce residual common-mode coupling. Together, these effects are the dominant engineering risk for the transverse static channels and set the required common-mode rejection and trim strategy.

The leading cross-talk classes are:

- x - y mixing from imperfect cosine/sine winding orthogonality, azimuthal placement errors, and unequal layer geometry;
- transverse-axial mixing from loop tilt, broken cylindrical symmetry, or common flux-transformer routing;
- spin-independent backgrounds from charge-current leakage, external fields, vibration in residual field gradients, and trapped flux;
- spin-correlated backgrounds from machine elements whose fields or beam conditions change with spin pattern.

These effects define the response-matrix problem below.

8.7 Systematic effects and cross-talk

In the ideal apparatus, the three pickup channels are mutually orthogonal by the angular-symmetry argument of Eqs. (95), and the relation between the reconstructed channel amplitudes and the local polarization vector is diagonal,

$$\begin{pmatrix} \hat{m}_x \\ \hat{m}_y \\ \hat{m}_z \end{pmatrix} = \begin{pmatrix} R_{xx} & 0 & 0 \\ 0 & R_{yy} & 0 \\ 0 & 0 & R_{zz} \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ P_z \end{pmatrix} + \begin{pmatrix} B_x^{\text{bg}} \\ B_y^{\text{bg}} \\ B_z^{\text{bg}} \end{pmatrix}. \quad (104)$$

In a realistic apparatus, geometric imperfections (winding-density errors, layer offsets, coil tilt, broken cylindrical symmetry), residual conductive structures (beam-pipe coatings, image currents, feedthroughs), and shared flux-transformer routing introduce off-diagonal cross-channel coupling, and the response matrix contains off-diagonal terms,

$$\hat{m}_i = \sum_j R_{ij} P_j + B_i^{\text{bg}}, \quad i, j \in \{x, y, z\}, \quad (105)$$

where the off-diagonal elements quantify geometric cross-talk, alignment errors, and deviations from the ideal basis. The calibration program must therefore determine the diagonal gains R_{ii} and bound or correct the off-diagonal terms R_{ij} .

A quantitative evaluation of the off-diagonal response matrix requires detailed mechanical, electromagnetic, and prototyping inputs that are beyond the scope of the present concept paper. We defer this evaluation, together with the related systematics analysis (mechanical alignment tolerances, common-mode rejection floor of the $\cos\theta$ and $\sin\theta$ windings, residual environmental backgrounds, and spin-correlated machine couplings), to a dedicated engineering paper. Equations (104) and (105) establish the framework within which the systematic-effect categories of Tab. 9 will be quantified.

9 Longitudinal component and axial gradiometer channel

9.1 Physical picture

The local polarization vector $\vec{P} = (P_x, P_y, P_z)$ introduced in Sec. 8 produces a bunch magnetic moment $\vec{m} = N_p \mu_p \vec{P}$ whose three orthogonal components are read out by the three pickup channels of Fig. 13. The cosine- θ and sine- θ saddle channels read out the transverse components

Table 9: Principal classes of systematic effects for static vector SQUID polarimetry. Quantitative evaluation of each category is deferred to a dedicated engineering paper; the present table serves as a forward pointer for the categories that must be addressed.

Category	Origin	Suppression or calibration strategy
Geometric cross-talk	Winding errors, layer offsets, loop tilt	Mechanical survey, bench calibration, response matrix
Charge-current leakage	Broken symmetry, image currents, feedthroughs	Gradiometry, trim coils, reference pickup, shielding
Environmental background	External fields, vibration, trapped flux	SC shield, quiet location, cooldown procedure
Spin-correlated background	Machine or bunch-pattern correlations	Spin-pattern reversal, bunch-subset tests, HJET/pC comparison
Gain and phase errors	SQUID/FLL/electronics response	Injected calibration signals, RF-clock phase reference

m_x and m_y via Eqs. (96)–(97); the axial gradiometer reads out the longitudinal component m_z via Eq. (98), and is the subject of this section.

At an arc pickup location with $\vec{n}_0 \simeq \vec{e}_y$, residual or deliberately excited in-plane and longitudinal components m_x and m_z precess about $\vec{n}_0$ at $\nu_s = 1/2$, so both the cosine- θ and the axial channels see oscillating signals at the spin-precession frequency $f_s = f_{\text{rev}}/2 \simeq 39.1$ kHz. The cosine- θ saddle is by construction insensitive to azimuthally symmetric fields and does not couple to m_z ; the coaxial-loop gradiometer geometry is required to resolve the longitudinal component. Two channel-specific properties developed in this section set the axial channel apart from the saddle channels: a geometric charge-current rejection mechanism that requires no winding-balance tolerancing (Sec. 9.2), and a flux-signal kernel g_z whose finite-bunch-length suppression differs qualitatively from the saddle template g_T (Sec. 9.3).

9.2 Charge current rejection

The beam charge current $I_{\text{beam}}(t)$ flowing along $\vec{e}_z$ produces only an azimuthal magnetic field,

$$\vec{B}_{\text{charge}}(\rho, t) = \frac{\mu_0 I_{\text{beam}}(t)}{2\pi\rho} \vec{e}_\phi, \quad (106)$$

where ρ is the transverse distance from the beam axis. This holds for every Fourier harmonic of the bunch train: whether the current is a DC component, an RF harmonic at $n f_{\text{rev}}$, or the transient field of a single bunch passage, the field is purely azimuthal by the cylindrical symmetry of the problem.

A circular pickup loop coaxial with the beam, with its area normal along $\vec{e}_z$, is threaded only by the axial field component B_z . Since $\vec{B}_{\text{charge}} \cdot \vec{e}_z = 0$ at every point on the loop, the charge-current flux is

$$\Phi_{\text{charge}} = \iint_{\text{loop}} \vec{B}_{\text{charge}} \cdot d\vec{A} = 0 \quad \text{exactly.} \quad (107)$$

This rejection is purely geometric: it requires no special winding balance, holds independently of the beam current magnitude, energy, or transverse beam profile, and applies at all frequencies. A plain circular loop is intrinsically charge-blind.

This is in marked contrast to the cosine- θ coil, where the charge current couples into the winding through any imbalance ϵ in the winding density, and a common-mode rejection ratio $\text{CMRR} = 1/\epsilon \gtrsim 10^{12}$ is required.

9.3 Flux signal from the longitudinal magnetic moment

In the following, we consider the time-dependent longitudinal magnetic moment $m_z(t)$ of a bunch arising from spin precession. A “bunch” denotes the ensemble of particles in a given RF bucket. The detector does not resolve individual particles, but measures the total magnetic moment, i.e., the sum over all particle spins. The observed signal is therefore proportional to the oscillating component $m_z(t)$ and reflects the bunch-averaged polarization.

9.3.1 Point-dipole result

After the small tipping pulse, the total magnetic moment of a bunch remains dominated by the main spin component, while only a small oscillating projection $m_z(t)$ appears along the detector axis $\vec{e}_z$. It is this oscillating longitudinal component, not the full bunch magnetic moment, that generates the signal considered here.

Treating this oscillating component as a point dipole located at the bunch centre z_b on the beam axis, the instantaneous flux threading a coaxial circular loop of radius r_{coil} centred at $z = z_{\text{loop}}$ is

$$\Phi_{\text{loop}}(z_b, t) = \frac{\mu_0}{2} m_z(t) \frac{r_{\text{coil}}^2}{[r_{\text{coil}}^2 + (z_b - z_{\text{loop}})^2]^{3/2}}. \quad (108)$$

For a fixed value of $m_z(t)$ this is a unipolar pulse in the passage coordinate: as the bunch centre approaches the loop the flux rises, reaches its maximum when the bunch centre is in the plane of the loop ($z_b = z_{\text{loop}}$),

$$\Phi_{\text{peak}}^{\text{pt}}(t) = \frac{\mu_0}{2} \frac{m_z(t)}{r_{\text{coil}}}, \quad (109)$$

and falls symmetrically as the bunch recedes. The sign of Φ_{loop} therefore does not change within a single passage. Since $\nu_s = 1/2$, the spin completes exactly half a precession cycle per revolution, so $m_z(t)$ changes sign on every successive bunch passage through the loop. This turn-by-turn sign alternation is the signature of the $\nu_s = 1/2$ spin tune and is what the matched filter locks onto at the sideband frequency $\nu_s f_{\text{rev}} = 39.1$ kHz.

9.3.2 Finite bunch length

At injection the rms bunch length $\sigma_L = 24$ cm exceeds $r_{\text{coil}} = 4$ cm by a factor of six, so the point-dipole approximation significantly overestimates the peak flux. For a Gaussian bunch of rms length σ_L with its centre at z_b , the total magnetic moment is distributed along z , and the instantaneous loop flux is obtained by convolution,

$$\Phi_z(z_b, t) = \frac{\mu_0 m_z(t)}{2\sqrt{2\pi} \sigma_L} \int_{-\infty}^{+\infty} \frac{r_{\text{coil}}^2 \exp(-z'^2/2\sigma_L^2)}{[r_{\text{coil}}^2 + (z_b - z')^2]^{3/2}} dz'. \quad (110)$$

In the limit $\sigma_L \gg r_{\text{coil}}$ the convolution gives $\Phi_z^{\text{peak}}(t) \approx \mu_0 m_z(t)/\sigma_L$: the same total magnetic moment is spread over a longer bunch, and the peak flux therefore falls as $1/\sigma_L$, suppressing the point-dipole result by a factor of ~ 6 at injection. At flat-top ($\sigma_L = 6$ cm, comparable to r_{coil}) the suppression is only a factor of ~ 2 . Equation (110) is evaluated numerically using the beam parameters of Tab. 1 and the axial-gradiometer geometry of Tab. 8; the resulting peak fluxes at the SQUID input are given in Eq. (112).

9.4 Gradiometer configuration

A single axial loop is immune to the beam charge current but would pick up any uniform external B_z (residual Earth’s field, stray fields from ring magnets, low-frequency environmental noise). A first-order axial gradiometer – two identical coaxial loops wound in series opposition

and separated by Δz – cancels any field component uniform over Δz , while retaining the bunch signal because the non-uniform dipole field differs between the two loop positions.

With loop 1 at $z = +\Delta z/2$ and loop 2 at $z = -\Delta z/2$, the gradiometer output is

$$\Phi_{\text{grad}}(z_b) = \Phi_z(z_b; z_{\text{loop}} = +\frac{\Delta z}{2}) - \Phi_z(z_b; z_{\text{loop}} = -\frac{\Delta z}{2}), \quad (111)$$

where each term is given by Eq. (110). As the bunch sweeps through the gradiometer, Φ_{grad} traces a bipolar waveform: it peaks positively near loop 1, passes through zero at the midpoint $z_b = 0$, and peaks negatively near loop 2. A uniform ambient B_z gives zero gradiometer output.

In the point-dipole limit the peak of $|\Phi_{\text{grad}}|$ is maximised at $\Delta z/2 = r_{\text{coil}}/\sqrt{2} \approx 2.8$ cm. We adopt $\Delta z = 30$ cm, constrained by the centre-to-centre separation of coils A and B in the cryostat. Despite being far from the point-dipole optimum, numerical evaluation of Eq. (111) with the full Gaussian bunch profile shows that the gradiometer retains $\sim 65\%$ of the single-loop peak signal: the long injection bunch ($\sigma_L = 24$ cm) already spans both loop positions, so the penalty for large Δz is modest. The resulting post-pickup peak fluxes are quoted in Eq. (112).

The per-bunch peak flux delivered to the SQUID input by the axial gradiometer evaluates to

$$\Phi_{\text{pickup}}^{\text{axial, inj}} \approx 174.9 \mu\Phi_0, \quad \Phi_{\text{pickup}}^{\text{axial, flat}} \approx 221.3 \mu\Phi_0, \quad (112)$$

to be compared with the cosine- θ values $1236 \mu\Phi_0$ (injection) and $309 \mu\Phi_0$ (flattop) of Eq. (60). The flux ratio $\Phi_{\text{pickup}}^{\text{axial}}/\Phi_{\text{pickup}}^{\cos\theta}$ is 0.142 at injection and 0.716 at flattop, consistent with the sensitivity-ratio discussion of Sec. 10.

9.5 Signal-to-noise ratio

9.5.1 Signal accumulation: how SNR is built up

Each time a bunch passes through the gradiometer it produces a flux signal of peak amplitude Φ_z^{peak} [Eq. (112)]. This per-bunch signal is well below the instantaneous noise floor of the SQUID: for the canonical white flux noise spectral density $S_{\Phi}^{1/2} = 0.4 \mu\Phi_0/\sqrt{\text{Hz}}$, consistent with the performance of modern wideband LTS DC SQUID systems [43] at the 39 kHz signal frequency where the noise is white and the $1/f$ corner of LTS devices lies well below, and a pulse bandwidth $\sim 1/\sigma_t \sim 1$ GHz, the single-pulse noise is $\sim 1.3 \times 10^4 \mu\Phi_0$, exceeding the per-bunch signal by several orders of magnitude. The signal is recovered by coherent summation: each bunch produces a pulse at a known time, and the pulses from successive revolutions carry an amplitude modulated at the spin-tune sideband $\nu_s f_{\text{rev}} = 39.1$ kHz. A matched filter locked to this sideband sums the contributions coherently, so the signal grows linearly with the number of bunch passages while the noise grows only as the square root. After N_{pass} coherently summed bunch passages,

$$\text{SNR} = \frac{\Phi_z^{\text{peak}} \sqrt{N_{\text{pass}}}}{S_{\Phi}^{1/2}}, \quad (113)$$

since $N_{\text{pass}} \propto T$, the SNR grows as $\sqrt{T}$. The two strategies below differ only in how many bunches are summed per revolution.

9.5.2 Strategy 1: all-bunch coherent summation

All N_b bunches in the fill are summed coherently on each revolution. The spin orientation of each bunch is known from the injection pattern and the spin tune $\nu_s = 1/2$: consecutive bunches have alternating spin signs, and the matched filter applies the correct phase weight to each. The number of bunch passages accumulated in time T is $N_{\text{pass}} = N_b f_{\text{rev}} T$, giving

$$\text{SNR}_z^{(1)}(T) = \frac{\Phi_z^{\text{peak}} \sqrt{N_b f_{\text{rev}} T}}{S_{\Phi}^{1/2}}. \quad (114)$$

9.5.3 Strategy 2: single-bunch tracking

A single tagged bunch is followed turn by turn, ignoring all other bunches. Here $N_{\text{pass}} = f_{\text{rev}} T$ (one bunch passage per revolution), and the SNR is lower than strategy 1 by $\sqrt{N_b}$:

$$\text{SNR}_z^{(2)}(T) = \frac{\Phi_z^{\text{peak}} \sqrt{f_{\text{rev}} T}}{S_{\Phi}^{1/2}} = \frac{\text{SNR}_z^{(1)}(T)}{\sqrt{N_b}}. \quad (115)$$

At injection $\sqrt{N_b} = \sqrt{290} \approx 17$. The same single-bunch matched-filter formulation extends to every bunch in the fill via parallel matched-filter channels on the digitized SQUID stream, and is not specific to the axial channel: the $\cos\theta$ and $\sin\theta$ channels admit the same per-bunch decomposition. The full bunch-resolved vector polarimetry capability (joint three-channel reconstruction of $\vec{P}(t)$ for each bunch, per-bunch precision as a function of averaging window, and the operational comparison with existing polarimetry) is developed as the joint deliverable of Sec. 10.

9.5.4 Coherence-time interaction and multi-cycle accumulation

The axial-gradiometer channel is subject to the same reversible coherence-time limit as the transverse FID and matched-filter measurements of Secs. 4–6.8. Tipped or residual perpendicular spin components dephase with the rms spin-tune spread σ_{ν_s} on the timescale $\tau_{\text{coh}} = 1/(2\pi f_{\text{rev}} \sigma_{\nu_s})$ defined in Eq. (42). The canonical operating-point value, $\sigma_{\nu_s}^{\text{eff}} = 10^{-3} \rightarrow \tau_{\text{coh}}^{\text{eff}} \simeq 2.04 \text{ ms}$ (Tab. 5), is used throughout the present analysis; the DLC and LC test scales of Tab. 3 give correspondingly shorter coherence times.

The axial channel is no different in this respect from the transverse matched-filter channel: both extract signals whose envelopes decay on τ_{coh} . Any matched-filter integration time $T > \tau_{\text{coh}}$ must therefore proceed as a sequence of $N \simeq T/\tau_{\text{coh}}$ tip- π -echo-restore cycles, each within one coherence window, using the spin-echo extension of Sec. 5. The per-cycle integration times tabulated below are noise-limited single-cycle values; their realization at $T > \tau_{\text{coh}}$ incurs the bounded per-cycle depolarization budget of Sec. 5.

9.5.5 Numerical performance and measurement strategies

The sensitivity of the axial P_z channel can be compared directly with the transverse cosine- θ reference. For fixed N_b , f_{rev} , and integration time T , the SNR is proportional to the peak flux. The ratio of the axial and transverse sensitivities is therefore simply

$$\frac{\text{SNR}_z}{\text{SNR}_x} = \frac{\Phi_z^{\text{grad}}}{\Phi_x}. \quad (116)$$

At injection this ratio is 0.142, while at flattop it is 0.716. The axial gradiometer is less sensitive than the cosine- θ coil at both energies, but the gap is much smaller at flattop than at injection. The reduction reflects the strong bunch-length dependence of the axial signal: in the long-bunch limit it scales approximately as $1/\sigma_L$, whereas the cosine- θ signal is much less sensitive to σ_L . The shorter flattop bunch, $\sigma_L = 6 \text{ cm} \simeq r_{\text{coil}}$, is therefore strongly favourable for the axial gradiometer.

The all-bunch matched-filter integration times for the axial channel are therefore obtained directly by rescaling the cosine- θ values of Sec. 6.8 by the inverse of these ratios. The resulting per-cycle times exceed the reversible coherence time τ_{coh} at every relevant σ_{ν_s} scale, so the axial measurement must be accumulated as a sequence of tip- π -echo-restore cycles using the spin-echo extension of Sec. 5, matching the multi-cycle strategy already established for the transverse matched-filter channel. Single-bunch tracking is one operational mode of a per-bunch matched-filter framework that applies equally to the $\cos\theta$ and $\sin\theta$ channels (Sec. 9.5.3); the bunch-resolved vector polarimetry deliverable and its comparison with existing storage-ring polarimetry are developed in Sec. 10.

9.5.6 SQUID readout bandwidth

The spin signal arrives as a train of bunch-passage signals at the revolution frequency $f_{\text{rev}} = 78 \text{ kHz}$, with the amplitude modulated at the spin-tune sideband $\nu_s f_{\text{rev}} = 39 \text{ kHz}$. A DC SQUID operated in a wideband flux-locked loop (FLL) readily handles this: modern LTS SQUID FLL systems achieve bandwidths of $\sim 1 \text{ MHz}$ [43], well above the 39 kHz signal frequency. The SQUID sensor itself is a DC device, but the FLL electronics extend the linear dynamic range and tracking bandwidth into the RF range required here.

9.6 Real-world effects not yet addressed

The arguments above establish the ideal geometric sensitivity of the axial gradiometer, but several non-ideal effects remain to be quantified before this channel can be promoted from a sensitivity estimate to an engineering design. The most important ones are listed here for later follow-up.

- **Mechanical alignment of the pickup loop.** As argued in Section 9.2, the charge-current rejection is exact for any $\vec{e}_z$ -directed current, regardless of the transverse beam position, beam shape, or current density profile. The only misalignment that breaks this rejection is a tilt of the pickup loop relative to the beam axis, which introduces a component of the loop normal in the azimuthal or radial direction that couples to $\vec{B}_{\text{charge}}$. The scaling of the residual flux with tilt angle should be worked out explicitly, preferably with a 3D field model that includes the actual beam-pipe geometry, in order to establish the mechanical alignment tolerance for the loop.
- **Image currents and nearby conductive structures.** The ideal argument assumes free space. In reality the beam tube, copper coating, bellows, transitions, and neighbouring superconducting leads reshape the field lines and can convert part of the beam-current field into a small but non-zero axial pickup. The copper layer must also be thick enough to guide the image currents over the relevant RF harmonics; a rough skin-depth estimate suggests that 50 μm to 100 μm of Cu should already be sufficient in the tens-of-MHz range, but this should be checked with a realistic field model of the local environment.¹
- **Ambient and non-uniform B_z backgrounds.** A first-order axial gradiometer cancels only the component of B_z that is uniform across the two-loop baseline Δz . Longitudinal field gradients, vibration-driven motion in residual fields, and stray fields from nearby magnets or pulsed devices can leak into the signal channel. The expected environmental rejection therefore needs to be estimated for the actual cryostat and ring location.
- **Cross-talk from the P_x and P_y channels.** In the full three-channel concept the cosine- θ , sine- θ , and axial loops share the same cryostat and former. Mutual inductive coupling between channels, imperfect orthogonality of the pickup geometries, and common flux-transformer routing may produce cross-talk. This should be quantified experimentally on a prototype former.
- **Bunch-to-bunch variations.** In the present estimate all bunches are assumed to have identical intensity and bunch length. In practice, the bunch charge may vary across the fill because the bunches are injected individually from the AGS. This leads primarily to a bunch-dependent signal amplitude, but not to a corresponding variation of the signal frequency. By contrast, substantial bunch-to-bunch variations of the bunch length are not expected on the relevant timescale, since all bunches follow the same RF program of the

¹The skin depth in copper is $\delta = \sqrt{2\rho/(\omega\mu_0)}$ with $\rho \approx 1.7 \times 10^{-8} \Omega\text{m}$, giving $\delta \approx 20 \mu\text{m}$ at 10 MHz and $\delta \approx 6.5 \mu\text{m}$ at 100 MHz. A coating of 5–10 skin depths, i.e. 50 μm to 100 μm , is therefore sufficient; a somewhat thicker coating of 0.1 mm to 0.2 mm would provide margin for fabrication tolerances and higher frequency content.

machine. Thus, the dominant effect of bunch-to-bunch non-uniformity is expected to be amplitude weighting in the coherent sum, rather than loss of phase coherence.

- **Absolute calibration.** The present note focuses on relative sensitivity. The full readout chain (loop geometry, flux-transformer coupling, SQUID gain, matched-filter normalization, and the effective FID tipping angle α) introduces multiplicative factors that propagate into the polarization scale and cannot be eliminated by internal characterization alone. Absolute polarimetry therefore requires cross-calibration of the SQUID against an external polarization standard; the atomic-hydrogen-jet (HJET) polarimeter, with its proven absolute-polarization capability, will provide this standard at the EIC [1].
- **Validity of the assumed coherence model.** The coherence estimate above is intentionally simple and is adequate for this brief writeup. A more realistic treatment should use the actual spin-tune distribution expected for the machine optics and bunch parameters, including any time dependence across the store.

9.7 Implementation

The axial gradiometer requires two plain circular NbTi loops wound directly on the same Macor ceramic former as the cosine- θ and sine- θ coils. The loops are placed at the axial positions of coils A and B ($\Delta z = 30$ cm centre-to-centre), with loop radius $r_{\text{coil}} = 40$ mm matching the saddle-pickup inner radius of the former. They are connected in series opposition and coupled through a superconducting flux transformer to one additional wideband FLL SQUID channel. No special winding geometry is required.

The present design assumption is that the beam tube itself must be cold and located inside the cryostat, together with the pickup loops. Proceeding radially outward from the beam axis, the layout is as follows: the beam travels through the hollow ceramic former; the superconducting NbTi pickup loops are wound on the outer surface of the former; a vacuum gap separates the former from the surrounding stainless steel beam tube, whose inner (vacuum-facing) surface carries a thin copper coating that guides the beam image currents in a well-defined way. This fully cryogenic layout is expected to reduce uncontrolled magnetic pickup, preserve the cylindrical symmetry underlying the charge-current rejection argument, and keep the electromagnetic environment of the gradiometer as simple as possible.

9.8 Conclusions for the axial channel

The axial loop gradiometer provides a viable and geometrically elegant route to P_z measurement. Its key properties are:

- **Exact charge current rejection by geometry.** B_ϕ from the beam current is tangential to the axial loop normal at every point, giving zero flux at all frequencies with no winding balance required.
- **Per-bunch flux of the same order as the cosine- θ coil at flattop.** At the SQUID input, the per-bunch peak flux is $\Phi_{\text{pickup}}^{\text{axial}} \approx 174.9 \mu\Phi_0$ (injection) and $221.3 \mu\Phi_0$ (flattop). Relative to the cosine- θ coil [Eq. (60)], this is $\sim 14\%$ at injection and $\sim 72\%$ at flattop: the axial channel is significantly weaker than the saddle channel at injection but approaches it at flattop, where the shorter bunch length ($\sigma_L = 6$ cm $\simeq r_{\text{coil}}$) is favourable for the axial geometry.
- **Per-cycle precision and multi-cycle accumulation (Strategy 1).** At the canonical operating point ($\alpha = 30$ mrad, $P = 0.70$, $S_\Phi^{1/2} = 0.4 \mu\Phi_0/\sqrt{\text{Hz}}$), the all-bunch matched-filter sensitivity is $K_z = 156.8 \sqrt{\text{s}}$ at injection and $198.4 \sqrt{\text{s}}$ at flattop, giving noise-limited 1% P_z measurement times $T_{1\%} \simeq 7.53$ min at injection (for a residual $P_z \simeq 0.03$ anchored

on the pC transverse-tilt observations of Ref. [44]) and $T_{1\%} \simeq 25.4\text{ s}$ at flattop (for a residual $P_z \simeq 0.10$). Both exceed the canonical coherence time $\tau_{\text{coh}}^{\text{eff}} \simeq 2.04\text{ ms}$ (Tab. 5) and are realized as multi-cycle tip- π -echo-restore sequences per Sec. 5.

- **Bunch-resolved vector polarimetry (Strategy 2).** The single-bunch matched-filter framework extends to all three channels and provides continuous bunch-by-bunch reconstruction of $\vec{P}(t)$ throughout the fill; the joint three-channel capability is developed in Sec. 10.
- **Minimal hardware.** One wideband FLL SQUID channel; two plain circular NbTi loops on the existing former; no cryostat modification.

Together with the cosine- θ coil (P_x) and the sine- θ coil (P_y), the axial gradiometer completes a three-component noninvasive spin vector polarimeter in a single cryostat.

10 Bunch-resolved vector polarimetry

10.1 Three-channel measurement modes

The three pickup channels of Sec. 7 admit both *static* (no kicker firing) and *dynamic* (with kicker firing) operating modes. In static mode, each component of the local polarization vector produces a signal at its natural frequency, fixed by the precession geometry: P_y , aligned with $\vec{n}_0 \simeq \vec{e}_y$, is constant in time and produces a bunch-passage pulse train at the revolution frequency f_{rev} ; the perpendicular components P_x and P_z precess about $\vec{n}_0$ at the spin tune $\nu_s = 1/2$ and produce oscillating signals at $f_s = f_{\text{rev}}/2$. The signal frequencies are physically determined by the precession and cannot be shifted by operational choice.

In dynamic mode, the longitudinal kicker fires a small-angle pulse synchronized to the phase-locked spin precession established in Sec. 4. The rotation $R_z(\alpha)$ mixes the (P_x, P_y) components in the $(\vec{e}_x, \vec{e}_y)$ plane and leaves P_z unchanged in the same instant. The free precession between successive kicks rotates the perpendicular components (P_x, P_z) continuously in the $(\vec{e}_x, \vec{e}_z)$ plane at f_s , so the in-plane component aligned with $\vec{e}_x$ at any chosen moment is determined by the firing phase $\varphi_k = 2\pi f_s t_k$ relative to the phase-locked precession. By selecting φ_k , the kicker tips either residual P_x ($\varphi_k = 0$) or residual P_z ($\varphi_k = \pi/2$) into the $\vec{e}_y$ direction, where it appears as a small additive contribution to the sin- θ channel signal at f_{rev} . Independently, the same kick rotates stored P_y into the $-\vec{e}_x$ direction with amplitude $P_y \sin \alpha$; this contribution precesses and is read out by the cos- θ channel at f_s regardless of firing phase, providing the Sec. 4 FID readout simultaneously with the in-plane projection extracted via sin- θ .

The six accessible (channel, mode) combinations are summarized in Table 10. Each row corresponds to a distinct matched-filter readout of one polarization component on one channel at one signal frequency. The static rows operate by default throughout the fill; the dynamic rows fire periodically and serve dual roles as α calibration and cross-checks against the static readouts.

A representative operating point can be established using the transverse-spin tilt measured at RHIC by the pC polarimeters [44]. The observed tilts of $\phi_{\text{pC}} \simeq 3^\circ$ at 100 GeV and $\simeq 8^\circ$ – 16° at 255 GeV imply residual transverse polarization components of order $P_x \simeq 0.03$ at injection and $P_x \simeq 0.10$ at flattop. Although no comparable measurement of P_z exists, we adopt the same numerical range as a conservative working estimate for the longitudinal residual; this should be regarded as a working assumption rather than a quantitative claim about P_z . The resulting operating-point integration times for the six (channel, mode) combinations are summarized in Tab. 10. All static-mode integration times lie at or below a few minutes; the longest, for P_z at injection, is $\sim 3\text{ min}$ and corresponds to many tip- π -echo-restore cycles per Sec. 5. Dynamic-mode integration times incur an additional $\sin^2 \alpha$ suppression and are tabulated for cross-calibration purposes. The static-mode saddle-channel signal is linear in the residual

Table 10: Three-channel vector polarimetry measurement modes. Static operation (no kicker firing) is the default and reads all three polarization components simultaneously and continuously, each on its dedicated channel at its natural signal frequency. Dynamic operation fires the longitudinal kicker (tip axis $\vec{e}_z$) at a chosen phase φ_k of the phase-locked spin precession to tip the selected in-plane component into the $\vec{e}_y$ direction, where it is read on the $\sin\theta$ channel at f_{rev} . The same firing simultaneously produces the FID readout on $\cos\theta$ at f_s via the $-P_y \sin\alpha$ contribution. Alternating firings between $\varphi_k = 0$ and $\varphi_k = \pi/2$ provides dynamic readouts of the entire perpendicular plane (P_x, P_z) from a single set of kicker pulses. Per-cycle integration times $T_{1\%}$ are quoted from the per-channel matched-filter sensitivities of Secs. 8–9, evaluated at the operating point $P = 0.7$, $\alpha = 30$ mrad, $S_\Phi^{1/2} = 0.4 \mu\Phi_0/\sqrt{\text{Hz}}$. Static-mode signal amplitudes for P_x and P_z adopt residual polarization values of 0.03 at injection and 0.10 at flattop, anchored on the RHIC pC tilt observations of Ref. [44]; the longitudinal residual is assumed equal to the transverse residual in the absence of any prior measurement.

comp.	mode	channel	sig. freq	kicker phase φ_k	$T_{1\%}$ (inj)	$T_{1\%}$ (flat)
P_y	static	$\sin\theta$	f_{rev}	—	16.6 ms	266 ms
P_x	static	$\cos\theta$	f_s	—	9.05 s	13 s
P_z	static	axial	f_s	—	7.53 min	25.4 s
P_y	dynamic	$\cos\theta$	f_s	any	18.5 s	4.93 min
P_x	dynamic	$\sin\theta$	f_{rev}	0	2.79 h	4.03 h
P_z	dynamic	$\sin\theta$	f_{rev}	$\pi/2$	2.79 h	4.03 h

in-plane polarization component, so a fixed integration time yields a fractional precision $\delta P_x/P_x$ that scales as $1/P_x$; in HSR configurations where the equilibrium $\vec{n}_0$ is tuned close to the vertical and the residual in-plane component is small by design, the integration time required for a target fractional precision grows accordingly, and the dynamic-mode kicker readout (which transfers a controlled fraction of the dominant P_y into the in-plane channel as a calibrated probe) becomes the natural complement for resolving the smallest residual components.

The duty cycle between static monitoring and dynamic calibration is an operational choice. A representative pattern is continuous static monitoring throughout the fill, interrupted periodically (e.g. every T_{cal} minutes) by short dynamic-mode bursts that refresh the f_s phase lock and update the in-situ α calibration. Each burst contains a small number of kicker firings at prescribed φ_k . The cumulative depolarization cost over the full store is bounded by the noninvasive tip- π -echo-restore budget established in Sec. 5, evaluated for the chosen T_{cal} .

10.2 Comparison with existing polarimetry

At the EIC, absolute proton-beam polarization will be measured by the same atomic-hydrogen-jet (HJET) and proton-carbon (pC) instruments developed and operated at RHIC [1]. Three-dimensional reconstruction of the stable spin direction in principle is possible with these instruments by combining the transverse measurements with controlled local spin rotations and spin transport, as demonstrated at RHIC [4], but only as a time-consuming dedicated procedure, providing a fill-averaged result. The three-channel SQUID polarimeter proposed in this paper supplements that suite by delivering continuous, noninvasive, bunch-resolved measurement of the full polarization vector $\vec{P}(t)$ throughout the store. The SQUID output is qualitatively distinct from the established scattering-based instruments, particularly for the longitudinal component P_z , as discussed below.

The pC polarimeter operates with an unpolarized carbon target. Only single-spin observables are accessible, and the longitudinal analyzing power $A_z \propto \vec{S} \cdot \vec{k}$ is a parity-violating pseudoscalar that vanishes identically under parity-conserving strong interactions. The longitudinal polar-

ization P_z is therefore inaccessible to pC polarimetry by parity conservation, irrespective of statistical precision.

The HJET polarimeter uses a polarized atomic-hydrogen jet target and provides access to single-spin and double-spin observables. In its present RHIC operating configuration the target is vertically polarized, so HJET measures the vertical beam polarization P_y via the analyzing power A_y and the spin-correlation parameter A_{yy} . The A_{yy} -based determination requires knowledge of both the target polarization Q_y and the analyzing power A_{yy} itself; the latter is itself a calibrated quantity [45]. Rotation of the target polarization to other directions gives access to P_x (via A_{xx} with a transversely polarized target) and to P_z (via A_{zz} with a longitudinally polarized target); these target-rotation techniques are established [5] and are expected to form part of the upgraded HJET planned for the EIC, but they were not the routine RHIC operating mode.

The SQUID polarimeter is not subject to either limitation. The longitudinal magnetic moment $m_z = N_p \mu_p P_z$ is sensed directly by the axial-gradiometer geometry of Sec. 9, with no scattering process and no target manipulation required. The output is continuous throughout the store and resolved bunch by bunch. To our knowledge, the SQUID concept provides the first practical route to continuous, bunch-resolved, noninvasive $P_z(t)$ monitoring of a stored proton beam.

A related development worth noting is the framework of Ref. [46], in which deliberately precessing in-plane deuteron polarizations are used as a means to search for T -violating asymmetries in pd transmission, exploiting Fourier analysis at f_s and $2f_s$ to expose oscillating signals against backgrounds. That approach relies on active RF spin rotation to generate the in-plane components and on beam-current attenuation to extract asymmetries. The SQUID polarimeter provides the natural detector counterpart: continuous direct measurement of the precessing in-plane and longitudinal magnetic moments, bunch-resolved, with no active manipulation needed. The two approaches are complementary, one generates oscillating signals to expose new physics in scattering, the other senses oscillating signals continuously, and the SQUID polarimeter substantially relaxes the operational requirements for the precessing-polarization-based searches developed in Ref. [46].

10.3 Bunch-resolved polarization history

The combination of the three-channel pickup geometry and the matched-filter analysis of Sec. 6.6 delivers a distinctive operational output: the time evolution of the full polarization vector $\vec{P}(t)$ for every bunch in the stored beam, continuously throughout the fill, without disturbing the beam. We illustrate this output with a synthetic but representative simulation of an eight-hour flattop fill ($N_b = 1160$ bunches, $T_{\text{avg}} = 180$ s averaging window per bunch and per time bin), in which P_y undergoes a slow T_1 -type decay, residual P_x and P_z follow few-percent drifts over the fill, and two localized anomalous events affect distinct subsets of bunches.

Fig. 17 shows the three resulting heatmaps. Panel (a) tracks the dominant stored vertical component P_y and exhibits the expected slow fill-averaged decay. Panels (b) and (c) track the residual transverse and longitudinal components and reveal features that no fill-averaged polarimeter can resolve: bunch-to-bunch variation in the steady-state offsets, slow drifts of the offsets over hours, and the two localized anomalous events as compact patches against the smoother background. Panel (c) in particular constitutes the first practical demonstration of bunch-resolved longitudinal polarization reconstruction; the same information is not currently extractable from any combination of pC and HJET measurements (see Sec. 10.2).

The achievable per-bunch precision is set by the SNR available to one specific bunch as a function of the averaging window T devoted to that bunch. Both the saddle and axial channels reach a 1% per-bunch precision in a few hundred seconds at flattop and in correspondingly shorter times at injection. Fig. 16 makes this explicit. The all-bunch matched filter (solid curves) uses every bunch passage and reaches $\delta P_{\text{min}} \sim 10^{-4}$ within seconds; the per-bunch




Figure 16: Per-bunch versus all-bunch matched-filter precision at flattop ($N_b = 1160$). Solid curves: all-bunch matched filter, which uses every bunch passage. Dashed curves: per-bunch matched filter, which uses only the passages of a single bunch and is therefore $\sqrt{N_b} \simeq 34$ worse for the same integration time. Blue: saddle channels (per-bunch $K = 8.13 \sqrt{s}$). Red: axial gradiometer (per-bunch $K = 5.83 \sqrt{s}$). The $\sqrt{N_b}$ vertical offset between the dashed and solid curves is the only factor by which bunch-resolved reconstruction is statistically penalized relative to fill-averaged measurement.

matched filter (dashed curves) uses only the passages of a single bunch and is therefore $\sqrt{N_b}$ worse for the same integration time. The $\sqrt{N_b} \simeq 34$ at flattop sets the vertical separation of the dashed and solid curves and is the only factor by which per-bunch reconstruction is statistically penalized relative to fill-averaged measurement. Selecting T_{avg} per bunch and per time bin is then a direct trade-off between bunch time-resolution and per-cell polarization precision: the 180s value used in Fig. 17 gives roughly 1% per-cell precision on the saddle channels and 1.5% on the axial channel at flattop, sufficient to resolve the illustrated drifts and anomalous events while still providing 160 independent time bins per fill.

11 Extensions and outlook

Although the quantitative treatment in this paper uses the polarized proton beam as its worked example, the concept is not specific to protons and extends naturally to other particle species and to other storage rings. Two directions are of particular interest: storage-ring searches for electric dipole moments, for which the JEDI collaboration has recently reported the first direct limit on the deuteron EDM at COSY [47], and self-polarized electron beams.

11.1 Storage-ring electric dipole moment searches

A frozen-spin storage ring stores light ions at the magic momentum $p_{\text{magic}} = mc/\sqrt{G} \approx 0.70 \text{ GeV}/c$ for protons, at which the in-plane $g-2$ precession is frozen and the spin tracks the velocity [18]. A nonzero electric dipole moment then tilts the spin slowly out of the ring




Figure 17: Bunch-resolved polarization history over an 8-hour EIC HSR fill at flattop. (a) Dominant component P_y . (b) Residual in-plane P_x . (c) Residual longitudinal P_z , the first-ever bunch-resolved measurement.

plane, so that a small vertical polarization $P_y(t)$ grows almost linearly in time: this growing out-of-plane projection is the EDM signal. This invariant-spin-axis tilt is precisely the observable that the JEDI collaboration recently determined for deuterons at COSY, combining a radio-frequency Wien filter, a Siberian snake, and an electron-cooler solenoid to set a first direct limit on the deuteron electric dipole moment, $|d_d| < 2.5 \times 10^{-17} e \text{ cm}$ at the 95% confidence level [47]. Counter-propagating clockwise and counter-clockwise beams share the same magnetic guide field but sample radial fields of opposite sign, so their difference cancels the dominant fake-EDM systematic associated with radial magnetic fields.

The SQUID readout is a natural match for this signal. The observable is a slowly growing out-of-plane moment, precisely the quantity read by the static transverse and axial channels of Secs. 8 and 9. Because the pickup is noninvasive and continuous, both counter-rotating beams can be read in real time, replacing the slow destructive scattering polarimetry otherwise used to follow the spin. The clockwise-minus-counter-clockwise difference isolates the EDM term, while the stable revolution frequency of a dedicated EDM ring enables the same coherent, phase-locked integration developed here. A quantitative treatment of the frozen-spin operating point, the dual-beam systematics, and the achievable EDM sensitivity is the subject of a dedicated follow-up paper.

11.2 Electron storage rings and self-polarized beams

The matched-filter reconstruction and the FID-echo cycle both rely on a stable revolution frequency, so that the turn-by-turn precession phase can be predicted and the per-bunch contributions summed coherently. Among electron machines, a storage ring with a stable f_{rev} is therefore the natural fit. For the EIC electron storage ring [48], the large electron magnetic moment ($\mu_e/\mu_p \approx 658$), together with a tighter pickup coil, yields orders of magnitude more flux per passage; since the integration time scales as $\Phi_{\text{pickup}}^{-2}$, the corresponding measurement times collapse to the millisecond range.

The electron spin dynamics differ in two respects. The spin tune is $\nu_s = G_e\gamma \approx 41$ at top energy, set directly by the beam energy with no Siberian snakes, so the half-integer working point used in the HSR does not apply. In addition, Sokolov-Ternov self-polarization [49] is fast and asymmetric between spin states: bunches anti-aligned with the guide field depolarize faster than the aligned ones, with a characteristic time of order minutes at 18 GeV and a steep E^{-5} energy dependence. Machines without a stable revolution period are not direct users of the coherent method: in a rapid-cycling synchrotron the ramp leaves no stable f_{rev} and the spin tune sweeps through resonances, leaving no coherent window, while a recirculating linac such as CEBAF has no revolution frequency at all and admits at best single-pass detection, aided by the large electron moment.

12 Summary

We have developed a quantitative design and feasibility analysis of a SQUID-based polarimeter for the EIC Hadron Storage Ring that exploits the six-snake $\nu_s = 1/2$ spin lattice to read out the macroscopic magnetic dipole moment of polarized bunches at the spin precession frequency $f_s = f_{\text{rev}}/2$. The pickup couples three local channels at a single station: a saddle cosine- θ loop sensitive to the horizontal transverse dipole component, a saddle sine- θ loop sensitive to the vertical component, and a coaxial axial gradiometer sensitive to the longitudinal component. The three channels are read out simultaneously by independent dc-SQUID front ends and reconstruct the full polarization vector (P_x, P_y, P_z) . The measurement proceeds as a sequence of phase-locked tip- π -echo-restore cycles synchronized to the SQUID reference oscillator, with the closing restore pulse returning the tipped in-plane component to the longitudinal stable spin axis so that the cumulative polarization cost across a full measurement is negligible.

At the canonical operating point of $\alpha = 30$ mrad, $P = 0.7$, $\sigma_{\nu_s}^{\text{eff}} = 10^{-3}$, and $S_{\Phi}^{1/2} = 0.4 \mu\Phi_0/\sqrt{\text{Hz}}$, the matched-filter sensitivities reach $K_x = K_y \simeq 1108 \sqrt{\text{s}}$ at injection and $\simeq 277 \sqrt{\text{s}}$ at flattop on the saddle channels, and $K_z \simeq 156.8 \sqrt{\text{s}}$ (injection) and $\simeq 198.4 \sqrt{\text{s}}$ (flattop) on the axial gradiometer. With residual transverse and longitudinal polarizations anchored on the RHIC pC tilt observations of Ref. [44], all static-mode integration times for 1% precision lie between ~ 17 ms (stored P_y at injection) and ~ 7.5 min (residual P_z at injection); at flattop, all three components reach 1% precision in under 30 s. Free-precession coherence is preserved over the canonical $\tau_{\text{coh}}^{\text{eff}} \simeq 2.04$ ms by phase-locked operation; longer integrations are realized through the synchronized π rephasing pulse delivered as $N \simeq 80$ per-turn pulses with per-pass integrated field $\int B_z dl|_{\pi, \text{per-pass}} \simeq 1.1 \text{ T m}$ at injection and 12.9 T m at flattop (Tab. 5).

Beyond the absolute calibration use case, the continuous and noninvasive character of the SQUID readout supports a class of measurements that no existing storage-ring polarimeter delivers. Per-bunch matched filtering on the saddle and axial channels reconstructs $\vec{P}(t)$ for every bunch throughout the fill, at a per-bunch precision set by $K/\sqrt{N_b}$. At flattop this is $\sim 8.1 \sqrt{\text{s}}$ on the saddle channels and $\sim 5.8 \sqrt{\text{s}}$ on the axial channel, sufficient to resolve few-percent bunch-by-bunch variations of all three components in averaging windows of order one minute. The longitudinal projection P_z , inaccessible to single-spin pC polarimetry as a matter of parity conservation and not routinely measured by HJET, becomes a routine monitoring

observable on every bunch and time bin (Sec. 10). The same readout provides a noninvasive diagnostic of the spin tune, its linewidth-derived spread, and the coherence envelope.

The matched-filter sensitivities quoted above derive from a conservative baseline operating point. Three modest hardware optimizations, each individually within current SQUID engineering practice, would compound to deliver substantially shorter integration times. From the scaling $K_{\text{template}} \propto (\eta \sqrt{N_{\text{SQUIDS}}})/S_{\Phi}^{1/2}$ [Eq. (83)], lowering the white flux-noise floor from $0.4 \mu\Phi_0/\sqrt{\text{Hz}}$ to $0.3 \mu\Phi_0/\sqrt{\text{Hz}}$ (via multi-stage SQUID arrays or operation at sub-K temperatures), improving the flux-transformer coupling efficiency from $\eta = 0.7$ to $\eta = 0.9$ (via optimized inductive matching of the pickup loop to the SQUID input coil), and doubling the number of independent SQUID channels per pickup from $N_{\text{SQUIDS}} = 4$ to 8, each contributes independently. The compound enhancement is $K_{\text{template}} \mapsto 2.4 K_{\text{template}}$, equivalent to a factor ~ 6 shorter integration time for the same target precision. Under this optimized scenario the matched-filter 1% integration times of Tab. 7 would shorten to ~ 3 s at injection and ~ 50 s at flattop, bounding the practical performance ceiling for the architecture proposed in this paper.

The same architecture applies without modification to deuteron and ^3He beams via species-specific spin-magnetic factors, with the storage-ring electric-dipole-moment and electron-beam extensions outlined in Sec. 11. Engineering details of the saddle-coil fabrication, the axial gradiometer mechanical layout, the common-mode-rejection budget, the SQUID front-end electronics, and cryogenic and shielding requirements are treated in a companion engineering paper. The present work establishes the principle, the sensitivity budget, and the operational concept, demonstrating that noninvasive, bunch-resolved, three-component polarimetry of the EIC proton beam is feasible with currently available SQUID technology.

Acknowledgments

The author thanks Helmut Soltner (Forschungszentrum Jülich, Germany) and Nikolai N. Nikolaev (L. D. Landau Institute for Theoretical Physics, Chernogolovka, Russia) for useful comments and discussions.

This manuscript has been authored by an employee of Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy. The publisher, by accepting the manuscript for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

- [1] F. Rathmann, A. Nass, K. O. Eyser, V. Shmakova, E. C. Aschenauer, G. Atoian, A. Cannavo, X. Chu, K. Hock, H. Huang, H. Lovelace, G. Mahler, N. N. Nikolaev, J. Ritter, G. Robert-Demolaize, V. Schoefer, P. Shanmuganathan, E. Shulga, H. Soltner, and Z. Zhang. Eliminating beam-induced depolarizing effects in the hydrogen jet target for high-precision proton beam polarimetry at the electron-ion collider. *Phys. Rev. Accel. Beams*, 29:021001, Feb 2026. URL: <https://link.aps.org/doi/10.1103/8nh5-163q>, doi:10.1103/8nh5-163q.
- [2] F. Rathmann, O. Eyser, M. Sangroula, P. Shanmuganathan, and V. Shmakova. Thermal and electromechanical response of ultra-thin carbon-strip polarimeter targets in relativistic bunched beams. Technical report, Brookhaven National Laboratory, Upton, New York 11973, USA, June 2026. Submitted as a BNL Technical Note (June 17, 2026).

- [3] Brookhaven National Laboratory / Jefferson Lab. EIC global system requirements, 2022. URL: <https://eic.jlab.org/Requirements/GLOBAL.html>.
- [4] V. Schoefer, E.C. Aschenauer, H. Huang, F. Mèot, V. Ptitsyn, and V. H. Ranjbar. Stable spin direction measurements at RHIC with polarized proton beams. Technical Report BNL-225293-2024-TECH C-A/AP/703, Brookhaven National Laboratory, January 2024.
- [5] F. Rathmann, B. von Przewoski, W. A. Dezarn, J. Doskow, M. Dzemidzic, W. Haeberli, J. G. Hardie, B. Lorentz, H. O. Meyer, P. V. Pancella, R. E. Pollock, T. Rinckel, F. Sperisen, and T. Wise. Complete angular distribution measurements of pp spin correlation parameters A_{xx} , A_{yy} , and A_{zz} and analyzing power A_y at 197.4 mev. *Phys. Rev. C*, 58:658–673, Aug 1998. URL: <https://link.aps.org/doi/10.1103/PhysRevC.58.658>, doi:10.1103/PhysRevC.58.658.
- [6] Ya. S. Derbenev. Rf-resonance beam polarimeter. In *The 11th International Symposium on High Energy Spin Physics*, volume 343 of *AIP Conference Proceedings*, pages 264–272. American Institute of Physics, 1995. doi:10.1063/1.48865.
- [7] M. Conte, M. Pusterla, and G. Pucacco. The stern-gerlach interaction between a traveling particle and a time varying magnetic field. 2000. [arXiv:physics/0003069](https://arxiv.org/abs/physics/0003069).
- [8] P. R. Cameron, M. Conte, N. D’Imperio, W. Franklin, D. A. Goldberg, A. Luccio, W. W. MacKay, M. Palazzi, M. Pusterla, K. Vetter, and T. Zwart. Design and test of a prototype cavity for a stern-gerlach polarimeter. Technical report, Brookhaven National Laboratory, Upton, NY, March 2005.
- [9] Richard Talman. Relativistic Stern-Gerlach deflection. *arXiv preprint*, arXiv:1611.03810, 2016. doi:10.48550/arXiv.1611.03810.
- [10] S. R. Mane. Relativistic Stern-Gerlach deflection: Hamiltonian formulation. *arXiv preprint*, arXiv:1611.07326, 2016. doi:10.48550/arXiv.1611.07326.
- [11] P. R. Cameron, A. U. Luccio, T. J. Shea, and D. A. Goldberg. Squids, snakes, and polarimeters: A new technique for measuring the magnetic moments of polarized beams. In *High Energy Spin Physics: 12th International Symposium*, volume 390 of *AIP Conference Proceedings*, pages 306–315. American Institute of Physics, 1997. doi:10.1063/1.51839.
- [12] Younggeun Kim, Themis Bowcock, Dmitry Budker, Giovanni Cantatore, Hooman Davoudiasl, Dmitry Denisov, Abhay Deshpande, Wolfram Fischer, Selcuk Haciomeroglu, Haixin Huang, David Kawall, Alexander Keshavarzi, On Kim, Ivan Koop, Valeri Lebedev, Jonathan Lee, William M. Morse, Cenap Ozben, Vincent Schoefer, Yannis K. Sermertzidis, Eleftherios Skordis, Edward Stephenson, Vladimir Tishchenko, Nicholas Tsoupas, Graziano Venanzoni, Joost Vosseveld, and Peter Winter. Continuous coherent spin-frequency metrology in storage rings via resonant beam-driven detection. 2026. [arXiv:2605.23013](https://arxiv.org/abs/2605.23013).
- [13] D. Eversmann, V. Hejny, F. Hinder, A. Kacharava, J. Pretz, F. Rathmann, M. Rosenthal, F. Trinkel, S. Andrianov, W. Augustyniak, Z. Bagdasarian, M. Bai, W. Bernreuther, S. Bertelli, M. Berz, J. Bsaisou, S. Chekmenev, D. Chiladze, G. Ciullo, M. Contalbrigo, J. de Vries, S. Dymov, R. Engels, F. M. Esser, O. Felden, M. Gaisser, R. Gebel, H. Glückler, F. Goldenbaum, K. Grigoryev, D. Grzonka, G. Guidoboni, C. Hanhart, D. Heberling, N. Hempelmann, J. Hetzel, R. Hipple, D. Hölscher, A. Ivanov, V. Kamerdzhev, B. Kamys, I. Keshelashvili, A. Khoukaz, I. Koop, H.-J. Krause, S. Krewald, A. Kulikov, A. Lehrach, P. Lenisa, N. Lomidze, B. Lorentz, P. Maanen, G. Macharashvili, A. Magiera, R. Maier, K. Makino, B. Mariański, D. Mchedlishvili, Ulf-G. Meißner, S. Mey, A. Nass, G. Natour,

- N. Nikolaev, M. Nioradze, A. Nogga, K. Nowakowski, A. Pesce, D. Prasuhn, J. Ritman, Z. Rudy, A. Saleev, Y. Semertzidis, Y. Senichev, V. Shmakova, A. Silenko, J. Slim, H. Soltner, A. Stahl, R. Stassen, M. Statera, E. Stephenson, H. Stockhorst, H. Straatmann, H. Ströher, M. Tabidze, R. Talman, P. Thörngren Engblom, A. Trzciński, Yu. Uzikov, Yu. Valdau, E. Valetov, A. Vassiliev, C. Weidemann, C. Wilkin, A. Wirzba, A. Wrońska, P. Wüstner, M. Zakrzewska, P. Zuprański, and D. Zyuzin. New method for a continuous determination of the spin tune in storage rings and implications for precision experiments. *Phys. Rev. Lett.*, 115:094801, Aug 2015. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.094801>, doi:10.1103/PhysRevLett.115.094801.
- [14] G. Guidoboni et al. Connection between zero chromaticity and long in-plane polarization lifetime in a magnetic storage ring. *Phys. Rev. Accel. Beams*, 21:024201, 2018. doi:10.1103/PhysRevAccelBeams.21.024201.
- [15] A. Saleev, N. N. Nikolaev, F. Rathmann, W. Augustyniak, Z. Bagdasarian, M. Bai, L. Barion, M. Berz, S. Chekmenov, G. Ciullo, S. Dymov, D. Eversmann, M. Gaisser, R. Gebel, K. Grigoryev, D. Grzonka, G. Guidoboni, D. Heberling, V. Hejny, N. Hempelmann, J. Hetzel, F. Hinder, A. Kacharava, V. Kamerdzhev, I. Keshelashvili, I. Koop, A. Kulikov, A. Lehrach, P. Lenisa, N. Lomidze, B. Lorentz, P. Maanen, G. Macharashvili, A. Magiera, D. Mchedlishvili, S. Mey, F. Müller, A. Nass, A. Pesce, D. Prasuhn, J. Pretz, M. Rosenthal, V. Schmidt, Y. Semertzidis, Y. Senichev, V. Shmakova, A. Silenko, J. Slim, H. Soltner, A. Stahl, R. Stassen, E. Stephenson, H. Stockhorst, H. Ströher, M. Tabidze, G. Tagliente, R. Talman, P. Thörngren Engblom, F. Trinkel, Yu. Uzikov, Yu. Valdau, E. Valetov, A. Vassiliev, C. Weidemann, A. Wrońska, P. Wüstner, P. Zuprański, and M. Zurek. Spin tune mapping as a novel tool to probe the spin dynamics in storage rings. *Phys. Rev. Accel. Beams*, 20:072801, Jul 2017. URL: <https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.072801>, doi:10.1103/PhysRevAccelBeams.20.072801.
- [16] N. Hempelmann et al. Phase measurement for driven spin oscillations in a storage ring. *Phys. Rev. Accel. Beams*, 21:042002, 2018. doi:10.1103/PhysRevAccelBeams.21.042002.
- [17] F. Rathmann, N. N. Nikolaev, and J. Slim. Spin dynamics investigations for the electric dipole moment experiment. *Phys. Rev. Accel. Beams*, 23:024601, 2020. doi:10.1103/PhysRevAccelBeams.23.024601.
- [18] F. Abusaif et al. *Storage ring to search for electric dipole moments of charged particles: Feasibility study*. CERN, Geneva, 6 2021. arXiv:1912.07881, doi:10.23731/CYRM-2021-003.
- [19] J. Slim, F. Rathmann, A. Andres, V. Hejny, A. Nass, A. Kacharava, P. Lenisa, N. N. Nikolaev, J. Pretz, A. Saleev, V. Shmakova, H. Soltner, F. Abusaif, A. Aggarwal, A. Aksentev, B. Alberdi, L. Barion, I. Bekman, M. Beyß, C. Böhme, B. Breitzkreutz, N. Canale, G. Ciullo, S. Dymov, N.-O. Fröhlich, R. Gebel, M. Gaisser, K. Grigoryev, D. Grzonka, J. Hetzel, O. Javakhishvili, V. Kamerdzhev, S. Karanth, I. Keshelashvili, A. Kononov, K. Laihem, A. Lehrach, N. Lomidze, B. Lorentz, G. Macharashvili, A. Magiera, D. Mchedlishvili, A. Melnikov, F. Müller, A. Pesce, V. Poncza, D. Prasuhn, D. Shergelashvili, N. Shurkhno, S. Siddique, A. Silenko, R. Stassen, E. J. Stephenson, H. Ströher, M. Tabidze, G. Tagliente, Y. Valdau, M. Vitz, T. Wagner, A. Wirzba, A. Wrońska, P. Wüstner, and M. Żurek. Proof-of-principle demonstration of a pilot bunch comagnetometer in a stored beam. *Phys. Rev. Res.*, 7:023257, Jun 2025. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.7.023257>, doi:10.1103/PhysRevResearch.7.023257.
- [20] V. Anastassopoulos, S. Andrianov, R. Baartman, S. Baessler, M. Bai, J. Benante, M. Berz, M. Blaskiewicz, T. Bowcock, K. Brown, B. Casey, M. Conte, J. D. Crnkovic, N. D’Imperio, G. Fanourakis, A. Fedotov, P. Fierlinger, W. Fischer, M. O. Gaisser,

- Y. Giomataris, M. Grosse-Perdekamp, G. Guidoboni, S. Hacıömeroğlu, G. Hoffstaetter, H. Huang, M. Incagli, A. Ivanov, D. Kawall, Y. I. Kim, B. King, I. A. Koop, D. M. Lazarus, V. Lebedev, M. J. Lee, S. Lee, Y. H. Lee, A. Lehrach, P. Lenisa, P. Levi Sandri, A. U. Lucio, A. Lyapin, W. MacKay, R. Maier, K. Makino, N. Malitsky, W. J. Marciano, W. Meng, F. Meot, E. M. Metodiev, L. Miceli, D. Moricciani, W. M. Morse, S. Nagaitsev, S. K. Nayak, Y. F. Orlov, C. S. Ozben, S. T. Park, A. Pesce, E. Petrakou, P. Pile, B. Podobedov, V. Polychronakos, J. Pretz, V. Ptitsyn, E. Ramberg, D. Raparia, F. Rathmann, S. Rescia, T. Roser, H. Kamal Sayed, Y. K. Semertzidis, Y. Senichev, A. Sidorin, A. Silenko, N. Simos, A. Stahl, E. J. Stephenson, H. Ströher, M. J. Syphers, J. Talman, R. M. Talman, V. Tishchenko, C. Touramanis, N. Tsoupas, G. Venanzoni, K. Vetter, S. Vlassis, E. Won, G. Zavattini, A. Zelenski, and K. Zioutas. A storage ring experiment to detect a proton electric dipole moment. *Review of Scientific Instruments*, 87(11):115116, 11 2016. [arXiv:https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/1.4967465/15852296/115116_1_online.pdf](https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/1.4967465/15852296/115116_1_online.pdf), doi:10.1063/1.4967465.
- [21] C. Liu, J. Kewisch, H. Huang, and M. Minty. Minimization of spin tune spread for preservation of spin polarization at rhic. *Phys. Rev. Accel. Beams*, 22:061002, 2019. doi:10.1103/PhysRevAccelBeams.22.061002.
- [22] Y. Luo and W. Fischer. Beam–beam observations in the rhic, 2014. [arXiv:1410.5936](https://arxiv.org/abs/1410.5936).
- [23] E. Hamwi et al. Snake matching the eic’s hadron storage ring. In *Proceedings of IPAC’24*, page TUPS33, 2024. doi:10.18429/JACoW-IPAC2024-TUPS33.
- [24] Eiad Hamwi and Georg H. Hoffstaetter. Polarization transmission in the Electron-Ion Collider’s Hadron Storage Ring. 2025. URL: <https://arxiv.org/abs/2509.18558>, [arXiv:2509.18558](https://arxiv.org/abs/2509.18558).
- [25] V. Ptitsyn and J. S. Berg. EIC hadron spin rotators. In *Proc. 13th Int. Particle Accelerator Conf. (IPAC’22)*, page WEPOST020, 2022. URL: <https://proceedings.jacow.org/ipac2022/papers/wepost020.pdf>.
- [26] N. J. Stone. Table of nuclear magnetic dipole and electric quadrupole moments. *At. Data Nucl. Data Tables*, 111–112:1–28, 2016. doi:10.1016/j.adt.2015.12.002.
- [27] P. J. Mohr, D. B. Newell, and B. N. Taylor. CODATA recommended values of the fundamental physical constants: 2014. *Rev. Mod. Phys.*, 88:035009, 2016. doi:10.1103/RevModPhys.88.035009.
- [28] V. Bargmann, L. Michel, and V. L. Telegdi. Precession of the polarization of particles moving in a homogeneous electromagnetic field. *Phys. Rev. Lett.*, 2:435, 1959. doi:10.1103/PhysRevLett.2.435.
- [29] Taskeshi Fukuyama and Alexander J. Silenko. Derivation of generalized thomas–bargmann–michel–telegdi equation for a particle with electric dipole moment. *International Journal of Modern Physics A*, 28(29):1350147, 2013. [arXiv:https://doi.org/10.1142/S0217751X13501479](https://doi.org/10.1142/S0217751X13501479), doi:10.1142/S0217751X13501479.
- [30] V. Ptitsyn and J. S. Berg. EIC hadron spin rotators. In *Proc. 13th International Particle Accelerator Conference (IPAC2022)*, pages 1734–1736, Geneva, Switzerland, 2022. JACoW Publishing. Paper WEPOST020. URL: <https://proceedings.jacow.org/ipac2022/papers/wepost020.pdf>, doi:10.18429/JACoW-IPAC2022-WEPOST020.
- [31] Alston S. Householder. Unitary triangularization of a nonsymmetric matrix. *J. ACM*, 5(4):339–342, October 1958. doi:10.1145/320941.320947.

- [32] S. Y. Lee and E. D. Courant. Tolerance of imperfections in high-energy circular accelerators for polarized protons. *Phys. Rev. D*, 41:292–302, Jan 1990. URL: <https://link.aps.org/doi/10.1103/PhysRevD.41.292>, doi:10.1103/PhysRevD.41.292.
- [33] Eiad Hamwi. *Spin Polarization, Modeling, and Beam Control in Hadron Accelerators: For RHIC and the EIC*. Ph.d. dissertation, Cornell University, December 2025.
- [34] S. Y. Lee. *Spin Dynamics and Snakes in Synchrotrons*. World Scientific, Singapore, 1997. doi:10.1142/3377.
- [35] J. Slim, R. Gebel, D. Heberling, F. Hinder, D. Hölscher, A. Lehrach, B. Lorentz, S. Mey, A. Nass, F. Rathmann, L. Reifferscheidt, H. Soltner, H. Straatmann, F. Trinkel, and J. Wolters. Electromagnetic simulation and design of a novel waveguide rf wien filter for electric dipole moment measurements of protons and deuterons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 828:116–124, 2016. URL: <https://www.sciencedirect.com/science/article/pii/S0168900216303710>, doi:10.1016/j.nima.2016.05.012.
- [36] G. Guidoboni et al. How to reach a thousand-second in-plane polarization lifetime with 0.97-GeV/c deuterons in a storage ring. *Phys. Rev. Lett.*, 117:054801, 2016. doi:10.1103/PhysRevLett.117.054801.
- [37] E. L. Hahn. Spin echoes. *Phys. Rev.*, 80:580–594, Nov 1950. URL: <https://link.aps.org/doi/10.1103/PhysRev.80.580>, doi:10.1103/PhysRev.80.580.
- [38] Alessandro Bravar. Proton polarimetry at RHIC. In *AIP Conference Proceedings*, volume 792, pages 1039–1042, 2005. doi:10.1063/1.2122177.
- [39] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Phys. Rev.*, 94:630–638, May 1954. URL: <https://link.aps.org/doi/10.1103/PhysRev.94.630>, doi:10.1103/PhysRev.94.630.
- [40] S. Meiboom and D. Gill. Modified spin-echo method for measuring nuclear relaxation times. *Review of Scientific Instruments*, 29(8):688–691, 08 1958. arXiv:https://pubs.aip.org/aip/rsi/article-pdf/29/8/688/19287064/688_1_online.pdf, doi:10.1063/1.1716296.
- [41] D. Drung, C. Abmann, J. Beyer, A. Kirste, M. Peters, F. Ruede, and Th. Schurig. Highly sensitive and easy-to-use squid sensors. *IEEE Transactions on Applied Superconductivity*, 17(2):699–704, 2007. doi:10.1109/TASC.2007.897403.
- [42] D. Xu et al. EIC beam dynamics challenges. Technical Report Fermilab-Conf-22-682-AD, Fermilab, 2022. URL: <https://lss.fnal.gov/archive/2022/conf/fermilab-conf-22-682-ad.pdf>.
- [43] R. L. Fagaly. Superconducting quantum interference device instruments and applications. *Rev. Sci. Instrum.*, 77:101101, 2006. doi:10.1063/1.2354545.
- [44] W. B. Schmidke. RHIC polarization for Runs 9–17. Technical Report BNL-209057-2018-TECH, C-A/AP/609, Brookhaven National Laboratory, September 2018. URL: <https://www.osti.gov/servlets/purl/1473643>.
- [45] A. A. Poblaguev, A. Zelenski, E. Aschenauer, G. Atoian, K. O. Eyser, H. Huang, Y. Makdisi, W. B. Schmidke, I. Alekseev, D. Svirida, and N. H. Buttimore. Precision small scattering angle measurements of elastic proton-proton single and double spin analyzing powers at the rhic hydrogen jet polarimeter. *Phys. Rev. Lett.*, 123:162001,

- Oct 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.162001>, doi: 10.1103/PhysRevLett.123.162001.
- [46] N.N. Nikolaev, F. Rathmann, A.J. Silenko, and Yu.N. Uzikov. New approach to search for parity-even and parity-odd time-reversal violation beyond the standard model in a storage ring. *Physics Letters B*, 811:135983, 2020. URL: <https://www.sciencedirect.com/science/article/pii/S0370269320307863>, doi:10.1016/j.physletb.2020.135983.
 - [47] A. Andres, V. Hejny, A. Nass, N. N. Nikolaev, J. Pretz, F. Rathmann, V. Shmakova, J. Slim, et al. First experimental limit on the permanent electric dipole moment of the deuteron. *Phys. Rev. Lett.*, 136(24):241801, 2026. [arXiv:2602.20828](https://arxiv.org/abs/2602.20828), doi:10.1103/ns3s-ld4k.
 - [48] R. Abdul Khalek, A. Accardi, J. Adam, et al. Science requirements and detector concepts for the Electron-Ion Collider: EIC Yellow Report. *Nucl. Phys. A*, 1026:122447, 2022. [arXiv:2103.05419](https://arxiv.org/abs/2103.05419), doi:10.1016/j.nuclphysa.2022.122447.
 - [49] A. A. Sokolov and I. M. Ternov. On polarization and spin effects in the theory of synchrotron radiation. *Sov. Phys. Dokl.*, 8:1203, 1964.
 - [50] M. Bai, T. Roser, C. Dawson, J. Kewisch, Y. Makdisi, P. Oddo, C. Pai, and P. Pile. Rhic spin flipper commissioning results. In *Proceedings of IPAC2012*, pages 1302–1304, New Orleans, Louisiana, USA, 2012. <https://proceedings.jacow.org/IPAC2012/papers/tuppc057.pdf>. URL: <https://proceedings.jacow.org/IPAC2012/papers/tuppc057.pdf>.
 - [51] H. Huang, J. Kewisch, C. Liu, A. Marusic, W. Meng, F. Méot, P. Oddo, V. Ptitsyn, V. Ranjbar, and T. Roser. High spin-flip efficiency at 255 GeV for polarized protons in a ring with two full Siberian snakes. *Phys. Rev. Lett.*, 120:264804, 2018. doi:10.1103/PhysRevLett.120.264804.

A Spin manipulation with an RF Wien filter

The longitudinal kicker introduced in Sec. 5 provides the tipping, π , and restore pulses required for the SQUID measurement cycle. The same spin-manipulation functions can in principle be realized with an RF Wien filter operated at the spin-resonance frequency, with the operational advantage that the Wien-filter configuration preserves the closed orbit by construction. A recent COSY experiment by Slim *et al.* [19] demonstrated that fast rf switches integrated into the Wien-filter driving circuit allow the device to be gated bunch by bunch. In their two-bunch demonstration with polarized deuterons, the signal bunch received the resonant RF and underwent multiple continuous spin flips with a fitted per-flip spin-flip efficiency $\varepsilon_{\text{SF}} = 0.9954 \pm 0.0037$, while the pilot bunch was shielded from the RF during its passage through the device with a gate efficiency $\varepsilon_{\text{gate}} = 0.9921 \pm 0.0136$ consistent with unity, leaving its idle precession unperturbed. The two efficiencies quantify distinct aspects of the demonstration: ε_{SF} measures how cleanly the resonant RF reverses the polarization of the driven bunch, while $\varepsilon_{\text{gate}}$ measures how completely the masked bunch is isolated from the RF field during the gated interval. That experiment was performed at the natural deuteron spin tune $\nu_s = \gamma G_d$ (with $G_d \simeq -0.143$) in a ring with no Siberian snakes; the two-branch resonance degeneracy that arises at $\nu_s = 1/2$ and which is the subject of Sec. A.1 below is therefore specific to the EIC HSR six-snake lattice considered here. A bunch-gated RF Wien filter pairs naturally with the bunch-resolved SQUID readout developed in Sec. 10: the combination of per-bunch (P_x, P_y, P_z) measurement and per-bunch polarization reversal provides the natural systematic-control mode for the precision spin-physics applications, including the storage-ring electric-dipole-moment searches and related beyond-the-Standard-Model studies treated in a dedicated follow-up paper.

Having established the spin lattice of the LC configuration, we now consider its response to an external RF perturbation.

To investigate controlled spin manipulation in the LC configuration introduced in Sec. 3.4.6, we consider the effect of a transverse RF Wien filter providing a horizontal magnetic field. The spin transport is described using the one-turn map of Eq. (14) together with the sector-wise rotations of Eqs. (13) and (12), with the RF Wien filter treated as a localized time-dependent perturbation.

The RF Wien filter is operated at the spin-resonance frequency,

$$f_{\text{RF}} = (k + \nu_s) f_{\text{rev}}, \quad k \in \mathbb{Z}, \quad (117)$$

such that the perturbation is phase-locked to the spin precession. For the present lattice with full Siberian snakes, $\nu_s = \frac{1}{2}$, and the resonance condition reduces to

$$f_{\text{RF}} = \left(k + \frac{1}{2}\right) f_{\text{rev}}. \quad (118)$$

Under these conditions, the spin motion can be described in a rotating frame in which the resonant component of the RF field appears as an effective static transverse field. A continuous sequence of small spin rotations then adds coherently, leading to a gradual reorientation of the polarization.

The driven motion can be expressed in terms of a turn-dependent spin map,

$$\mathbf{M}_{\text{tot}}(n) = \mathbf{M}_{\text{ring}} \mathbf{R}_{\text{WF}}(\phi_n), \quad (119)$$

with RF phase

$$\phi_n = 2\pi \frac{f_{\text{RF}}}{f_{\text{rev}}} n + \phi_0. \quad (120)$$

Tracking the spin turn-by-turn allows one to monitor the projection onto the invariant spin axis,

$$P_n = \vec{S}_n \cdot \vec{n}_0. \quad (121)$$

On resonance, the RF field generates a coherent buildup of spin rotation, leading to a continuous change of P_n and, for sufficient interaction time, a reversal of the polarization. Off resonance, the phase coherence is lost and the oscillation amplitude is reduced, preventing full spin flip.

To illustrate the resonant mechanism, we perform a numerical tracking calculation in which a small RF-induced rotation is applied once per turn. Starting from a vertically polarized beam, $\vec{P} = (0, 1, 0)$, the spin is propagated using the same one-turn map as in the unperturbed case, supplemented by the RF kick with time-dependent phase.

This simplified model represents the idealized resonant component of the spin motion and assumes a well-defined effective rotation axis in the rotating frame.

The result is shown in Fig. 18. The vertical polarization component P_y decreases continuously and changes sign, demonstrating a full reversal of the spin direction. The longitudinal component P_z oscillates while the radial component P_x remains essentially constant, consistent with a rotation about the radial axis in the rotating frame. The observed evolution confirms that, under resonant conditions, an RF Wien filter can provide coherent spin manipulation, with polarization reversal arising from the accumulation of small rotations over many turns.

A.1 Field structure

The origin of the two resonance conditions can be understood by expressing a linearly oscillating RF magnetic field in terms of its rotating components. For a transverse field,

$$B_x(t) = B_0 \cos(\omega t) = \frac{B_0}{2} (e^{+i\omega t} + e^{-i\omega t}), \quad (122)$$




Figure 18: Turn-by-turn evolution of the spin components under the action of a resonant RF Wien filter. The vertical polarization P_y (black) is gradually reversed, while the longitudinal component P_z (red) oscillates and the radial component P_x (blue) stays essentially constant. The RF field is synchronized with the spin precession according to $f_{\text{RF}} = (k + \nu_s)f_{\text{rev}}$, enabling coherent buildup of the spin rotation.

which shows explicitly that the field consists of two circularly rotating components with frequencies $+\omega$ and $-\omega$, corresponding to opposite rotation senses.

The spin precesses with angular frequency

$$\omega_s = 2\pi\nu_s f_{\text{rev}}. \quad (123)$$

Resonance occurs when one of the two rotating components becomes stationary in the rotating spin frame. This leads to the conditions

$$\omega - \omega_s = k \omega_{\text{rev}} \quad \text{or} \quad -\omega - \omega_s = k \omega_{\text{rev}}, \quad (124)$$

where $\omega_{\text{rev}} = 2\pi f_{\text{rev}}$. Dividing by ω_{rev} gives

$$\frac{\omega}{\omega_{\text{rev}}} = k + \nu_s \quad \text{or} \quad \frac{\omega}{\omega_{\text{rev}}} = k + (1 - \nu_s). \quad (125)$$

In terms of the RF tune $Q_{\text{osc}} = f_{\text{RF}}/f_{\text{rev}}$, this corresponds to the two resonance conditions

$$\nu_s = Q_{\text{osc}} \quad \text{and} \quad \nu_s = 1 - Q_{\text{osc}}. \quad (126)$$

Away from $\nu_s = \frac{1}{2}$, the two branches are separated and only one of the two rotating components is resonant. For the present lattice with $\nu_s = \frac{1}{2}$, however, the two conditions become degenerate. In this case, both the $+\omega$ and $-\omega$ components are simultaneously resonant, and a purely linearly oscillating field does not define a unique rotation sense in the rotating frame.

A well-defined coherent spin rotation is obtained only if the RF field provides a single rotating component. This can be achieved by combining two orthogonal magnetic-field components with a relative phase shift of 90° ,

$$B_x(t) \propto \cos(\omega t), \quad B_z(t) \propto \sin(\omega t), \quad (127)$$

so that

$$B_x(t) + iB_z(t) \propto e^{+i\omega t}, \quad (128)$$

which corresponds to selecting only one of the two rotation senses. In this case, only a single resonance branch is excited and the spin experiences a well-defined rotation in the rotating frame. A different implementation of this rotation-sense selection, using multiple AC dipoles arranged in two closed vertical orbit bumps interleaved with four DC spin rotators rather than a single rotating-field device, was developed for the RHIC spin flipper. First commissioning measurements with this design demonstrated significant suppression of the unwanted resonance branch though without full cancellation [50]; subsequent operation of the device in the snake ring at 255 GeV achieved a high single-flip efficiency [51], although routine operational use for systematic-error reduction during physics running has remained limited.

A.2 Idealized spin reversal in the rotating frame

Once the rotation-sense degeneracy is broken and a single rotating component is selected, the driven spin motion takes a simple closed form. In the frame rotating at the spin-precession frequency $\omega_s = 2\pi\nu_s f_{\text{rev}}$, the resonant component of the combined Wien-filter/longitudinal field is stationary and acts as a static effective field along a fixed axis $\vec{e}$, while the counter-rotating component averages to zero over many turns. For the field phasing $B_x(t) \propto \cos\omega t$, $B_z(t) \propto \sin\omega t$ of Sec. A.1, this effective axis lies along the radial direction, $\vec{e} = \vec{e}_x$.

The spin then undergoes uniform rotation about $\vec{e}_x$,

$$\vec{P}(N) = \mathbf{R}_{\vec{e}_x}(\Theta(N)) \vec{P}(0), \quad \Theta(N) = N \delta\theta, \quad (129)$$

where N is the number of turns the combined field is gated on and $\delta\theta$ is the spin rotation accumulated per turn. The per-turn rotation angle is set by the integrated transverse field through the same T-BMT relation as the tipping kick from Eq. (38),

$$\delta\theta = \frac{(1+G)}{B\rho} \int B_\perp dl, \quad (130)$$

with $B\rho = p/e$ the magnetic rigidity and $\int B_\perp dl$ the integrated resonant transverse field per revolution.

Writing the initial polarization as $\vec{P}(0) = (P_x, P_y, P_z)$, the rotation about $\vec{e}_x$ leaves the radial projection unchanged and rotates the (P_y, P_z) pair,

$$\mathbf{R}_{\vec{e}_x}(\Theta) \vec{P} = (P_x, P_y \cos \Theta + P_z \sin \Theta, -P_y \sin \Theta + P_z \cos \Theta). \quad (131)$$

A clean reversal of the stored vertical polarization is obtained at $\Theta = \pi$,

$$\mathbf{R}_{\vec{e}_x}(\pi) (P_x, P_y, P_z) = (P_x, -P_y, -P_z), \quad (132)$$

so that $P_y \rightarrow -P_y$ with its magnitude preserved, the longitudinal component is likewise reversed, and the radial component on the rotation axis is conserved. During the buildup, P_z oscillates as Θ advances from 0 to π while P_y reverses monotonically, as shown in Fig. 18; the radial component P_x remains constant throughout. A single fixed-axis rotation therefore reverses the two components transverse to $\vec{e}_x$ but not the on-axis component, which is the operationally relevant behavior here: the goal is to reverse the stored vertical polarization P_y , and hence, through the interaction-region spin rotators that map the stable vertical direction onto the longitudinal direction at the interaction point, to reverse the sign of the longitudinal polarization delivered to the ePIC detector for a selected bunch. Reversing P_y on a single bunch in this way provides a powerful handle on the systematics of spin-dependent observables, since bunches of opposite longitudinal polarization sign traverse the same detector under otherwise identical conditions.

The reversal is dosed by the accumulated angle $\Theta = N \delta\theta$. The combined field is gated on for the number of turns required to reach $\Theta = \pi$,

$$N_\pi = \frac{\pi}{\delta\theta} = \frac{\pi B\rho}{(1+G) \int B_\perp dl}, \quad (133)$$

which fixes the relation between the integrated field per turn and the number of turns of operation: a larger integrated field per turn reaches the reversal in fewer turns, and the product $N_\pi (1+G) \int B_\perp dl / B\rho = \pi$ is held fixed. Operating for $\Theta > \pi$ rotates the polarization back past the reversal, while $\Theta < \pi$ leaves a partial flip with a residual in-plane projection; the reversal fidelity is therefore set by how accurately $N \delta\theta$ is tuned to π , the same quantity characterized by the per-flip efficiency ε_{SF} reported in the COSY bunch-gated demonstration [19].

A.3 Implementation options

In contrast to an RF dipole flipper, which necessarily generates a transverse Lorentz force and drives coherent betatron oscillations, the RF Wien filter preserves the beam orbit by construction,

$$\vec{F}_L = q(\vec{E} + \vec{v} \times \vec{B}) = 0. \quad (134)$$

For an RF dipole, the driven orbit motion couples the spin dynamics to the transverse beam motion, leading to amplitude-dependent resonance strengths and phase decoherence across the beam, which limits the achievable spin-flip efficiency, as observed in RHIC spin-flipper studies [50]. By eliminating orbit excitation, the RF Wien filter decouples spin manipulation from the transverse dynamics and provides a more uniform and controlled spin rotation across the beam ensemble.

Three practical implementation options can be considered. The first is to combine a conventional RF Wien filter providing the transverse field $B_x(t)$ with a second device generating a longitudinal RF magnetic field $B_z(t)$, phase-locked with a relative phase shift of 90° . The second is to develop an integrated RF structure that generates both $B_x(t)$ and $B_z(t)$ within a single device. The third is to operate the RF Wien filter in a pulsed, phase-synchronous mode:

1. **RF Wien filter plus longitudinal-field device.** A conventional RF Wien filter provides the transverse magnetic field $B_x(t)$ together with the compensating electric field $E_y(t) = -v_z B_x(t)$ that enforces $\vec{F}_L = q(\vec{E} + \vec{v} \times \vec{B}) = 0$ on the nominal closed orbit, so that the device acts purely on the spin without exciting orbit motion. A second device, phase-locked to the first with a relative phase of 90° , provides the longitudinal magnetic field $B_z(t)$. Because the longitudinal field is parallel to the beam velocity, $\vec{v} \times \vec{B}_z = 0$ and no compensating electric field is required for the second device: the orbit is preserved automatically. The combination

$$B_x(t) \propto \cos(\omega t), \quad B_z(t) \propto \sin(\omega t)$$

realizes a single rotating field component $e^{+i\omega t}$ in the horizontal plane, which selects one of the two resonance branches and produces a unique direction of cumulative spin rotation, as required to break the $\nu_s = 1/2$ degeneracy of Sec. A.1. Each device is a well-characterized stand-alone component; the cost of this configuration is the mechanical layout and phase-stability requirement of operating two synchronized RF assemblies. A natural hardware-reuse possibility is to operate the longitudinal kicker of Sec. 5 as the second device, driven in continuous-wave mode at the spin-resonance frequency rather than in the pulsed mode used for the tip- π -restore cycle.

2. **Integrated multi-electrode structure:** A dedicated RF structure can be designed to generate both $B_x(t)$ and $B_z(t)$ within a single device, for example using an extended multi-electrode geometry with independently driven field components.
3. **Pulsed RF Wien filter.**

A.4 Pulsed-mode RF Wien filter

In the pulsed mode, the RF Wien filter is operated using short, phase-synchronized excitations rather than a continuous wave. In continuous operation, as discussed in Sec. A.1, a linearly oscillating field $B_x(t) \propto \cos(\omega t)$ contains two counter-rotating components $e^{+i\omega t}$ and $e^{-i\omega t}$. At $\nu_s = \frac{1}{2}$, both components satisfy the resonance condition $f_{\text{RF}} = (k + \nu_s)f_{\text{rev}}$, and therefore no unique rotation sense is defined.

In contrast, the pulsed scheme samples the RF field at discrete times $t_n = n/f_{\text{rev}}$ and enforces a controlled phase relation with respect to the spin precession. The phase of the excitation,

$$\phi_n = 2\pi Q_{\text{osc}}n + \phi_0, \quad Q_{\text{osc}} = k + \nu_s, \quad (135)$$

is matched to the spin phase advance such that the kicks add coherently from turn to turn. Under this condition, only the component co-rotating with the spin motion contributes to the cumulative rotation, while the counter-rotating component does not build up.

The spin manipulation is thus realized as a sequence of discrete kicks applied once per turn, whose coherent sum produces a net rotation over many turns. In this way, the pulsed scheme effectively selects a single resonance branch in the time domain, without requiring a physically rotating magnetic field or additional field components such as $B_z(t)$.

The numerical example shown in Fig. 18 corresponds to this idealized situation in which a single effective rotation sense is selected. The figure therefore illustrates the coherent spin evolution obtained when either the $+\omega$ or the $-\omega$ branch is effectively isolated. A purely linearly oscillating RF field at $\nu_s = \frac{1}{2}$ does not automatically realize this condition, and additional measures such as the pulsed operation described here are required.