New Journal of Physics
The open access journal at the forefront of physics
Deutsche Physikalische Gesellschaft DPG
IOP Institute of Physics
Published in partnership with: Deutsche Physikalische Gesellschaft and the Institute of Physics

PAPER

OPEN ACCESS

RECEIVED
4 February 2026

REVISED
18 June 2026

ACCEPTED FOR PUBLICATION
20 August 2026

PUBLISHED
1 September 2026

Diagnosing antihydrogen energy distributions using spatio-temporal annihilation patterns

R Ahemeh[1], R Akbari[1], L O de Araujo Azevedo[2], C J Baker[3], W Bertsche[4,5], N M Bhatt[3], G Bonomi[6], A Capra[7], I Carli[7], C L Cesar[2], M Charlton[3], A Cridland Mathad[8], A Del Vincio[6], D Duque Quiceno[1,7], S Eriksson[3], A Evans[1,7], J Fajans[9,21], T Friesen[10], M C Fujiwara[7], L M Golino[3], M B Gomes Gonçalves[3], J S Hangst[11], M E Hayden[12], D Hodgkinson[8,9,*], C A Isaac[3], A J Uribe Jimenez[7,10], S A Jones[13], S Jonsell[14], N Madsen[3], V R Marshall[11], J T K McKenna[4], T Momose[1,7,15], J Nauta[8], A Olin[7,20], A N Oliveira[4], A Powell[8], C Ø Rasmussen[8], T Robertson-Brown[3], F Robicheaux[16], R L Sacramento[2], E Sarid[17,18], J Schoonwater[3], D M Silveira[2], J Singh[4], G Smith[1,7], C So[7], S Stracka[6], J Suh[10], A G Swadling[10], T D Tharp[19], K A Thompson[3], R I Thompson[10], E Thorpe-Woods[3], M Urioni[6], D P van der Werf[3], P Woosaree[10], J S Wurtele[9] and The ALPHA Collaboration[22]

1 Department of Physics and Astronomy, University of British Columbia, Vancouver, BC, Canada
2 Instituto de Fisica, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil
3 Department of Physics, Faculty of Science and Engineering, Swansea University, Swansea, United Kingdom
4 School of Physics and Astronomy, University of Manchester, Manchester, United Kingdom
5 Sci-Tech Daresbury, Cockcroft Institute, Warrington, United Kingdom
6 University of Brescia, Brescia and INFN Pavia, Pavia, Italy
7 TRIUMF, Vancouver, BC, Canada
8 Experimental Physics Department, CERN, CH-1211 Geneva 23, Switzerland
9 Department of Physics, University of California at Berkeley, Berkeley, CA, United States of America
10 Department of Physics and Astronomy, University of Calgary, Calgary, AB, Canada
11 Department of Physics and Astronomy, Aarhus University, Aarhus, Denmark
12 Department of Physics, Simon Fraser University, Burnaby, BC, Canada
13 Van Swinderen Institute for Particle Physics and Gravity, University of Groningen, Groningen, The Netherlands
14 Department of Physics, Stockholm University, Stockholm, Sweden
15 Department of Chemistry, University of British Columbia, Vancouver, BC, Canada
16 Department of Physics and Astronomy, Purdue University, West Lafayette, IN, United States of America
17 Ben Gurion University of the Negev Be'er Sheva, Israel
18 Soreq NRC, Yavne, Israel
19 Physics Department, Marquette University, Milwaukee, WI, United States of America
20 Department of Physics and Astronomy, University of Victoria, Victoria BC, V8P 5C2, Canada
21 Deceased.
22 The contact for the collaboration is alpha-contact@cern.ch.
* Author to whom any correspondence should be addressed.

E-mail: danielle.louise.hodgkinson@cern.ch

Keywords: antihydrogen, antimatter, neutral atom trapping, nonlinear dynamics, chaotic dynamics, octupolar ioffe-pritchard magnetic traps, energy diagnostic

Abstract

Precision measurements of the fundamental properties of trapped antihydrogen in the (ALPHA) experiment at CERN require an accurate determination of the antihdyrogen energy distribution. We introduce a new diagnostic technique, termed Spatio-Temporal Annihilation Mapping (STAM), which measures this distribution by exploiting the dynamics of antihydrogen in adiabatically time-varying magnetic fields. By analyzing the spatial and temporal patterns of annihilations recorded during adiabatic release from magnetic confinement and comparing them with detailed simulations, we infer the underlying energy distributions. STAM enables the simultaneous measurement of axial and transverse energy components, achieving a relative precision of approximately 10% over a mean energy range of ∼10–100 mK in the ALPHA trap. The technique is applied to both uncooled and laser-cooled antihydrogen populations, with results benchmarked against established laser-based diagnostics. While laser-based methods provide high-resolution measurements of the transverse energy and limited sensitivity to the axial component, an accurate determination of the axial energy remains challenging. In addition to complementing existing techniques, STAM is expected to access energy regimes that are currently inaccessible to laser-based diagnostics due to intrinsic limitations of those methods. These capabilities make STAM broadly applicable to future high-precision spectroscopic and gravitational studies with antihydrogen. Furthermore, STAM provides a novel probe of off-axis magnetic field structures, addressing a critical gap in current magnetometry techniques within the ALPHA experiment.

1. Introduction

The antihydrogen laser physics apparatus (ALPHA) experiment at CERN measures the fundamental properties of magnetically-trapped antihydrogen ($\overline{\mathrm{H}}$), e.g. [1–4]. An insufficient understanding of the trapped $\overline{\mathrm{H}}$ energy distribution is often a significant source of measurement uncertainty in experiments conducted at ALPHA; for example, in the first direct observation that $\overline{\mathrm{H}}$ is gravitationally attracted to the Earth [5]. In spectroscopic measurements, the kinetic energy distribution strongly influences the observed lineshape: when Doppler broadening dominates [1, 2], the lineshape is sensitive to the component of the velocity parallel to the spectroscopy laser beam; when transit-time broadening dominates [3, 6], it is instead sensitive to the perpendicular component. This work was conducted in the ALPHA-2 apparatus, which is primarily used for spectroscopic measurements of $\overline{\mathrm{H}}$. Gravitational experiments are carried out in the ALPHA-g apparatus, in which the trap axis is rotated by 90° relative to ALPHA-2 such that Earth's gravitational field points along the trap axis. However, due to the similarities between the $\overline{\mathrm{H}}$ traps, all of the work discussed here is directly applicable to ALPHA-g, aside from the fact that the spatial resolution with which $\overline{\mathrm{H}}$ annihilations can be reconstructed is significantly lower in ALPHA-g. We do not expect this to prevent the technique from being useful in ALPHA-g, but assessing its performance will be a topic of future study.

In ALPHA, $\overline{\mathrm{H}}$ are confined in a magnetic minimum trap, with radial confinement provided by an octupole magnet and axial confinement by five axially-spaced short-solenoids, which we refer to as mirror coils, and two end solenoids, all shown in figure 1. The total energy of trapped $\overline{\mathrm{H}}$ motion is the sum of its magnetic potential and kinetic energy. Due to the cylindrical geometry of the trap and the partial separation of kinetic energy parallel and transverse to the trap axis (discussed in detail in [7]), it is useful to decompose the $\overline{\mathrm{H}}$ energy into components along these degrees of freedom. Throughout this article, we refer to these as energy components, noting that this does not constitute a vector decomposition. We define the axial component along the z direction, chosen to be coaxial with the trap axis, as $E_{\parallel}$ and the component in the transverse plane as $E_{\perp}$. The quantity $E_{\parallel}$ is nearly, though not exactly, aligned with the spectroscopy laser beam used for precision measurements, as the beam enters the trap at an angle of 2.3° relative to the axis (see figure 1).

Existing techniques currently employed in ALPHA to diagnose the energy of trapped $\overline{\mathrm{H}}$ fall broadly into two categories: laser-based methods and diagnostics based on controlled magnetic release. Diagnostics based on controlled magnetic release have previously been used in ALPHA primarily to extract information about the total energy of the trapped $\overline{\mathrm{H}}$ population [8–10], but thus far have not provided access to the individual axial and transverse energy components. Laser-based techniques use short-pulse laser light at 121 nm to excite $\overline{\mathrm{H}}$ into untrapped states, allowing the reconstruction of the energy component perpendicular to the laser beam by inferring the time-of-flight from on-axis excitation to detected annihilations on the trap wall [11]. In principle, the energy component parallel to the laser beam can be inferred by sweeping the laser frequency across resonance and fitting the resulting Doppler-broadened spectrum. In practice, however, the presence of multiple line-broadening mechanisms complicates this approach, preventing a straightforward determination of the parallel energy component [11]. Moreover, the method is intrinsically limited at low energies, since decay to the untrapped state imparts a recoil velocity of 3.3 $\mathrm{m\,s^{-1}}$, corresponding to 0.6 mK in energy units. Throughout this article, energy, E, is expressed in temperature-equivalent units, i.e. $T = E/k_{\mathrm{B}}$, where k_{B} is the Boltzmann constant.

In this work, we present a novel technique that extends previous applications of magnetic release in ALPHA to diagnose the energy of trapped antihydrogen by analyzing the spatial and temporal distributions of $\overline{\mathrm{H}}$ annihilations during a tailored form of release from magnetic confinement in which radial magnetic confinement is removed while axial confinement is maintained. We refer to this technique as *Spatio-Temporal Annihilation Mapping* (STAM). Although this release configuration has been previously performed [10], we introduce a key insight that substantially increases both the information extracted and the achievable precision of this established release protocol. Specifically, we recognize that the dynamics governing the radial adiabatic release give rise to strong, energy-dependent correlations

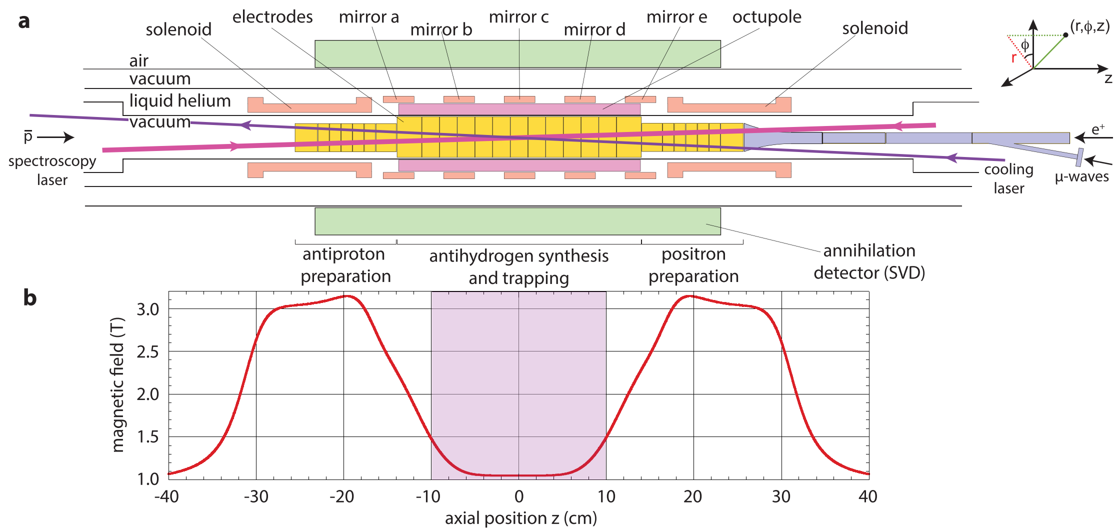

Figure 1. (a) Diagram of the ALPHA-2 Apparatus. Antiprotons ($\overline{\mathrm{p}}$) and positrons (e^+) enter the apparatus from the left and right respectively. (b) Magnetic field (in Tesla) along the trap axis in the flat configuration (see text). The pink shaded region marks the maximum axial region occupied by $\overline{\mathrm{H}}$ during the STAM release. The electrodes of the Penning–Malmberg trap used to confine charged particles are shown in yellow, and sit inside the superconducting magnets of the $\overline{\mathrm{H}}$ trap. The axial scale in (a) and (b) is the same. The radial extent of the SVD is larger than illustrated. An external solenoid (not shown), located radially outside the SVD, generates a uniform axial magnetic field of 1 T within the trap.

in the spatial and temporal annihilation distributions. By exploiting these correlations through detailed modeling of the release dynamics, STAM enables accurate reconstruction of both axial and transverse energy components from the observed annihilation data.

At its core, STAM exploits the fact that higher-energy $\overline{\mathrm{H}}$ access stronger magnetic fields, meaning they escape earlier during the release. Additionally, because the magnetic field strength increases with distance from the trap minimum, higher-energy $\overline{\mathrm{H}}$ atoms access more distant regions of the trap. Together, these effects ensure that the magnetic field strength at the time and location of annihilation encodes information about the $\overline{\mathrm{H}}$'s axial and transverse energy.

Because STAM is implemented during the controlled removal of the magnetic trapping potential, it constitutes a fully destructive measurement in which the trapped $\overline{\mathrm{H}}$ population is released and annihilates. Laser-based diagnostics likewise rely on inducing transitions to untrapped states, leading to annihilation of the affected atoms. In contrast to STAM, however, laser excitation can be applied in a partially destructive manner, removing only the subset of atoms that pass through the laser beam. The two approaches are therefore complementary, providing access to different aspects of the trapped antihydrogen ensemble. STAM exploits the unique diagnostic potential of antimatter annihilation; however, the insights into the escape dynamics may be transferable to similar magnetic minimum traps.

In this paper, we demonstrate the utility of STAM for diagnosing typical energy distributions both before (mean energies on the order of 100 mK) and after (mean energies on the order of 10 mK) Doppler laser cooling.

2. Experimental context

In the ALPHA experiment, $\overline{\mathrm{H}}$ is synthesized by merging cold plasmas of antiprotons and positrons in a Penning–Malmberg trap (figure 1). The resulting $\overline{\mathrm{H}}$ are confined in a superimposed magnetic minimum trap. Owing to the ultra-high vaacuum conditions ($<10^{-14}$ mbar), the lifetime of trapped $\overline{\mathrm{H}}$ due to collisions with background gas exceeds 66 hours [12]. To build up a population of trapped $\overline{\mathrm{H}}$, the magnetic trap remains energized while new antihydrogen is synthesized every four minutes over several hours. This accumulation process, which has produced up to 15 000 trapped $\overline{\mathrm{H}}$, is detailed in [13], where recent improvements in trapping efficiency were achieved through sympathetic cooling of the positron plasma with co-trapped, laser-cooled beryllium ions. We leverage these advances to obtain high-statistics datasets, enabling a detailed analysis using the energy diagnostic technique described below. Given the low density (<100 cm^{-3}) of trapped $\overline{\mathrm{H}}$, collisions between $\overline{\mathrm{H}}$ are negligible.

In the synthesis procedure, $\overline{\mathrm{H}}$ is formed via three-body combination, producing $\overline{\mathrm{H}}$ in Rydberg states which decay in a cascade to the ground state on $\sim$millisecond timescales [14]. The magnetic moment of the $\overline{\mathrm{H}}$ remains adiabatically locked, either parallel or antiparallel depending on the positron spin, to the trap magnetic field lines due to the small variation in magnetic field over a single Larmor precession cycle. The $\overline{\mathrm{H}}$ are produced in a mixture of hyperfine states, but only those with positron spins aligned antiparallel to the magnetic field, known as low-field-seeking states, remain confined and are governed by the Hamiltonian

$$H(\mathbf{x},\mathbf{v}) = \frac{m|\mathbf{v}|^2}{2} + U(\mathbf{x}), \tag{1}$$

where $\mathbf{x} \equiv \{r,\phi,z\}$ and $\mathbf{v} \equiv \{v_r, v_\phi, v_z\}$ denote the position and velocity in a cylindrical basis (see figure 1). The magnetic potential term is given by

$$U(\mathbf{x},t) = |\boldsymbol{\mu}|\left\{|\mathbf{B}(\mathbf{x},t)| - |\mathbf{B}(\mathbf{x}_{\min}(t),t)|\right\}, \tag{2}$$

where $\mathbf{x}_{\min}(t)$ is the spatial location of the global magnetic minimum, t is time, $\mathbf{B}(\mathbf{x},t)$ is the magnetic field vector, and $\boldsymbol{\mu}$ is the $\overline{\mathrm{H}}$ magnetic moment vector. Throughout, we assume $\boldsymbol{\mu}$ is equal to the hydrogen magnetic moment vector which is one bohr magneton (μ_{B}) to a very good approximation for the low-field-seeking ground states. The potential is expressed as a function of time because dynamic magnetic field ramps are used in this work to diagnose the $\overline{\mathrm{H}}$ energy.

For the initial $\overline{\mathrm{H}}$ confinement, the depth of the magnetic trapping potential is typically around 0.5 K, which is bounded by the magnetic potential at the nearest material surfaces to the confinement region: the inner surface of the Penning–Malmberg electrodes. $\overline{\mathrm{H}}$ will annihilate if they meet the material walls of the apparatus. These annihilation events can be resolved in space and time through analysis of events from ALPHA-2's Silicon-Vertex-Detector (SVD) [15, 16].

Prior to $\overline{\mathrm{H}}$ formation, the antiprotons are assumed to equilibrate with the temperature of the positron plasma [17, 18]. Hence, the $\overline{\mathrm{H}}$ is assumed to be formed in a thermal distribution at this temperature [7]. Currently employed measurement techniques [19] limit our ability to measure the positron plasma temperature to $\lesssim 7$ K [13, 20]. Most $\overline{\mathrm{H}}$ with energy exceeding the trap depth annihilate shortly after formation, resulting in a truncated energy distribution. A small fraction of $\overline{\mathrm{H}}$ with energies above the trap depth, referred to as quasi-trapped $\overline{\mathrm{H}}$ [8], can remain confined if they occupy trajectories that do not reach the trap walls.

3. $\overline{\mathrm{H}}$ escape model

In this section, we describe our model underlying the energy-dependent escape of $\overline{\mathrm{H}}$ during magnetic release, and how the $\overline{\mathrm{H}}$ energy distribution can be reconstructed from these escape events. The discussion largely follows the experimental procedure described in section 6. We first consider the simple case in which $\overline{\mathrm{H}}$ is confined in a magnetic field configuration known as the *harmonic configuration*. In this setup, axial confinement of $\overline{\mathrm{H}}$ is provided by the outer mirror coils and solenoids. The release protocol involves ramping down only the octupole from its full field at time $t = 0$ to zero field at $t = T_{\mathrm{ramp}}$, while the magnetic fields of the mirror coils and solenoids remain fixed. After discussing this simple case, we generalize the framework to diagnose the $\overline{\mathrm{H}}$ energy distribution for an arbitrary initial magnetic configuration. The escape mechanism is initially described in simplified terms, followed by a discussion of effects that complicate the underlying dynamics.

The release protocol is carried out adiabatically: T_{ramp} is sufficiently long such that the change in magnetic potential, U, during a single $\overline{\mathrm{H}}$ oscillation is small compared to the magnetic potential itself i.e.

$$\tau \frac{\mathrm{d}U(\mathbf{x},t)}{\mathrm{d}t} \ll U(\mathbf{x},t) \tag{3}$$

where τ is the characteristic ($\sim$ millisecond) bounce time of an $\overline{\mathrm{H}}$ oscillation. Motivated by this slow release and the observation that the $\overline{\mathrm{H}}$ follow chaotic trajectories [7], we make an *ergodic assumption* that $\overline{\mathrm{H}}$ have sufficient time to move throughout all regions of the trap accessible to them based on their axial and transverse energies.

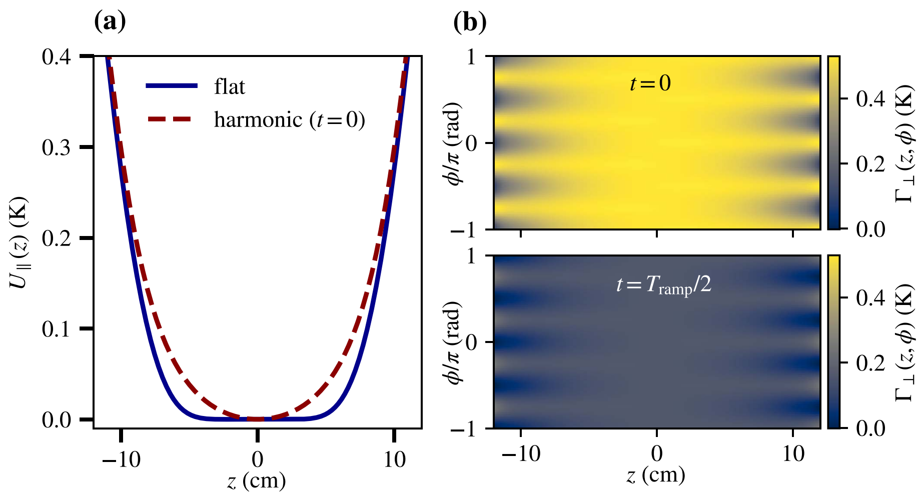

Figure 2. (a) Modeled axial potential, $U_{\parallel}(z)$, plotted for the flat configuration (blue solid curve) and in the harmonic configuration (evaluated at $t=0$, red dashed curve). See text for configuration definitions. (b) Modeled transverse potential on the inner electrode surface (transverse trap depth, $\Gamma_{\perp}(z,\phi,t)$) at two different times, t, during an octupole rampdown. During the octupole rampdown in (b), the axial potential remains approximately static and corresponds to the harmonic configuration plotted in (a). See appendix B for details of the magnetic field calculation.

Due to the geometry of the magnetic trapping potential and the separation of the energy definition into $E_{\parallel}$ and $E_{\perp}$, we decompose the magnetic potential into axial and transverse components. The reference axis is defined by locating the minimum $|\mathbf{B}(\mathbf{x})|$ in the transverse plane at each value of z and connecting these minima. We then define the *axial potential* as

$$U_{\parallel}(z,t) \equiv U(\mathbf{x} = \{r_{\min}(z,t), \phi_{\min}(z,t), z\}, t), \tag{4}$$

and the *transverse potential* as

$$U_{\perp}(r,\phi,z,t) \equiv U(\mathbf{x} = \{r,\phi,z\}, t) - U(\mathbf{x} = \{r_{\min}(z,t), \phi_{\min}(z,t), z\}, t), \tag{5}$$

where $r_{\min}(z,t)$ and $\phi_{\min}(z,t)$ denote the coordinates of the minimum $|\mathbf{B}(\mathbf{x})|$ in the transverse plane. At $t=0$, this minimum lies a few millimetres off axis due to the competing radial dependences of the mirror-coil magnetic potential, which decreases with radius, and the octupole magnetic potential, which increases with radius. As the octupole magnetic field is reduced, the transverse minimum moves radially outward toward the trap wall.

The axial potential is shown in figure 2(a) and represents the axial energy required for an $\overline{\mathrm{H}}$ orbit to reach a given axial position z while remaining on the minimum magnetic potential in the transverse plane. The axial potential $U_{\parallel}(z,t)$ remains largely unchanged throughout the octupole rampdown; however, it varies by up to approximately 10 % within $|z| < 10\,cm$, which corresponds to the approximate maximum axial extent occupied by trapped $\overline{\mathrm{H}}$. Critically, the harmonic configuration provides a monotonic mapping between axial potential and z.

The transverse potential represents the transverse energy required for an $\overline{\mathrm{H}}$ orbit to reach a given radial displacement from the transverse potential minimum at a fixed axial position z. At time t, a trapped $\overline{\mathrm{H}}$ is confined to a volume defined by the conditions $U_{\parallel}(z,t) \leqslant E_{\parallel}(t)$ and $U_{\perp}(r,\phi,z,t) \leqslant E_{\perp}(t)$. The energy components are expressed as explicit functions of time, as they may evolve during the release process due to adiabatic expansion, which is discussed in section 3.1.

Evaluating the transverse potential at the radius of the inner electrode surface ($r_{\max} = 2.2175$ cm in ALPHA-2) defines a quantity we refer to as the *transverse trap depth*,

$$\Gamma_{\perp}(z,\phi,t) \equiv U_{\perp}(r = r_{\max}, \phi, z, t), \tag{6}$$

which represents the minimum transverse energy required for an $\overline{\mathrm{H}}$ to reach the trap wall and annihilate at a given axial position z, and is shown in figure 2(b). Under the ergodic assumption, annihilation occurs when the transverse energy equals the local transverse trap depth,

$$E_{\perp}(t_{\mathrm{a}}) = \min_{\phi} \Gamma_{\perp}(z_{\mathrm{a}}, \phi, t_{\mathrm{a}}), \tag{7}$$

where the subscript a denotes the annihilation coordinates. As the octupole rampdown progresses, the transverse trap depth decreases at all z and ϕ locations and exhibits eight local minima of nearly equal magnitude in the z–ϕ plane, corresponding to regions where the radial components of the mirror-coil and octupole magnetic fields oppose one another. Consequently, this model predicts eight preferred annihilation locations in ϕ, each occurring with approximately equal probability.

Since the minimal transverse trap depth decreases with increasing $|z|$ (see figure 2(b)), the ergodic assumption further implies that annihilation occurs at the maximum z accessible to the atom. Thus,

$$E_{||}(t_{\mathrm{a}}) = U_{||}(z_{\mathrm{a}}, t_{\mathrm{a}}). \tag{8}$$

The deterministic nature of this relationship–rather than a probabilistic one, with a finite chance of annihilation over an extended region of the trap–is a key reason the precision of this technique surpasses previous implementations. We note that this model is an effective description and does not capture all aspects of the coupled axial–transverse dynamics. Its validity is examined in section 4, where we find it to be a good approximation for the system considered here.

Thus far, we have considered octupole release from the harmonic configuration. For ALPHA, STAM must also diagnose $\overline{\mathrm{H}}$ energies across other magnetic-field setups. A common example is the *flat configuration*, which increases trap uniformity by energizing mirror coils B, C, and D with small opposite currents relative to the outer coils, thereby maximizing the resonant volume accessible to the spectroscopy laser (figures 1(b) and 2(a)).

Directly ramping down the octupole from the flat configuration causes $\overline{\mathrm{H}}$ spanning a wide range of axial energies to annihilate over a narrow z range, due to the sharp axial potential gradient near the trap ends. This limits the achievable axial-energy resolution, given the resolution with which z_{a} can be detected by the SVD. To mitigate this effect, $\overline{\mathrm{H}}$ is first confined in the flat configuration, after which the field is adiabatically morphed to the harmonic configuration by ramping down coils B, C, and D. The theoretical model remains applicable under this transformation. This method can be adapted to diagnose energies in other configurations by tailoring the initial potential prior to ramping down the mirrors.

3.1. Effects beyond the simple escape picture

Several important complications have been omitted from the simplified picture described above, all of which are incorporated into the simulation-based energy reconstruction method detailed in section 5. One complication is that the dynamic magnetic fields used in the release technique can themselves induce changes in $\overline{\mathrm{H}}$ energy via adiabatic expansion or compression [10, 21]. In this process, total energy is not conserved, in contrast to mechanisms that merely redistribute energy between dimensions. This effect occurs during both the octupole rampdown and the transition from the flat configuration to the harmonic configuration. The former leads primarily to radial adiabatic cooling, while the latter results in axial adiabatic heating.

Energy exchange between the axial and transverse degrees of freedom during magnetic manipulations presents an additional complication. As studied in [7], this arises from gradients in the magnetic trapping potential and depends on the initial axial-to-transverse energy ratio of each $\overline{\mathrm{H}}$ trajectory. Owing to the chaotic nature of $\overline{\mathrm{H}}$ motion in ALPHA traps [7], $\overline{\mathrm{H}}$ that undergo energy exchange no longer retain a unique, invertible correspondence between their axial and transverse energies and their annihilation coordinates. Simulations predict that most $\overline{\mathrm{H}}$ exchange energy within a few seconds in the flat configuration [7], eliminating a one-to-one mapping between individual $\overline{\mathrm{H}}$ energies in the flat configuration and those at annihilation. While this precludes accurate determination of axial and transverse energy components for individual $\overline{\mathrm{H}}$, the ensemble energy distributions remain statistically recoverable. Because energy exchange influences the degree of adiabatic cooling or heating during magnetic ramps [10], it must be accounted for when interpreting the results.

Finally, some $\overline{\mathrm{H}}$ may not satisfy the ergodic assumption during the octupole rampdown. If $\overline{\mathrm{H}}$ fail to fully explore the trap volume accessible to them, based on their axial and transverse energies, this can introduce errors in the reconstructed energy distributions.

4. Testing the model via simulation

In this section, we test the model presented in section 3 for mapping energy-dependent escape from the flat configuration, as an example of an arbitrary initial magnetic-field configuration used for energy diagnosis. We perform simulations to investigate energy-dependent $\overline{\mathrm{H}}$ escape during the STAM technique, using techniques based on those described in [7]. The simulation is performed as follows:

(i) Initial position, $\mathbf{x}$, and velocity, $\mathbf{v}$, are Monte-Carlo sampled as described below.

(ii) At each simulation time t, the magnetic fields generated by the trap magnets are computed as described in appendix B. The currents in the individual magnets are varied as a function of time to simulate the STAM release described in section 3. The specific timings of the magnetic ramps are configured to match those of the experimental procedure which will be discussed in section 6.

(iii) At each timestep, the equations of motion, $m\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = -\mu_{\mathrm{B}}\nabla|\mathbf{B}(\mathbf{x},t)|$ and $\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \mathbf{v}$, are advanced using a symplectic Leapfrog integrator with a fixed timestep of $10\,\mu$s. The gradient $\nabla|\mathbf{B}(\mathbf{x},t)|$ is evaluated using a central finite-difference approximation.

(iv) Steps (ii)and(iii) are repeated until the $\overline{\mathrm{H}}$ atom reaches the material walls of the electrode stack, at which point annihilation is assumed to occur.

Initial conditions for the $\overline{\mathrm{H}}$ simulations fall into two categories, which we refer to throughout this work as the *uncooled* and *thermal 10 mK* cases. For the uncooled population, $\overline{\mathrm{H}}$ kinetic energies are sampled from a Maxwell–Boltzmann distribution at the positron plasma temperature, which we take to be 3 K as an illustrative example. In our analysis of the experimental data, we use simulations spanning a range of positron plasma temperatures bounded by the limits of currently employed plasma diagnostics, as discussed in section 7.1. Initial positions are sampled uniformly within a coaxial, ellipsoidal volume of length 10 mm and radius 0.8 mm, corresponding to the spatial extent of the positron plasma where $\overline{\mathrm{H}}$ is formed. $\overline{\mathrm{H}}$ are assumed to be initially in high principal quantum number Rydberg states and undergo a simulated cascade to the ground state via circular transitions, following the model described in [14].

In the thermal 10 mK case, used to approximate a laser-cooled population, $\overline{\mathrm{H}}$ are initialized in thermal equilibrium at 10 mK. There is no physical justification for assuming thermal equilibrium for this laser-cooled population; this assumption is made for convenience, in order to populate an approximately correct energy range. The cascade dynamics are not applied, since the laser cooling occurs after the $\overline{\mathrm{H}}$ reach the ground state. Initial positions are sampled from the full trapping volume using a rejection method: positions are drawn uniformly over the trap volume and accepted with a probability proportional to $\exp[-(U-U_{\mathrm{min}})/kT]$, where U is the magnetic potential energy and $T = 10$ mK is the temperature. Velocities are drawn independently from a Maxwell–Boltzmann distribution at the same temperature.

The energy of simulated $\overline{\mathrm{H}}$ is evaluated in the flat configuration in order to establish a mapping between an $\overline{\mathrm{H}}$'s axial and transverse energies and its annihilation coordinates. Following [7], the axial and transverse energy components are determined by evaluating the atom's motion at the location of the axial magnetic potential minimum, $z_{\mathrm{min}}(\approx 0$ mm). At this location, the axial energy is purely kinetic,

$$E_{\parallel} = \frac{1}{2}mv_z^2, \tag{9}$$

and the transverse energy includes both kinetic and magnetic potential contributions,

$$E_{\perp} = \frac{1}{2}m\left(v_r^2 + v_\phi^2\right) + U(r,\phi,z = z_{\mathrm{min}}). \tag{10}$$

This decomposition is only valid at z_{min}; at other points along the $\overline{\mathrm{H}}$'s trajectory, both axial and transverse directions include magnetic potential contributions, making it ambiguous how the total energy is partitioned between kinetic and potential terms.

The spatial and temporal distributions of simulated $\overline{\mathrm{H}}$ annihilations for the uncooled and thermal 10 mK populations are shown in figures 3(a) and (b), respectively, for a 180 s octupole rampdown. The ramp-down duration is chosen to prevent saturation of the SVD in the presence of the large antihydrogen populations, of order 10^4, that are now routinely accumulated by ALPHA [13]. The detector spatial resolution in z is well described by a mixture of two Gaussian distributions [15]. By fitting this double-Gaussian model to reconstructed antiproton annihilation vertices obtained from point-like, held sources of antiprotons, we find that one Gaussian has a standard deviation of 4 mm and the other 14 mm, with each component applied with equal probability. These antiproton measurements do not provide information about the detector angular resolution. We therefore determine this independently using simulations [15] of the scattering of charged annihilation products through the apparatus material and magnetic field. These simulations predict that the angular resolution follows a Gaussian distribution with

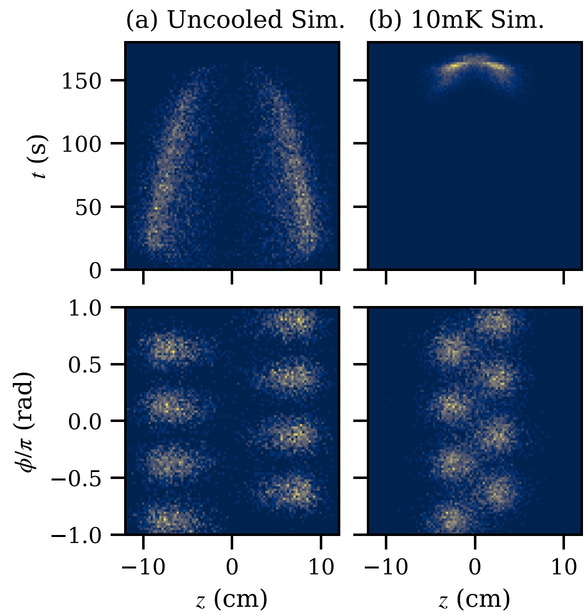

Figure 3. Simulated distributions of $\overline{\mathrm{H}}$ annihilation locations on the inner electrode surface, where $t = 0$ s marks the start of the 180 s octupole rampdown. Each distribution contains 10 000 simulated events for the (a) uncooled and (b) thermal 10 mK populations, initialized as described in the text. The distributions are displayed as two-dimensional histograms. The color scale represents the fraction of events in each bin, with yellow indicating the highest occupancy and dark blue the lowest.

a standard deviation of 0.26 radians. To model the SVD spatial resolution, we randomly smear the simulated annihilation positions according to these resolution models. The SVD reconstructs the annihilation time with a resolution of approximately 2 μs, which is negligible on the timescale of the rampdown.

As expected from the presence of eight nearly-equal minima in the transverse magnetic potential (see figure 2(b)), the simulated annihilation distribution exhibits eight distinct peaks in the azimuthal coordinate. In the axial coordinate, two peaks are observed, consistent with annihilations occurring preferentially near the turning points of the axial motion. The thermal 10 mK population shows a pronounced shift toward later annihilation times and lower axial annihilation positions compared to the uncooled case. This behavior reflects the lower average energy of the cooled population and highlights the STAM technique's sensitivity to variations in the underlying $\overline{\mathrm{H}}$ energy distribution.

Figures 4(a)–(c) compare the simulated axial, transverse, and total energies in the flat configuration (the well in which we diagnose the energy) against the model predictions given by equations (7) and (8) for the uncooled simulation. The equivalent plot for the thermal 10 mK population is provided in appendix A. Significant scatter is observed in the axial and transverse benchmarks, which is notably reduced in the total energy comparison. This behavior reflects the inability of the model to accurately predict an individual $\overline{\mathrm{H}}$'s axial and transverse energy components due to energy exchange between these degrees of freedom during the release sequence (see section 3.1). Additional deviations from the model arise from unaccounted-for adiabatic changes in energy during the magnetic field evolution.

To test the model without the confounding effects of energy exchange and adiabatic changes, figures 4(d)–(f) also shows the equivalent comparison using the axial and transverse energy components evaluated at each $\overline{\mathrm{H}}$'s last crossing of the axial magnetic minimum prior to annihilation. In this case, the deviation between the simulation results and the model is significantly reduced. This can be understood from the fact that energy exchange and adiabatic evolution tend to be much smaller over the course of a single axial oscillation than over the full magnetic manipulation sequence. The population of $\overline{\mathrm{H}}$ with total energies significantly exceeding the predicted values likely corresponds to $\overline{\mathrm{H}}$ that did not sufficiently sample the accessible phase space, i.e. those that do not satisfy the ergodic assumption (see section 3), and are therefore destined to annihilate at a later time once they find their way out.

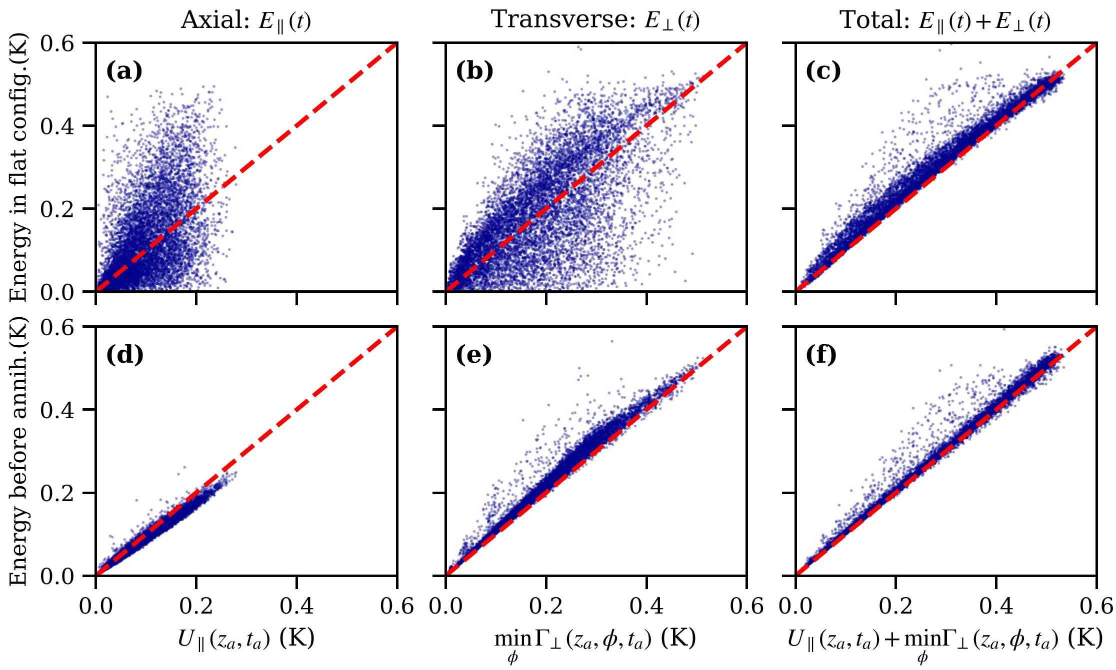

Figure 4. Demonstration of the energy content encoded in the annihilation data. Each blue point represents one of 10 000 simulated $\overline{\mathrm{H}}$ atoms, initialized according to the uncooled Monte Carlo distribution described in the text and subsequently subjected to the magnetic manipulations of the STAM diagnostic (shown in figure 2) prior to annihilation. The spread of points reflects both the distribution of initial conditions sampled by the Monte Carlo initialization and the subsequent dynamical evolution of the $\overline{\mathrm{H}}$ in the trapping potential. Panels a–c compare the energy components in the flat trap configuration, extracted directly from the simulation, with the corresponding model predictions calculated from the simulated annihilation coordinates using equations (7) and (8). Perfect agreement between model and simulation is indicated by the red dashed line. Panels d–f show the same comparison, but using the energy components extracted from the simulation at each simulated $\overline{\mathrm{H}}$ atom's final crossing of the magnetic minimum prior to annihilation ($t \approx t_\mathrm{a}$).

5. Energy reconstruction method

Because the model does not account for energy exchange or adiabatic energy changes, we instead estimate the distribution of axial and transverse energies in the flat configuration, $f_E(E_\parallel, E_\perp)$, from annihilation measurements using a variation of the technique in [9]. We define the integral

$$f_E\left(E_\parallel, E_\perp\right) = \int_0^\infty \mathrm{d}t_\mathrm{a} \int_{-\infty}^\infty \mathrm{d}z_\mathrm{a}\, P\left(E_\parallel, E_\perp \mid z_\mathrm{a}, t_\mathrm{a}\right) f(z_\mathrm{a}, t_\mathrm{a}), \tag{11}$$

where $f(z_\mathrm{a}, t_\mathrm{a})$ is the observed distribution of annihilations during the release, and $P(E_\parallel, E_\perp \mid z_\mathrm{a}, t_\mathrm{a})$ is the simulated conditional probability map that connects $(E_\parallel, E_\perp)$ in the flat configuration to the annihilation coordinates $(z_\mathrm{a}, t_\mathrm{a})$. The simulations are used to construct $P(E_\parallel, E_\perp \mid z_\mathrm{a}, t_\mathrm{a})$, allowing the energy reconstruction to incorporate the effects discussed in section 3.1 that are not captured by the simple analytic model. The dependence on ϕ_a has been omitted. While the annihilation distribution exhibits azimuthal structure, the transverse trapping potential contains azimuthal minima (figure 2(b)) toward which escaping $\overline{\mathrm{H}}$ atoms are preferentially directed prior to annihilation (see equation (7)). Therefore, the azimuthal annihilation location provides information about the relative strength of the magnetic field at these minima, but not about an $\overline{\mathrm{H}}$ atom's energy.

The key difference from the technique described in [9] is that the integral now includes both axial position and time, z_a and t_a, whereas the original method integrated over time only. This extension is motivated by the insights into the $\overline{\mathrm{H}}$ dynamics discussed in section 4, which show that both annihilation position and time exhibit strong, energy-dependent correlations with the initial axial and transverse energies. Crucially, this enables the axial and transverse energy components to be reconstructed separately, rather than only their sum.

We approximate the integral in equation (11) by adapting the method from [9] as follows: (a) For each observed annihilation event, we construct a narrow two-dimensional selection band of spatial width, dz, and temporal width, dt; centered on the event's axial position and annihilation time,

$(z_{\mathrm{a}}, t_{\mathrm{a}})$. (b) From this band, we randomly sample N_{samples} simulated annihilation events with replacement. (c) The corresponding $(E_{\|}, E_{\perp})$ values associated with these sampled events–evaluated in the flat configuration–are then aggregated. This sampling effectively integrates over t_{a} and z_{a}, weighted by $f(z_{\mathrm{a}}, t_{\mathrm{a}})$. (d) Finally, a histogram of the aggregated $(E_{\|}, E_{\perp})$ values is constructed, yielding the reconstructed energy distribution in the flat configuration. This method implicitly incorporates the effects of energy exchange and adiabatic changes to the energies, which are both challenging to model analytically, by relying on their natural inclusion in the simulation. As a result, it enables an accurate reconstruction of the underlying energy distribution.

Due to the binning-based method described above, the inferred energy distribution is independent of the specific simulated initial energy distribution used to sample $P(E_{\|}, E_{\perp} \mid z_{\mathrm{a}}, t_{\mathrm{a}})$, provided that the full $(E_{\|}, E_{\perp})$ space is sufficiently populated. This independence arises because a fixed number of samples are randomly drawn from each two-dimensional band, rendering the result insensitive to the absolute number of events within the band and therefore to the particular initial distribution. To ensure adequate population of the space, the probability map is constructed such that, for each annihilation event, a sufficient number of nearby simulated events are included within the corresponding band. In practice, this requires the average band population, N_{band}, to be large enough that N_{samples} events can be drawn without significant repetition; we therefore impose the condition $N_{\mathrm{band}} > 4N_{\mathrm{samples}}$.

The bandwidths, dz and dt, must also be chosen small enough that they do not smear out relevant features in the experimental data. As will be discussed in section 7.1, energy reconstruction using STAM in ALPHA is limited by systematic contributions. Consequently, the bandwidths are required to be small compared to systematic effects relevant to the experimental and simulated $(z_{\mathrm{a}}, t_{\mathrm{a}})$. The sampling uncertainty associated with the finite size of N_{samples} is quantified by repeating the binning-based reconstruction multiple times and evaluating the resulting spread in reconstructed energy distributions. Our requirement on N_{samples} is that this contribution be negligible compared to statistical uncertainties in the data.

To benchmark the binning-based technique, we reconstruct the energy distributions of separately simulated uncooled and thermal 10 mK $\overline{\mathrm{H}}$ populations. To ensure adequate coverage across the relevant energy range for both populations, we construct a combined probability map by aggregating data from the uncooled and thermal 10 mK simulations, each with 10 000 statistics. The benchmark samples are subsets of the full simulation data used to construct the probability map. These benchmark samples are independently smeared using the detector resolution model (see section 4), ensuring a fair comparison. We use a temporal bandwidth of $dt = 1.5\,\mathrm{s}$, a spatial bandwidth of $dz = 5\,\mathrm{mm}$, and $N_{\mathrm{samples}} = 5$ for the uncooled population and $N_{\mathrm{samples}} = 20$ for the laser cooled population. These values are identical to those used to reconstruct the energy distributions corresponding to the experimental samples that will be presented in section 7.

In figure 5, we compare the reconstructed energy distributions to the true distributions determined directly from the simulations. The same figure also includes reconstructed distributions obtained using the model (equations (7) and (8)). Because the model neglects the effects of energy exchange and adiabatic changes during the magnetic manipulations, discrepancies arise between the reconstructed and true energy distributions. In contrast, these effects are fully accounted for in the binning-based technique, which shows no statistically significant difference from the simulated distributions. This demonstrates that the STAM technique can accurately reconstruct the energy components of $\overline{\mathrm{H}}$, given the available resolution.

6. Experimental method

To investigate STAM experimentally, we performed dedicated experimental runs in which populations of several thousand $\overline{\mathrm{H}}$ were deliberately released from magnetic confinement in order to diagnose their energy. These runs were conducted both with and without laser cooling.

$\overline{\mathrm{H}}$ accumulation was performed over a period of several hours in the flat configuration (shown in figure 2(a)). In one set of runs, the trapped $\overline{\mathrm{H}}$ was laser cooled following the methods described in [11], while in the other set the population was left uncooled. We refer to these as the *laser-cooled* and *uncooled* populations, respectively.

Following laser cooling, or after a short hold of several minutes for the uncooled population, the small (few-ampere) currents in mirror coils B, C, and D were ramped down over 18 s to adiabatically morph the magnetic field into the harmonic configuration (also shown in figure 2(a). As described

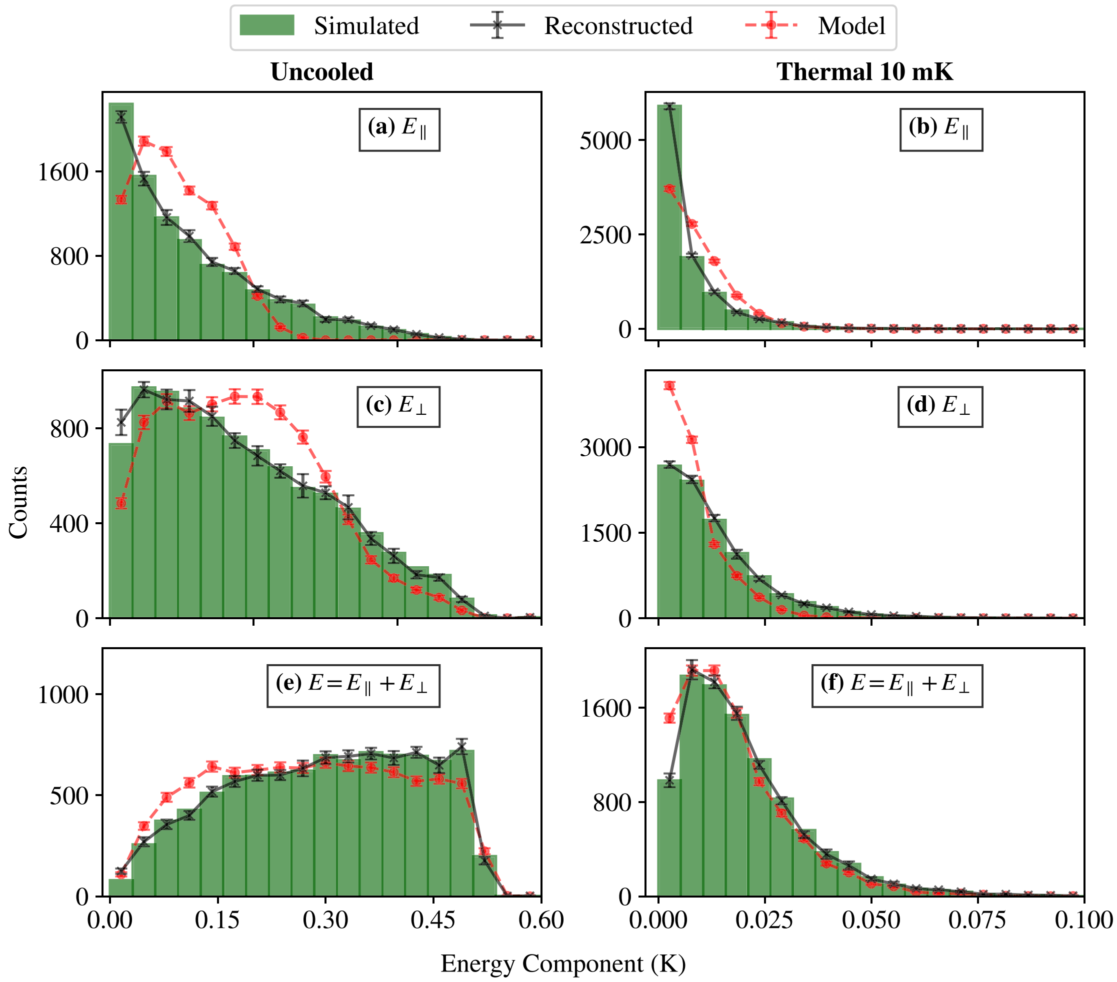

Figure 5. Simulation benchmark of the STAM technique. The green histogram represents the true energy distributions extracted from the simulation in the flat configuration. The red curve shows the reconstructed energy distributions using the model (equations (7) and (8)) applied to the (z_a, t_a) annihilation distributions from each simulated population. The black curve corresponds to the reconstructed distributions obtained by numerically approximating the integral in equation (11), as described in the text. Black error bars include statistical uncertainties and the contribution from finite sampling in the reconstruction. Red and black points are connected with a straight line to guide the eye. Red error bars represent statistical uncertainty. The error bars on the green histogram are omitted for clarity but are of the same order as those on the reconstructed distributions.

in section 3, this step increases the precision with which the axial energy can be determined given the available spatial resolution of the SVD.

Subsequently, the octupole current was ramped down quasi-linearly over 180 s, during which annihilation events were recorded by the SVD. Although the ramp is programmed to be linear, non-linearities can arise from details of the octupole current control circuitry. These non-linearities are incorporated into the simulations using measurements of the octupole current during the rampdown obtained with a direct-current current transducer (DCCT).

As a test of the robustness of the technique, additional experimental runs were performed, both with and without laser cooling, in which the duration of the octupole ramp was reduced to 18 s.

7. Experimental results and discussion

In figure 6, we present the observed annihilation distributions from the experiment. In general, the experimental data qualitatively validate the expected energy-dependent loss of $\overline{\mathrm{H}}$. However, it is immediately apparent that two of the eight annihilation hotspots predicted by the model (section 3) and nominal simulations (figure 3) are missing from both the uncooled and laser-cooled datasets. In the regions of discrepancy, a few percent of annihilation counts remain, but the signal is clearly suppressed relative to expectations.

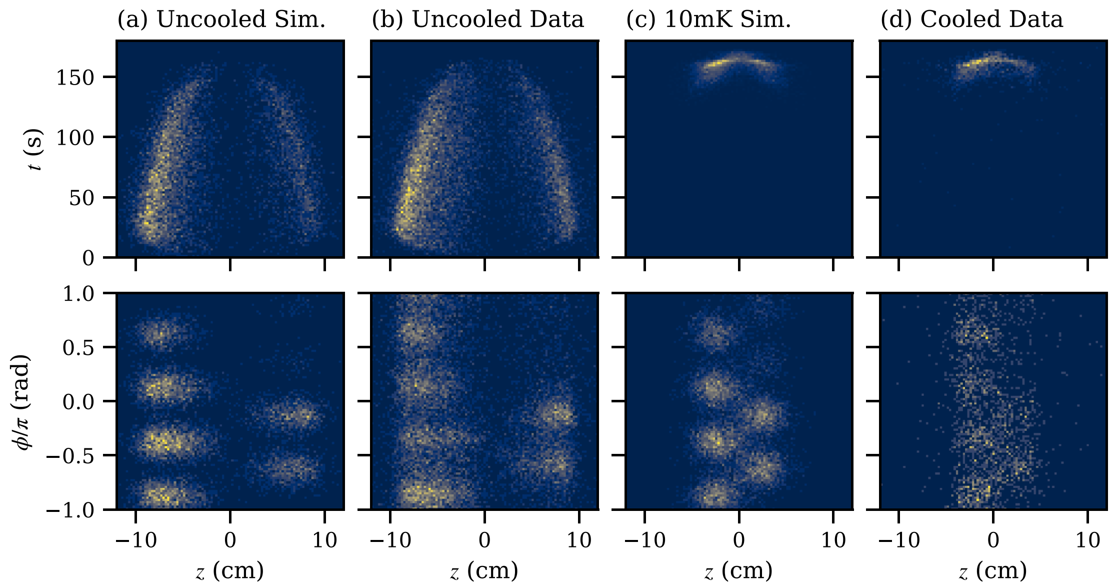

Figure 6. Distribution of $\overline{\mathrm{H}}$ annihilation locations on the inner electrode surface, where $t = 0$ s marks the start of the 180 s octupole rampdown. (a) Simulation of 10 000 $\overline{\mathrm{H}}$ atoms initialized under the uncooled conditions (see section 4). (b) Experimental data from 14 297 detected $\overline{\mathrm{H}}$ atoms without laser cooling. (c) Simulation of 10 000 $\overline{\mathrm{H}}$ atoms initialized under the thermal 10 mK conditions (see section 4). (d) Experimental data from 2070 detected $\overline{\mathrm{H}}$ atoms with laser cooling. Both simulations shown model possible winding errors in the octupole as described in the text. The distributions are displayed as two-dimensional histograms. The color scale represents the fraction of events in each bin, with yellow indicating the highest occupancy and dark blue the lowest.

Other experiments conducted in this apparatus release $\overline{\mathrm{H}}$ via different mechanisms, such as positron spin-flip transitions [22], in which the $\overline{\mathrm{H}}$ are actively accelerated toward the wall. This contrasts with the method discussed here, where the $\overline{\mathrm{H}}$ adiabatically sample the available magnetic potential before escaping. Notably, those other experiments have observed annihilations at the same azimuthal locations as the missing hotspots. This shows that the observed suppression does not result from non-uniform angular detection efficiency.

A plausible explanation for the missing hotspots is a discrepancy between the modeled and actual magnetic trapping potentials. Since the eight annihilation hotspots predicted by the model arise due to eight minima of nearly equal magnitude in the transverse potential (see figure 2), removing two of these minima requires raising the transverse potential at the corresponding locations. Through simulation, we investigated possible manufacturing deviations in the mirror coils, including shifts of their axis relative to the rest of the apparatus and elliptical perturbations to the coil geometry. However, within manufacturing tolerances, these effects were insufficient to reproduce the magnitude of the observed asymmetry. In contrast, the annihilation distributions are considerably more sensitive to geometric distortions of the octupole. This enhanced sensitivity arises from the steep radial dependence of the octupole magnetic field, which scales as r^3.

Within the allowed manufacturing tolerances, the observed asymmetry in the annihilation distribution can be reproduced by a combination of a conical squeeze of the octupole and an azimuthally asymmetric perturbation to the transverse trapping potential. Such an azimuthal asymmetry could arise either from an azimuthally asymmetric deformation of the octupole windings or from a small tilt between the axis of the Penning–Malmberg trap and that of the octupole. In the latter case, the tilt produces an equivalent modification of the magnetic field at the inner surface of the Penning–Malmberg electrodes, where the transverse trap depth is evaluated.

In figure 6, we present simulation results incorporating such a perturbation. For simplicity, the effect is modeled by modifying the Biot–Savart representation of the octupole, adjusting the radial positions of the winding segments as a function of axial position z. This introduces a conical squeeze of the octupole along the trap axis, while the perturbation is applied more strongly over one azimuthal sector to reproduce the observed azimuthal asymmetry. The maximum radial displacement of the windings from the nominal geometry is 240 μm, resulting in a local magnetic field deviation of up to 0.07 T at the inner electrode surface for the full octupole field strength (approximately 7 % of the full octupole field). The resulting change in the on-axis magnetic field is less than 0.001 T at all values of z with the maximum octupole current, ensuring consistency with ECR measurements.

Table 1. Summary of priors applied to experimental and simulated annihilation coordinates (z_a, t_a) during Monte Carlo error propagation.

Source of uncertainty	Prior applied to	Prior description
Axial detector resolution	Simulated z_a	Mixture of Gaussians: $\sigma = 14$ mm and 4 mm, equal weighting
Global detector alignment offset	Experimental z_a	Gaussian prior: mean 3 mm, $\sigma = 1$ mm; same shift applied to all z_a
Octupole magnetic field strength	Simulated t_a	Scale t_a according to 5% uncertainty in octupole field strength
Electrode alignment	Simulated z_a	Global axial scaling: $z_a \to \alpha z_a$, $\alpha \in [0.9, 1.1]$ drawn from uniform prior; applied to all z_a.

The distorted field qualitatively reproduces the main features of the experimental annihilation pattern while remaining within the manufacturing tolerances of the octupole. By applying the binning-based reconstruction (see section 5) using a full simulation map both with and without the azimuthal perturbation to the octupole, we find no statistically significant difference in the reconstructed energy. We therefore conclude that this effect does not contribute a resolvable systematic uncertainty.

Another discrepancy visible in figure 6 is that the angular annihilation pattern is more diffuse in the data than in the simulations. Since the detector angular resolution is estimated from simulations of charged-particle scattering (section 4), this likely reflects imperfect modeling of the scattering material. Since the STAM reconstruction does not use the annihilation azimuthal coordinate, this discrepancy has negligible impact.

7.1. Systematic effects

Our goal is to use the binning-based method described in section 5 to reconstruct the energy distribution of the experimental data from the measured annihilation coordinates (z_a, t_a). This reconstruction must account for systematic uncertainties that affect the observed annihilation coordinates. In this section, we describe the dominant sources of systematic uncertainty influencing (z_a, t_a).

To incorporate these uncertainties into the binning-based reconstruction, we employ a Monte Carlo error-propagation technique. In this approach, the reconstruction procedure is repeated over many iterations. In each iteration, the annihilation coordinates (z_a, t_a) in both the experimental and simulated datasets are randomly varied according to their prior probability distributions, producing a reconstructed energy distribution for that iteration. The ensemble of reconstructed distributions obtained across all iterations constitutes the posterior distribution of the reconstructed energy spectrum. This posterior ensures that systematic uncertainties, including their correlations, are properly propagated to the final reconstructed energy distributions. The priors, which quantify the uncertainty associated with each systematic effect, are summarized in table 1.

The finite axial resolution of the detector introduces a systematic uncertainty in the reconstruction. As described in section 4, this effect is modeled using a two-component Gaussian mixture with equal weights, characterized by standard deviations of 14 mm and 4 mm. Within the Monte Carlo reconstruction, this resolution is incorporated by applying the corresponding smearing to the simulated z_a values at each iteration using an independent random seed.

Additional uncertainty is introduced by a global axial offset between the SVD detector and the magnetic trap, which has been independently measured to be 3 ± 1 mm. This misalignment is accounted for by applying a Gaussian prior to the experimental z_a values. In each Monte Carlo iteration, a single offset is drawn from this prior and applied uniformly to all annihilation positions, thereby modeling the effect of detector–trap misalignment.

Finally, the reconstructed energy distribution is sensitive to uncertainty in the overall strength of the octupole field. While ECR magnetometry constrains the field on axis, an ideal octupole field vanishes there, meaning it cannot be used to accurately constrain the octupole field. This uncertainty propagates into the simulated annihilation times t_a, which depend strongly on the transverse trap depth (see equation (7)). To constrain this uncertainty, we compare simulated and experimental annihilation time distributions using a Kolmogorov–Smirnov (KS) test to identify scaling factors of the octupole field that yield statistically consistent distributions. However, a degeneracy exists: discrepancies in t_a may result from either incorrect modeling of the field strength or inaccuracies in the assumed $\overline{\mathrm{H}}$ energy distribution. To break this degeneracy, we use a reference population with a well-understood energy distribution.

For this purpose, we use the uncooled $\overline{\mathrm{H}}$ population, which is assumed to follow a truncated Maxwell–Boltzmann distribution at the positron plasma temperature, as discussed in section 2. Because the trap depth is much less than the characteristic formation temperature, the shape of this truncated distribution is largely insensitive to the precise value of the formation temperature, making it a robust reference. To capture the effect of the uncertainty in the positron plasma temperature on this reference population, we simulate uncooled populations for several values within the experimentally constrained range ($\lesssim 7$ K [13, 20]). The comparison of these simulations with experimental t_{a} distributions constrains the octupole field strength.

When comparing simulations to data for this reference population, we find that the axial annihilation distributions are not fully consistent, even when using the uncooled simulation variant that includes octupole wire perturbations that reproduce the observed missing hotspots. The inconsistency manifests in the shape of the axial annihilation distributions: rather than a disagreement in the number of $\overline{\mathrm{H}}$ annihilating at positive versus negative z, there is a difference in the distribution shape. In particular, the experimental data show more $\overline{\mathrm{H}}$ annihilating at smaller $|z|$ than predicted by the simulation.

A plausible source of this discrepancy is misalignment of the Penning–Malmberg electrodes. Electrode misalignments can alter the axial annihilation distribution by modifying the relative transverse trap depth over the axial extent of neighboring electrodes, thereby changing the probability that an escaping $\overline{\mathrm{H}}$ atom annihilates at a given z. Indeed, good agreement is recovered when small misalignments of the Penning–Malmberg electrodes, within mechanical tolerances, are introduced into the simulation. In particular, a radial shift of up to 500 μm is required to reproduce the observed experimental z_{a} distribution.

This observation motivates the inclusion of a fourth systematic uncertainty in the Monte Carlo error propagation to account for possible misalignments between individual electrodes in the Penning–Malmberg stack. To incorporate this effect, the simulations used to construct the probability map for the binning-based reconstruction include electrode configurations that best reproduce the observed z_{a} distribution. Since identifying the optimal configuration requires full simulations, the electrode offsets are restricted to the y direction in order to limit the dimensionality of the parameter space. Consequently, the resulting configurations should be regarded as representative electrode misalignments rather than a unique reconstruction of the true electrode geometry. No attempt is made to optimize the agreement with the azimuthal annihilation distribution, as doing so would require exploration of a substantially larger parameter space. Although the electrode misalignments do modify the azimuthal annihilation pattern, this has negligible impact on the reconstructed energy because the azimuthal annihilation coordinate is not used in the reconstruction.

To account for the residual uncertainty associated with electrode alignment, we assign an additional systematic uncertainty by applying a global axial stretch to the simulated annihilation coordinates. Distortions that effectively produce an axial stretch have the largest impact on the reconstructed energy. Specifically, we introduce a stretch parameter α and transform the simulated coordinates according to $z_{\mathrm{a}} \to \alpha z_{\mathrm{a}}$. A uniform prior is applied to $\alpha \in [0.9,, 1.1]$, corresponding to a 10% axial scaling, which conservatively exceeds the level of discrepancy observed between the experimental and simulated z_{a} distributions.

Within the range of these uncertainties, we find that the simulated and experimental annihilation time and z distributions for the uncooled reference population are statistically consistent. Consequently, the reconstructed axial, transverse, and total energy distributions of the experimental uncooled population agree with those of the corresponding simulated distributions.

Although the trapped $\overline{\mathrm{H}}$ energy distribution is relatively insensitive to the assumed formation temperature (taken to be the positron plasma temperature) due to truncation by the shallow trap depth, the KS comparison between simulated and experimental annihilation time distributions remains highly sensitive to this parameter. This indicates that the annihilation time structure retains sensitivity to the $\overline{\mathrm{H}}$ formation temperature despite the truncation. The observed sensitivity highlights both the excellent statistical precision of the simulations and experimental data and the remarkably low temperature at which $\overline{\mathrm{H}}$ is formed [13]. The consistency observed within the allowed temperature range provides further support for the assumption that the $\overline{\mathrm{H}}$ energy distribution is well described by a truncated Maxwell–Boltzmann distribution.

As table 1 lists two distinct sources of uncertainty affecting the simulated z_{a}, the effective prior on the simulated z_{a} is given by the convolution of these contributions. We do not quote individual uncertainty magnitudes for each source, as they are found to be highly correlated. We note, however, that the dominant uncertainty currently limiting the STAM technique arises from the uncertainty in the octupole magnetic field strength.

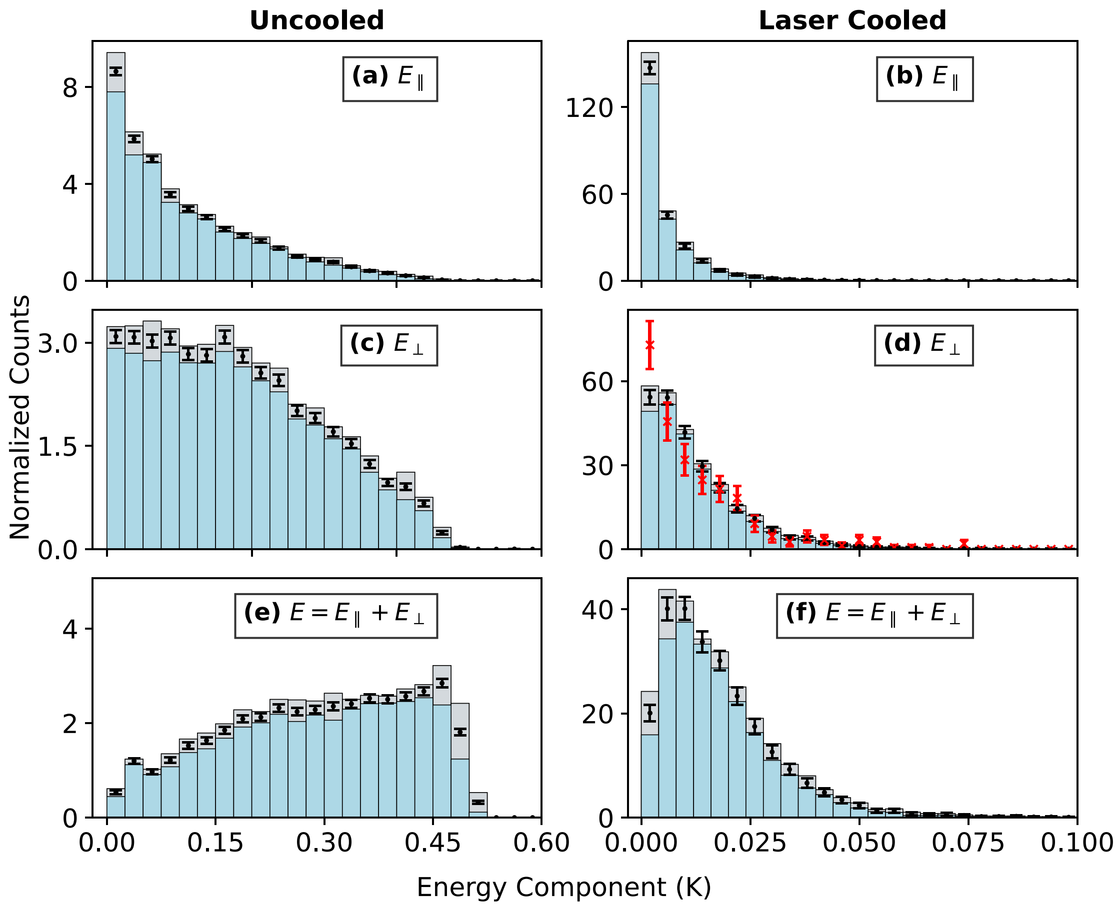

Figure 7. Reconstructed energy distributions of experimental $\overline{\mathrm{H}}$ populations. The reconstruction is performed on the experimental annihilation distributions shown in figures 6(b) and (d). Histogram bins are generated using the Monte Carlo error-propagation technique described in the text to produce the posterior energy distributions, which incorporate systematic effects. The grey band represents the 68 % credible interval of the posterior distributions (systematics only). The black error bars indicate the statistical uncertainty and are centered on the mean of the posterior distribution containing only systematic contributions. Uncertainties due to finite sampling in the binning-based method are negligible. The red crosses show the transverse energy measured using laser-based time-of-flight methods. The associated error bars are statistical only.

7.2. Reconstructed energy distributions of experimental data

In figure 7, we show histograms of the reconstructed energy distributions obtained from posterior samples generated using the Monte Carlo error-propagation procedure, which incorporates the systematic effects described in section 7.1. The reconstruction uses the binning-based method outlined in section 5, with a temporal bandwidth $\mathrm{d}t = 1.5\,\mathrm{s}$ and a spatial bandwidth $dz = 5\,\mathrm{mm}$. Within these two-dimensional bands, the average number of simulated events is 21 for the uncooled data and 170 for the laser-cooled data, reflecting the greater sparsity of the uncooled sample. Accordingly, we randomly sample $N_{\mathrm{samples}} = 5$ and $N_{\mathrm{samples}} = 20$ simulated events per band for the uncooled and laser-cooled cases, respectively. We find that the uncertainty associated with the finite random sampling is small compared to the statistical uncertainty in both cases.

For the uncooled dataset, we obtain mean axial, transverse, and total energies of $109 \pm 1(\mathrm{stat})^{+3}_{-4}(\mathrm{sys})\,\mathrm{mK}$, $180 \pm 1(\mathrm{stat})^{+4}_{-5}(\mathrm{sys})\,\mathrm{mK}$, and $290 \pm 1(\mathrm{stat})^{+8}_{-7}(\mathrm{sys})\,\mathrm{mK}$, respectively. We find that the reconstructed energy distributions corresponding to the uncooled dataset are consistent with those of the uncooled simulation within the measurement uncertainty. For the laser-cooled dataset, the corresponding mean energy components are $6.3 \pm 0.3(\mathrm{stat})^{+0.6}_{-0.6}(\mathrm{sys})\,\mathrm{mK}$, $14.3 \pm 0.5(\mathrm{stat})^{+0.7}_{-0.7}(\mathrm{sys})\,\mathrm{mK}$, and $20.6 \pm 0.7(\mathrm{stat})^{+1.5}_{-1.2}(\mathrm{sys})\,\mathrm{mK}$ respectively. The quoted systematic uncertainties correspond to the one–sigma credible intervals of the posterior energy distributions, while the statistical uncertainties are given by the standard error on the mean.

7.3. Robustness tests and cross-validation with laser-based diagnostics

To further test the robustness of the adiabatic release diagnostic, we performed additional experimental runs in which the energies of both uncooled and laser-cooled $\overline{\mathrm{H}}$ populations were diagnosed using the same procedure, but with a reduced octupole rampdown time of 18 s instead of the nominal 180 s. This shorter rampdown duration still ensures that the magnetic potential evolves adiabatically, according to equation (3). We find that the results obtained from the 18 s and 180 s datasets are consistent.

We also compared the energies diagnosed using the STAM technique with those obtained from established laser-based time-of-flight measurements (see [11]). An independent experimental run was performed in which a population of $\overline{\mathrm{H}}$ was trapped in the flat configuration and laser cooled using the same protocol as the population diagnosed via STAM. Figure 7(d) compares the transverse energy distributions reconstructed using the STAM and time-of-flight methods. The mean transverse energy determined from the time-of-flight measurement is $15.4 \pm 1.2\,\mathrm{mK}$. The two diagnostics yield statistically consistent distributions.

8. Concluding remarks

We have introduced the STAM technique for measuring the energy distribution of trapped $\overline{\mathrm{H}}$ using the time and location of annihilations during adiabatic reduction of the octupole field, with its reliability evaluated through comparison with established laser-based diagnostics. Experimentally, STAM achieves relative precision around 10% for both transverse and axial mean energies, validated across a wide energy range from $\sim 10\,\mathrm{mK}$ to $100\,\mathrm{mK}$, corresponding to laser-cooled and uncooled $\overline{\mathrm{H}}$ populations, respectively.

Because STAM is implemented during the removal of magnetic confinement, it is inherently destructive, whereas laser-based diagnostics can be applied in a largely non-destructive manner. The two approaches are therefore complementary: STAM provides a comprehensive characterization of the energy distribution for a given preparation protocol, while laser-based methods enable *in situ* monitoring of its stability and reproducibility.

As precision increases in $\overline{\mathrm{H}}$ experiments, this method is expected to play a central role in characterizing progressively lower-energy $\overline{\mathrm{H}}$ populations produced using advanced cooling techniques such as laser cooling [11] and adiabatic expansion [10, 21]. The combined application of these techniques is expected to reach the sub-millikelvin regime, where laser-based diagnostics lose sensitivity due to photon recoil effects, a limitation to which STAM is not subject. Although these techniques substantially reduce both the axial and transverse energies of the trapped $\overline{\mathrm{H}}$, adiabatic expansion intrinsically increases the spatial scale of the $\overline{\mathrm{H}}$ trajectories. In the implementation envisioned for future measurements, the expansion is also accompanied by a substantial reduction in trap depth. Consequently, despite the significantly lower energy scale, the characteristic annihilation positions and release times during STAM are expected to remain comparable to those observed here for $10\,\mathrm{mK}$ to $100\,\mathrm{mK}$ $\overline{\mathrm{H}}$. As a result, the systematic uncertainties presented in table 1 are expected to remain at a similar fractional level in this new energy regime. However, future applications in this regime will require careful characterization of the substantially reduced trap depths, and the associated uncertainties in the trapping potential will need to be incorporated into the systematic error budget.

Importantly, because STAM depends on the spatial distribution of the magnetic-field strength, it can also be applied in reverse: if the $\overline{\mathrm{H}}$ energy distribution is known independently with low uncertainty (for instance, through independent measurement using laser-based time-of-flight methods [11]), STAM can be used to constrain the absolute magnetic-field structure itself. However, even without an independently calibrated reference population, STAM remains sensitive to *relative* variations in the magnetic-field strength at the trap wall between different spatial regions. This capability is particularly valuable in the ALPHA experiment, where conventional magnetometry is largely limited to on-axis measurements, leaving off-axis fields poorly constrained. In ALPHA-g, for example, STAM-based release measurements can therefore provide a direct experimental constraint on off-axis magnetic-field systematics that dominate the uncertainty in measurements of the gravitational acceleration of $\overline{\mathrm{H}}$ [5].

Acknowledgments

We dedicate this paper to the memory of Joel Fajans, whose contributions to this work and to the ALPHA collaboration were profound, and whose friendship we will always remember. DH and JSW thank Andrew Charman for his support and encouragement during the preparation of this manuscript. This work was supported by: CNPq, FAPERJ, RENAFAE (Brazil); NSERC, NRC/TRIUMF, EHPDS/EHDRS, CFI, DRAC (Canada); FNU (Nice Centre), Carlsberg Foundation (Denmark); STFC, EPSRC, the Royal Society and the Leverhulme Trust (UK); DOE, NSF (USA); ISF (Israel); and the Swedish Research Council (VR) through Grant 2021–04005.

Data availability statement

The data cannot be made publicly available upon publication because they are not available in a format that is sufficiently accessible or reusable by other researchers. However, data may be made available upon reasonable request where appropriate contextual information and support can be provided.

Author contributions

All authors are members of the ALPHA Collaboration at CERN The work is based on data collected with the ALPHA-2 antihydrogen trapping apparatus, which was designed, constructed, and operated by the ALPHA Collaboration. All authors contributed to data taking through shift work on the ALPHA experiment. D H initiated the study, led the experimental investigation, and carried out the primary data analysis. **Conceptualization**: D H, E T W, W A B, J F. **Methodology**: D H, E T W, W A B, J F. **Software**: D H, W A B, J T K M, F R. **Analysis and validation**: D H, F R, J S H, W A B, J F, C R, M F, N M, S E, J N, A P, A G S, A O, C S, I C, L G, J S, M C, J S W, A J U J, A E, D P V D W. **Supervision**: W A B, J F, J S W. The manuscript was written by D H, edited by J S W, M C, C R, J S H, N M, and W A B, and refined and improved through collaboration-wide review.

Appendix A. Testing the model for the thermal 10 mK population

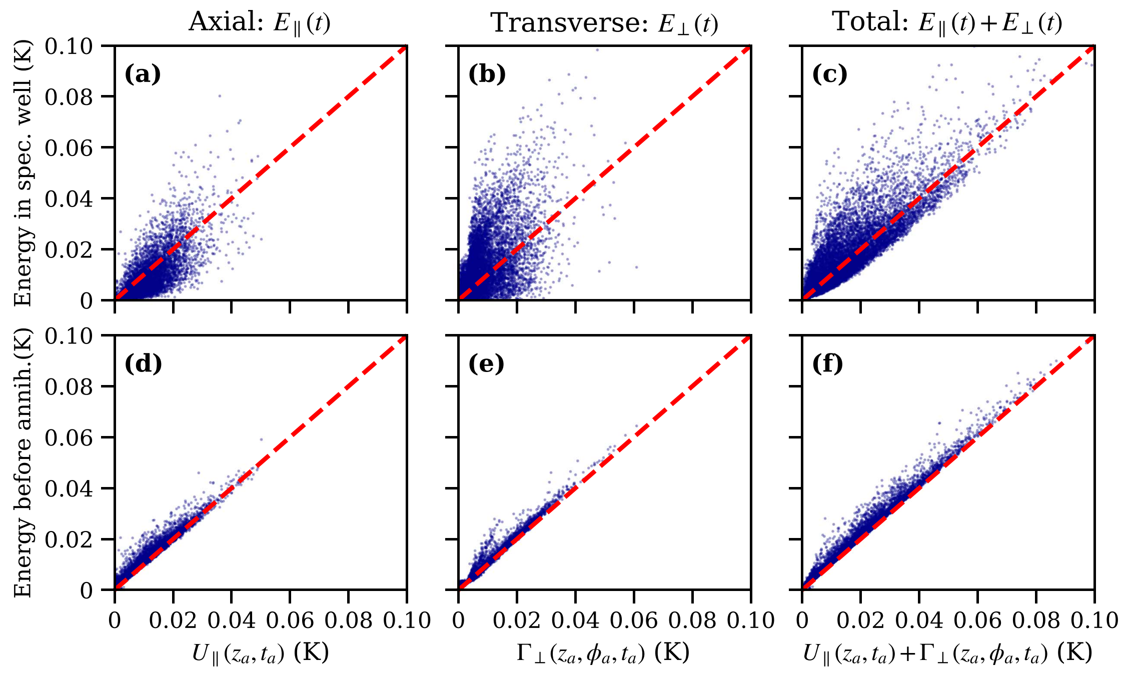

Figure A1. Equivalent to figure 4 for the thermal 10 mK population.

Appendix B. Computation of simulated magnetic fields

The magnetic fields of the octupole and mirror coils are computed using a Biot–Savart model based on the as-designed magnet geometries. The magnetic fields from the end solenoids are computed separately using an off-axis expansion technique [23]. The simulated magnetic fields are matched to the experiment by comparing the on-axis magnetic-field profile with measurements obtained using electron-cyclotron-resonance (ECR) magnetometry [24]. This comparison is restricted to the region $|z| \lesssim 6$ cm, where the relatively low magnetic-field gradient enables high-precision ECR measurements. Agreement between simulation and experiment is optimized by making small (few-millimeter) adjustments to the axial positions of the modeled solenoids and mirror coils, and by applying an overall scaling factor to the magnetic-field amplitudes. After this tuning, the simulated on-axis fields agree with the ECR measurements to within 5×10^{-4} T, corresponding to an uncertainty of approximately 3×10^{-4} K in the axial magnetic potential.

To make the simulations computationally viable, the magnetic fields from the Biot–Savart calculation of the octupole and mirror coils are precomputed on a grid and calculated at the $\overline{\mathrm{H}}$ location using a tricubic interpolation. This interpolation introduces a maximum error of 0.5 % in the trap magnetic fields within the $\overline{\mathrm{H}}$ confinement region relative to the direct Biot-Savart calculation.

ORCID iDs

C J Baker 0000-0002-9448-8419
M Charlton 0000-0002-9754-1932
D Hodgkinson 0000-0003-3410-5540
S A Jones 0000-0001-8205-2186
S Jonsell 0000-0003-4969-1714
N Madsen 0000-0002-7372-0784
T Momose 0000-0001-8976-1938
J Nauta 0000-0002-3477-7104
A N Oliveira 0000-0003-1567-7981
C Ø Rasmussen 0000-0002-6029-1730
F Robicheaux 0000-0002-8054-6040
J Schoonwater 0000-0001-5152-4830

References

[1] Ahmadi M, Alves B X R, Baker C J, Bertsche W, Capra A, Carruth C, Cesar C L, Charlton M, Cohen S and Collister R 2020 Investigation of the fine structure of antihydrogen *Nature* **578** 375–80
[2] Ahmadi M, Alves B X R, Baker C J, Bertsche W, Capra A, Carruth C, Cesar C L, Charlton M, Cohen S and Collister R 2018 Observation of the 1S–2P Lyman-α transition in antihydrogen *Nature* **56** 211–5
[3] Ahmadi M, Alves B X R, Baker C J, Bertsche W, Capra A, Carruth C, Cesar C L, Charlton M, Cohen S and Collister R 2018 Characterization of the 1S-2S transition in antihydrogen *Nature* **557** 71–75
[4] Ahmadi M, Baquero-Ruiz M, Bertsche W, Butler E, Capra A, Carruth C, Cesar C L, Charlton M, Charman A E and Eriksson S 2016 An improved limit on the charge of antihydrogen from stochastic acceleration *Nature* **529** 373–6
[5] Anderson E K, Baker C J, Bertsche W, Bhatt N M, Bonomi G, Capra A, Carli I, Cesar C L, Charlton M and Christensen A 2023 Observation of the effect of gravity on the motion of antimatter *Nature* **621** 716–22
[6] Rasmussen C Ø, Madsen N and Robicheaux F 2017 Aspects of 1s-2s spectroscopy of trapped antihydrogen atoms *J. Phys. B: At. Mol. Opt. Phys.* **50** 184002
[7] Zhong M, Fajans J and Zukor A F 2018 Axial to transverse energy mixing dynamics in octupole-based magnetostatic antihydrogen traps *New J. Phys.* **20** 053003
[8] Andresen G B, Ashkezari M D, Baquero-Ruiz M, Bertsche W, Bowe P D, Butler E, Cesar C L, Charlton M, Deller A and Eriksson S 2011 Confinement of antihydrogen for 1,000 seconds *Nat. Phys.* **7** 558–64
[9] Amole C, Andresen G B, Ashkezari M D, Baquero-Ruiz M, Bertsche W, Butler E, Cesar C L, Chapman S, Charlton M and Deller A 2012 Discriminating between antihydrogen and mirror-trapped antiprotons in a minimum-B trap *New J. Phys.* **14** 015010
[10] Ahmadi M, Alves B X R, Baker C J, Bertsche W, Capra A, Cohen S, Torkzaban C, Cesar C L, Charlton M and Collister R 2024 Adiabatic expansion cooling of antihydrogen *Phys. Rev. Res.* **6** L032065
[11] Baker C J, Bertsche W, Capra A, Carruth C, Cesar C L, Charlton M, Christensen A, Collister R and Mathad A C 2021 Laser cooling of antihydrogen atoms *Nature* **592** 35–42
[12] Capra A 2019 Lifetime of magnetically trapped antihydrogen in alpha *Hyperfine Interact.* **240** 1–6
[13] Akbari R, de Araujo Azevedo L O, Baker C J, Bertsche W, Bhatt N M, Bonomi G, Capra A, Carli I, Cesar C L and Charlton M 2025 Be+ assisted, simultaneous confinement of more than 15000 antihydrogen atoms *Nat. Commun.* **16** 1
[14] Taylor C L, Zhang J and Robicheaux F 2006 Cooling of Rydberg antihydrogen during radiative cascade *J. Phys. B: At. Mol. Opt. Phys.* **39** 4945–59
[15] Andresen G B, Ashkezari M D, Bertsche W, Bowe P D, Butler E, Cesar C L, Chapman S, Charlton M, Deller A and Eriksson S 2012 Antihydrogen annihilation reconstruction with the ALPHA silicon detector *Nucl. Instrum. Methods Phys. Res. A* **684** 73–81
[16] Amole C, Andresen G B, Ashkezari M D, Baquero-Ruiz M, Bertsche W, Burrows C, Butler E, Capra A, Cesar C L and Chapman S 2013 Silicon vertex detector upgrade in the ALPHA experiment *Nucl. Instrum. Methods Phys. Res. A* **732** 134–6
[17] Hurt J L, Carpenter P T, Taylor C L and Robicheaux F 2008 Positron and electron collisions with anti-protons in strong magnetic fields *J. Phys. B: At. Mol. Opt. Phys.* **41** 165206
[18] Robicheaux F 2008 Atomic processes in antihydrogen experiments: a theoretical and computational perspective *J. Phys. B: At. Mol. Opt. Phys.* **41** 192001
[19] Zhong A, Fajans J and Wurtele J S 2024 Shot-noise-induced lower temperature limit of the nonneutral plasma parallel temperature diagnostic *EPJ Tech. Instrum.* **11** 6

[20] Baker C J, Bertsche W, Capra A, Cesar C L, Charlton M, Mathad A C, Eriksson S, Evans A and Evetts N 2021 Sympathetic cooling of positrons to cryogenic temperatures for antihydrogen production *Nat. Commun.* **12** 6139
[21] Hodgkinson D 2022 On the dynamics of adiabatically cooled antihydrogen in an octupole-based ioffe-pritchard magnetic trap *PhD* (University of Manchester)
[22] Ahmadi M, Alves B X R, Baker C J, Bertsche W, Butler E, Capra A, Carruth C, Cesar C L, Charlton M and Cohen S 2017 Observation of the hyperfine spectrum of antihydrogen *Nature* **548** 66
[23] Mcdonald K Expansion of an axially symmetric, static magnetic field in terms of its axial field (available at: http://kirkmcd.princeton.edu/examples/axial.pdf)
[24] Amole C, Ashkezari M D, Baquero-Ruiz M, Bertsche W, Butler E, Capra A, Cesar C L, Charlton M, Deller A and Evetts N 2014 In situ electromagnetic field diagnostics with an electron plasma in a penning– malmberg trap *New J. Phys.* **16** 013037