Aura: Inside-out Electromagnetic Controller Tracking

Eric Whitmire*
Farshid Salemi Parizi*
emwhit@cs.washington.edu
farshid@cs.washington.edu
Paul G. Allen School of Computer Science & Engineering
Seattle, Washington

Shwetak Patel
shwetak@cs.washington.edu
Paul G. Allen School of Computer Science & Engineering
Seattle, Washington

ABSTRACT

The ability to track handheld controllers in 3D space is critical for interaction with head-mounted displays, such as those used in virtual and augmented reality systems. Today's systems commonly rely on dedicated infrastructure to track the controller or only provide inertial-based rotational tracking, which severely limits the user experience. Optical inside-out systems offer mobility but require line-of-sight and bulky tracking rings, which limit the ubiquity of these devices. In this work, we present Aura, an inside-out electromagnetic 6-DoF tracking system for handheld controllers. The tracking system consists of three coils embedded in a head-mounted display and a set of orthogonal receiver coils embedded in a handheld controller. We propose a novel closed-form and computationally simple tracking approach to reconstruct position and orientation in real time. Our handheld controller is small enough to fit in a pocket and consumes 45 mW of power, allowing it to operate for multiple days on a typical battery. An evaluation study demonstrates that Aura achieves a median tracking error of 5.5 mm and 0.8° in 3D space within arm's reach.

CCS CONCEPTS

• **Human-centered computing** → **Interaction devices**; **Mobile devices**; • **Hardware** → **Sensor devices and platforms**.

KEYWORDS

Electromagnetic tracking, virtual reality, mixed reality, controller, head-mounted display

ACM Reference Format:
Eric Whitmire, Farshid Salemi Parizi, and Shwetak Patel. 2019. Aura: Inside-out Electromagnetic Controller Tracking. In *The 17th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '19), June 17–21, 2019, Seoul, Republic of Korea.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3307334.3326090

*Both authors contributed equally to this research.

MobiSys '19, June 17–21, 2019, Seoul, Republic of Korea

ACM ISBN 978-1-4503-6661-8/19/06…$15.00
https://doi.org/10.1145/3307334.3326090

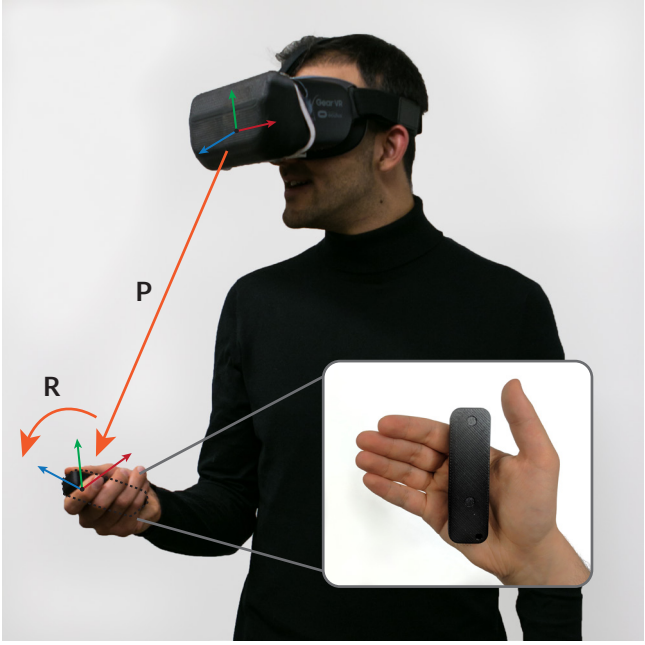

Figure 1: Aura is a 6-DoF electromagnetic tracking system for small, handheld controllers. It estimates the pose of the controller with respect to the head-mounted display.

1 INTRODUCTION

Head-mounted displays (HMDs) for virtual (VR) and augmented (AR) reality represent a promising direction for next-generation computing platforms. Although early iterations of such devices are intended for use in fixed environments, the promise of mobility offers exciting new opportunities for the future of personal computing. The handheld controller is a useful input device and offers advantages such as physical buttons with tactile feedback, a platform for rich haptic feedback [5, 35–37], and, most importantly, a mechanism for precise tracking. Controller tracking is classified as either 3-DoF rotational tracking (i.e. roll, pitch, yaw) or 6-DoF rotational and positional tracking (i.e. roll, pitch, yaw, x, y, z). The use of 6-DoF positional tracking enables an additional class of spatial computing applications by allowing the controller to serve as a virtual tool or a proxy for the user's hand.

Today, there is a significant divide between handheld controllers for mobile VR/AR platforms and those for high-end desktop VR systems. Mobile systems, such as the Samsung Gear VR or Google Daydream, use a small handheld controller that relies on 3-DoF

inertial orientation tracking, with limited positional tracking support. On the other hand, desktop VR systems like the Oculus Rift, use a larger controller with external light emitting diodes (LEDs) for 6-DoF optical tracking. These external elements are tracked using cameras placed in the environment. Though these controllers are much larger, the positional tracking they offer results in a substantially more immersive experience compared to the simple orientation tracking on mobile VR. Efforts to remove the need for environmental infrastructure often move the tracking cameras to the head. This approach leads to additional HMD power consumption, limits cameras' line of sight, and still requires bulky tracking elements on the controller.

If head-mounted displays are to become a compelling personal computing platform with applications beyond gaming, they must support mobility and robust usage. Handheld controllers must rely on inside-out tracking to enable mobility while maintaining low power consumption for extended use on battery power. The form factor should be small enough to fit within the hand during use and in a pocket or bag when not in use. They should also support a robust set of interaction scenarios and maintain usability outside of the view of a head-mounted camera—for example, with the hands at the side on a crowded bus or under a table during a meeting. In this work, we seek to close the gap between the attractive form-factor and mobility of 3-DoF inertial controllers and the performance and usability of high-end 6-DoF controllers.

We propose Aura, a novel low-power electromagnetic tracking technique to bring high-precision 6-DoF controllers to any head-mounted mixed-reality system without the need for line-of-sight, bulky tracking rings, or environmental sensors. Our proposed tracking system uses three coils embedded in a head-mounted display that each generate a unique magnetic field oscillating at 100 kHz. The generated fields induce a sinusoidal voltage in orthogonal receiver coils embedded in a handheld controller. As the user moves the controller, the signal in each receiver coil varies depending on its position and orientation within the field.

In a traditional electromagnetic tracking system, one would carefully construct a 3-axis dipole transmitter and place it far away from any other metallic objects. This significantly limits the design space of HMDs, which strive to be small and lightweight. Unlike traditional electromagnetic tracking systems, Aura makes no assumptions about the size, shape, or position of the transmit coils and no assumptions about nearby ferromagnetic material on the headset. This freedom enables custom coil shapes and configurations that open the design space for small and lightweight form factors. To demonstrate this capability, Aura foregoes the use of precisely manufactured orthogonal transmit coils and explicitly uses irregularly shaped coils designed to fit the contours of a Samsung Gear VR headset. The use of irregularly-shaped coils significantly complicates the pose estimation task. To solve this challenge, Aura uses a hybrid tracking approach that leverages a physics model to calibrate the sensor coils and a closed-form data-driven model using neural networks to directly estimate pose from the calibrated sensor data. While electromagnetic tracking has a rich history [4, 9, 16, 17, 20, 24, 29], to our knowledge, Aura is the first system to demonstrate 6-DoF pose estimation using irregularly-shaped, non-orthogonal coils.

Aura is inside-out; that is, it tracks the position of the controller with respect to the head. For high-end AR or VR systems with positional head-tracking, e.g. using inside-out SLAM-based tracking [11], Aura provides an upgraded controller tracking experience by enabling more robust and subtle usage without the need for line-of-sight or bulky tracking rings. On low-end VR systems that rely on inertial head-tracking, the ability to locate the hand with respect to the eyes still enables many new interactive experiences. On such devices, Aura serves as a snap-on upgrade that provides positional tracking of the controller. Positional head-to-controller tracking allows a VR system to render the hand or virtual objects within the hand at the correct visual position and orientation no matter how the user moves their head. This capability would allow the user to perform any of the pointing-based interactions common in today's VR applications, without having to frequently recalibrate to compensate for drift. Moreover, it would allow the user to directly manipulate objects locked or loosely locked to the body. For example, a user might reach out and type on a virtual keyboard placed directly in front of them, a task that is nearly impossible with inertial tracking. We note that on low-end systems with only inertial head-tracking, the ability to directly grasp other objects located in the world is still limited by the system's head-tracking capabilities.

In the following sections, we provide a brief overview of electromagnetic tracking techniques and design decisions we made in Aura, implementation and calibration details for the Aura system, and results from a system evaluation and characterization. Our results demonstrate that Aura can track a handheld controller with millimeter accuracy.

Specifically, our contributions include:

(1) An efficient, novel, closed-form tracking algorithm that works with arbitrary transmitter coil shapes and configurations and accounts for static magnetic field distortions
(2) A low-power hardware architecture for a 6-DoF handheld controller
(3) A prototype implementation of the Aura system and evaluation of tracking accuracy that demonstrates a median 3D error of 5.5 mm and 0.8°.

2 MAGNETIC TRACKING BACKGROUND

In general, magnetic tracking systems rely on two types of sources: the permanent magnet [19], which generates a static magnetic field, and alternating current (AC) electromagnetic coils [32, 33]. While permanent magnet systems are affected by the Earth's geomagnetic field, AC electromagnetic coils generate a magnetic field at a particular frequency that can be more easily measured by filtering out all other frequencies.

A typical AC tracking system consists of a three-axis magnetic field generator that produces an oscillating magnetic field and a sensor that measures the local magnetic flux density, from which the sensor's location can be estimated. Using an alternating current to drive an electromagnet formed by a wire coil is an effective method of generating a magnetic field that oscillates at a particular frequency. According to Maxwell's equations, an electric current flowing along a wire coil will generate a magnetic field. The oscillating magnetic flux from these generator coils intersects the

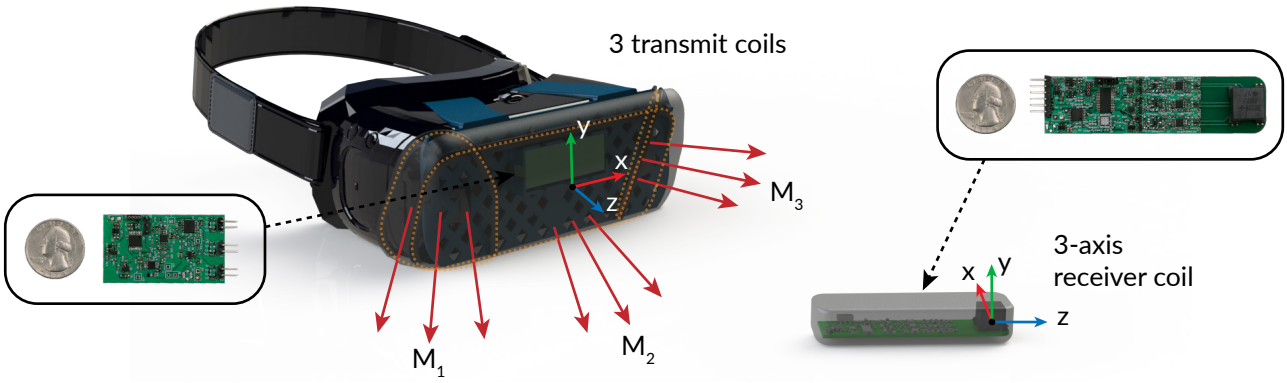

Figure 2: Three transmitter coils embedded in the headset produce three magnetic fields (M_1, M_2, M_3). A three-axis coil inside the handheld controller measures the fields.

sensor coils, inducing a voltage of the same frequency according to Faraday's law of induction. The voltage induced in the coil is proportional to the rate of change of the magnetic flux through the surface bounded by the coil and the number of windings in the coils.

The flux through the coil depends on its orientation within the magnetic field. If the coil is aligned with the field—that is, the normal vector to the coil is aligned with the field—then the flux and the magnitude of the induced voltage will be maximal. As the coil rotates away from the field, the induced voltage decreases to zero. If the coil flips around, the voltage acquires a 180° phase shift, which could manifest as a negative amplitude, depending on the measurement technique.

To track position and orientation in real time, researchers have historically relied on tracking pose changes [29] or on iterative algorithms that find the pose which best explains the observed sensor values [26]. These systems are usually realized by approximating the magnetic sources as dipoles and, in multi-axis systems, as orthogonal dipoles. These pose estimation approaches rely on analytical or numerical analysis of forward models, which describe the magnetic field at the sensor as a function of pose. More recently, closed-form solutions [13] have been developed that analytically invert the forward models. In all of these approaches, deviations from these ideal models cause inaccuracies in the estimated pose.

These assumptions lead to traditional electromagnetic tracking systems that use large high-power transmitter coils that can be approximated as dipoles and are intended to be placed in the environment away from any metallic elements. In contrast, Aura is designed for use with arbitrary coil shapes embedded into the frame of a head-mounted display. In the Aura system, not only can the coils not be modeled as dipoles due to their shape, the presence of nearby electronics in the HMD causes static distortions to the magnetic field. Because expressing an analytic form for either the forward or reverse model is intractable, we develop closed-form data-driven models that directly approximate the reverse model.

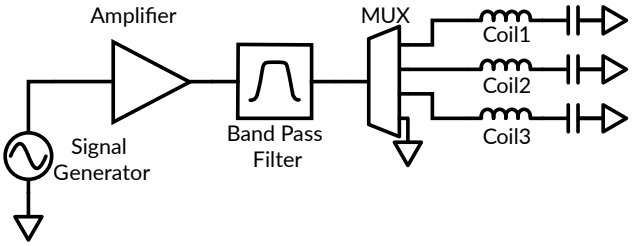

Figure 3: The transmitter filters the output of a waveform generator and feeds the signal through each of the transmitter coils sequentially.

3 SYSTEM IMPLEMENTATION

The Aura system consists of an HMD attachment that generates magnetic fields and a controller that measures the resulting fields. Figure 2 depicts the Aura system components. The following sections provide details of the Aura hardware and explore the capabilities and design challenges of the devices.

3.1 Transmitter

Aura's transmitter consists of three low-profile generator coils embedded in a head-mounted display that sequentially emit a magnetic field oscillating at 100 kHz. Each of the side generator coils consists of 30-40 turns of 22 AWG magnet wire wound around a 3D printed ABS frame. The central coil is wrapped in an oval shape of size 20 cm × 7.5 cm. The inductances for the three coils are measured as 358 µH, 476 µH and 279 µH.

One of the contributions of our work is the use of non-orthogonal coils for the transmitter. This allows Aura to have a configuration of coils that can be fit into any HMD. The coils are rigidly mounted on a 3D-printed support structure such that the magnetic flux is directed toward the user's hand. Our implementation is designed to fit on a Samsung Gear VR (Figure 2) but could easily be modified to fit other HMD designs.

Figure 3 shows the block diagram of Aura's transmitter. The transmitter uses a programmable waveform generator (AD9833)

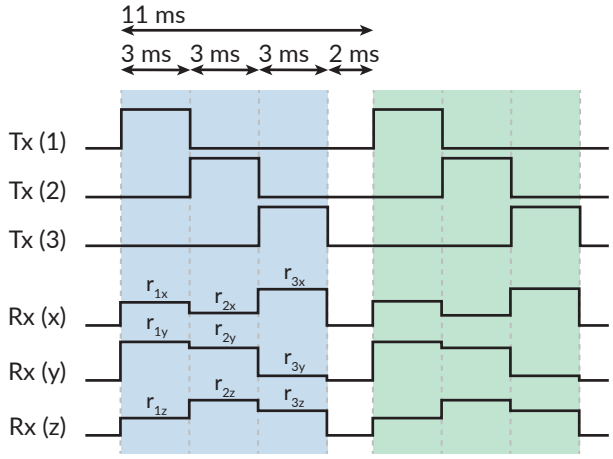

Figure 4: Each transmitter coil is activated for 3 ms, during which each of the three orthogonal receiver coils measures the magnetic flux.

to generate a 100 kHz square wave. This frequency was chosen to enhance sensitivity without approaching any of the coils' self-resonant frequencies or inducing troublesome eddy currents in nearby metallic objects. This signal is passed through a two-stage active bandpass filter (AD8616) with a Q-factor of 15.9 and gain of 0.5 dB. The resulting sinusoid is time-multiplexed (ADG1604) and fed through each of the generator coils.

Each coil generates a magnetic field for 3 ms in sequence; once the cycle is completed, all of the channels are turned off for 2 ms to synchronize the transmitter and receiver (Figure 4). Because Aura cycles through each coil every 11 ms, the resulting frame rate is 91 Hz. An ultra-low power microprocessor (MSP430FR2100) controls the multiplexing and interfaces with the components on the transmitter board. We use passive matching networks to tune the coils impedance, resulting in improved power transmission efficiency at 100 kHz.

3.2 Receiver

While the HMD-based transmitter uses hand-wound coils to generate three unique AC magnetic fields, the controller uses an off-the-shelf three-axis orthogonal receiver coil (Grupo Premo 3DCC08) to reconstruct the 3D magnetic field vector. Because the tracking system relies on interpreting the demodulated sensor measurements as field vectors, it is important to maximize precision and orthogonality in this procedure. The signal from each of the coils is fed to an amplifier (INA826) with a gain of 33.9 dB. The resulting amplified signal is fed to a two-stage active bandpass filter (AD8616) with a Q-factor of 10.2 and a gain of 33.5 dB. Then, we use a low-noise and low-voltage drop Schottky diode network (SMS7630) in a full-wave bridge rectifier configuration to demodulate each of the channels. Finally, each channel is sampled at 4 kHz using the 24-bit sigma-delta ADCs of an MSP430i2031. Figure 5 (Top) summarizes

This method effectively captures the magnitude of the magnetic fields but does not resolve the phase of the oscillating fields. Reconstructing position or rotation would be very difficult from this approach since rotating the controller 180° along one of its axes would result in the same overall magnitude for each of the coils but with a 180° phase shift on two of the axes. Since Aura uses a three-axis receiver and each axis could be in- or out- of phase with the transmitter, there are $2^3 = 8$ possible vectors that would deliver the same sensor values. For each frame of three vector fields (one from each transmit coil), there are a total of $8^3 = 512$ possible rotation states given the same channel magnitudes.

Aura uses a low-power solution to reduce this sign ambiguity. First, comparators binarize each of the amplified signals from the receiver coils prior to rectification. The comparator outputs logic low when the amplified signal of the channel is less than its common-mode voltage ($V_{CM} = 1.2$ V) and logic high otherwise. An XOR gate estimates the relative phase between a receiver channel and a reference signal. The output of the XOR gate is low when the two signals are in-phase and high when the two are out-of-phase. One can produce a referenced signal locked to the transmitter using a phase locked loop, but to save power and simplify the design, we have chosen channel 1 of the receiver as the reference. The resulting logic signal is low-pass filtered by an RC network to remove any glitches due to phase mismatch. These digital signals are then sampled by GPIO pins of the microcontroller unit (MCU). Using this low-power approach, we reduce the ambiguity of the signs to $2^3 = 8$ possible solutions since the phase of the reference channel with respect to the transmitter is still unknown. The exact signs can be determined by placing the controller in a known start state and temporally filtering to ensure consistency over time. Figure 5 (bottom) summarizes Aura's sign detection capabilities.

With the amplitude demodulation and phase estimation components, the MCU now has access to a signed magnitude for each channel every 250 µs (4 kHz). This data stream contains the time

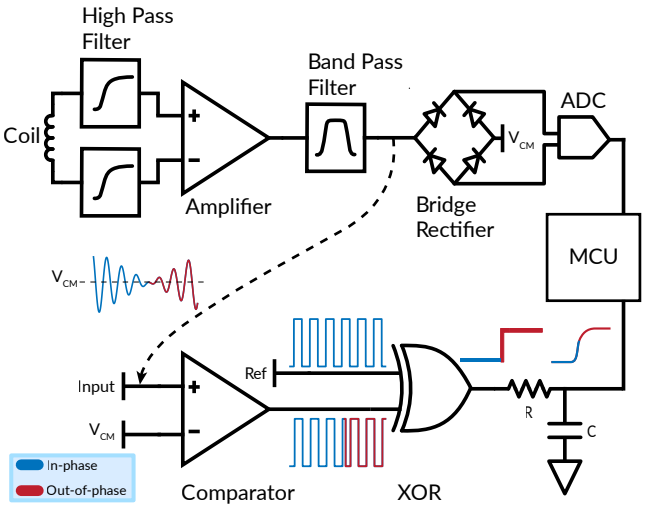

Figure 5: On the controller (single-channel depicted), the the received signal.

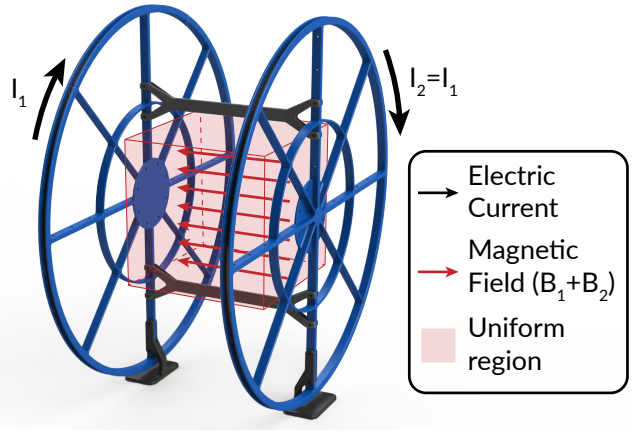

Figure 6: For calibration, Helmholtz coils produce a small volume where the magnetic field is insensitive to movement. The coil radius and separation distance is 15 cm.

multiplexed transmit signals and resembles the diagram in Figure 4. A segmentation step recovers each of the 9 measurements from this data stream. The segmentation algorithm will search for a local minimum to find the "off state" and synchronize itself with the transmitter. It then uses the known timing of the signal to extract the 9 mean values of each coil measurement ($r_{1\{xyz\}}$, $r_{2\{xyz\}}$, and $r_{3\{xyz\}}$). The MCU then sends the nine reconstructed signals to a PC over USB. In software, a digital second-order Butterworth filter with a 10 Hz cutoff frequency is applied to the raw signals for further noise reduction.

4 CALIBRATION

In order to treat the measurements from the ADC as magnetic field strengths, a calibration procedure must be performed to account for imperfections in the signal processing chain and channel gains. This calibration is intended to be performed only once per device, i.e. through a factory calibration step. While calibration in magnetic tracking systems often refers to modeling magnetic field distortions, Aura inherently accounts for this in the tracking algorithm described in Section 5.

To assist in the calibration process, we construct a set of Helmholtz coils as shown in Figure 6. With this device, current through two parallel coils (I_1 and I_2) generate AC magnetic fields. Because of the spacing of the coils, the tangential components of the fields cancel, and the coil pairs create a small volume with a uniform magnetic field, indicated by the red shaded region in Figure 6. Within that region, the magnetic field direction and magnitude are relatively insensitive to the receiver coil's position. By controlling the amplitude of the AC current through the coils with a function generator, we can precisely control the magnetic field strength within this volume.

4.1 Signal chain modeling

In an ideal sensor, the measurement for any axis would be linearly correlated with strength of the magnetic field along that axis. To measure this linearity, we place the Aura controller within the Helmholtz coil and use the function generator to step up the magnetic field strength linearly in small increments and record the Aura measurements. Figure 7 shows the observed signal and the ideal linear response. Though the signal is linear for much of the observed field strengths, significant nonlinearities were observed in the presence of weak magnetic fields.

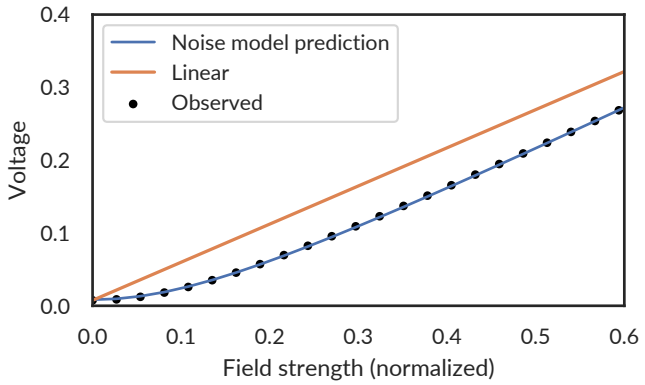

Figure 7: Our proposed model outperforms a linear prediction of the observed signals.

We model these observations using three parameters: a gain term (g), Gaussian noise (n) inherent to the signal chain, and a bias term (b) due to the forward voltage drop across the diodes. Equation 1 summarizes how these effects influence magnetic field strength (f) to produce the Aura measurement (r).

$$r = g\sqrt{f^2 + n^2} - b \tag{1}$$

After collecting data from the Helmholtz coil, we use an optimization procedure (SciPy) to fit this model to the observed data. As shown in Figure 7, this model is a good fit for the observed data. We then invert this model to derive an expression, as shown in Equation 2, for the desired magnetic field strength (f) as a function of the ADC measurements (r). We preserve the sign of the original signal as described in Section 3.

$$f = \sqrt{\left(\frac{r+b}{g}\right)^2 - n^2} \tag{2}$$

At runtime, we further improve the device performance by assuming the bias and gain terms remain constant and dynamically adjusting the noise term of this model based on the signal observed during the 2 ms off period of the device, when $f = 0$. The dynamic noise term is derived according to Equation 3.

$$n = \frac{r_0 + b}{g} \tag{3}$$

4.2 Channel gains

Due to component tolerances, each channel has a slightly different gain. To measure these gains, we place the controller within the Helmholtz coils and rotate it while recording data. In a properly calibrated system, the magnitude of the magnetic field measurement would remain constant as the device is rotated. We use a second

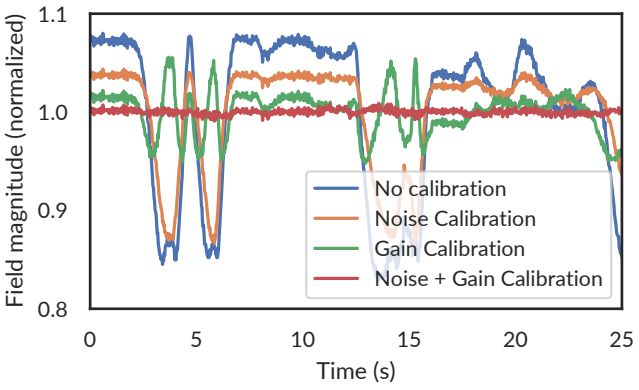

Figure 8: Within the Helmholtz coil, the measured field magnitudes should be constant as the controller is rotated. Without calibration (blue line), there are significant deviations in the field magnitude. As we apply different calibration steps (orange, green, red lines), the calibrated signal significantly improves.

optimizer to learn the optimal gains such that the magnitude of the calibrated field remains constant.

Figure 8 shows the observed magnitude (blue) and calibrated magnitudes after accounting for the signal chain model (orange), channel gains without the signal chain model (green), and both models (red). The constant magnitude obtained after calibration validates these models and shows that both are critical in obtaining correct magnetic field measurements.

5 TRACKING ALGORITHM

Because it is infeasible to construct an analytic formulation of the magnetic fields, there is no closed-form analytic solution for the inverse problem. Instead, we rely on a closed-form data-driven solution based on machine learning techniques. Our proposed closed-form tracking approach is summarized in Figure 9. We treat the process of position and orientation estimation separately. At a high level, we train a neural network to regress from extracted rotation-invariant features to position. We then train a separate network to estimate the magnetic field vectors from position and use SVD to compute orientation.

Unlike other approaches, this pipeline makes no assumptions about the shape of the transmit coils or the resulting field (other than it being separable by the features we extract). This frees us from some of the restrictions of a model-based approach, which excels only when the magnetic field model accurately represents the empirical data. The training methods we describe are intended to be a one-time factory calibration.

Aura must estimate the position of the controller from the three measurements of magnetic field vectors (f_1, f_2, f_3). Because the controller can exist at the same position with any rotation, we first extract six features that are invariant to the rotation of the controller. The first three features (Equation 4) include the magnitude of all measured fields. As the controller rotates, the *magnitude* of each of the three fields remains constant. The second three features (Equation 5) relate to the angle between two fields. As the controller rotates, the direction of each measured field in controller space, f_i will change, but the angle between any two fields will remain constant. Note that in Equation 5, we take the absolute value of the dot product. This is because we do not have an absolute sign reference between our transmitter and receiver coils. That is, we do not know whether an ideal sensor would have measured f_i or $-f_i$. By removing the sign of the dot product, we remain invariant to this ambiguity.

$$|f_i|, \forall i \in \{1, 2, 3\} \tag{4}$$

$$\frac{|f_i \cdot f_j|}{|f_i| \times |f_j|}, \forall (i, j) \in \{(1, 2), (2, 3), (1, 3)\} \tag{5}$$

These features summarize the relative directions and strengths of the three magnetic fields. A cross-sectional slice of two of these features from the simulated dipole dataset are depicted in Figure 10.

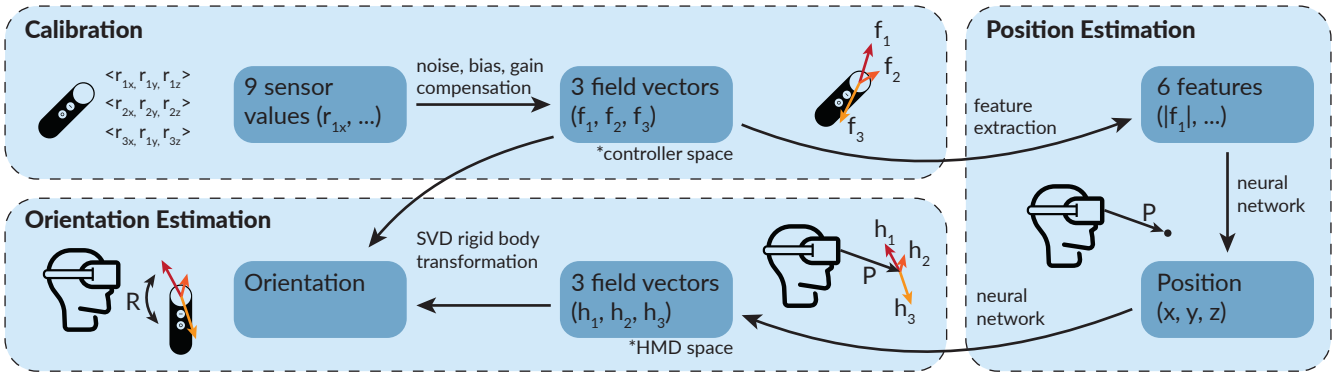

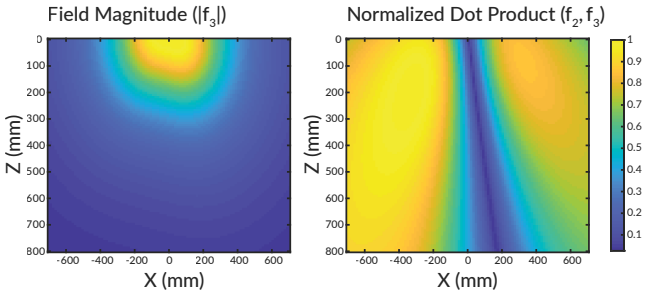

Figure 10: A cross-section at y = -300 mm of two of rotation-invariant features from the simulated dipole dataset. The left plot shows the magnitude of the magnetic field from the side one of the side coils, per Equation 4, and the right plot shows the dot product between the center and side coil, per Equation 5.

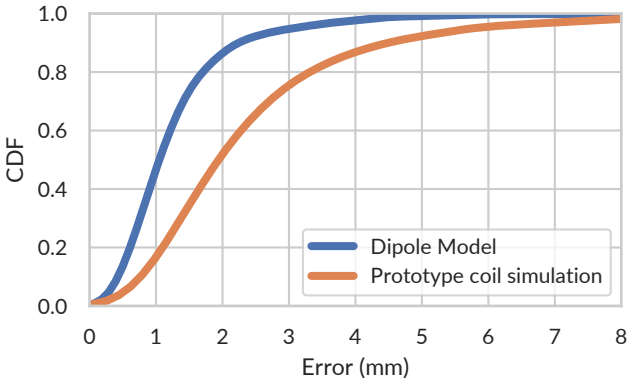

Figure 11: CDF of simulated tracking error for the dataset generated using dipole field models and numeric integration of the actual prototype coils.

5.1 Position

We propose the use of computationally simple models to regress to a position vector. In order to keep the models small, we split the tracking volume into four subspaces along two dimensions: 1) left ($x < 0$ mm) / right ($x > 0$ mm) and 2) near ($|P| < 200$ mm) / far ($|P| > 200$ mm). These volumes were empirically determined to balance model performance and complexity. By reducing the tracking volume for each model, we can train much smaller models than we could if the entire tracking volume were lumped together. For each of the four subspaces, we train a computationally simple neural network model with a single hidden layer of 32 nodes to fit a function that maps the six rotationally invariant features to a 3-dimensional position vector. Training is performed using the Levenberg-Marquardt algorithm. At runtime, position estimation equates to two matrix multiplications (6×32 and 32×3), that can easily run on a mobile processor.

It must be noted that this approach adds a dependency on knowing the position before one of these four models can be chosen (to then estimate position). In this work, we use the ground truth position to pick the correct subspace model in order to validate each model separately. However, since the controller will be tracked over time, the next position can easily be estimated from temporal extrapolation or Kalman filter prediction with an IMU. This rough estimate can be used to choose the appropriate model for that frame. Additional redundancy can be added by training intermediate models that straddle two subspaces.

We simulated our approach using magnetic field simulation tools that rely on quasi-static assumptions to generate two datasets. First, we model our transmitter coils as ideal dipole models and use physics calculations to compute the magnetic field due to each of the three coils at many locations around the head. As a comparison, we also use the BSMag toolbox [28] in MATLAB to model the specific shape of our transmitter coils and create a similar dataset of magnetic field locations.

Each dataset consists of 100,000 points generated within a tracking volume of 1.6 m×0.8 m×0.8 m. Our model is able to estimate the position of the synthetic datasets with a median error of 1.04 mm for the dipole model and 1.9 mm for the simulation of the Aura coil design. Figure 11 shows the positional tracking error CDF for simulated data on the dipole model and Aura prototype model.

5.2 Orientation

We derive the orientation, or attitude, of the controller by comparing the measured magnetic field vectors in controller-space ($\hat{f}_1$, $\hat{f}_2$, $\hat{f}_3$), with estimates of the magnetic field vectors in head-space ($\hat{h}_1$, $\hat{h}_2$, $\hat{h}_3$), derived from a forward model. In some electromagnetic tracking systems, one can compute the magnetic fields ($\hat{h}$) as a function of position using a magnetic field model, often a dipole model [13]. In our system, we anticipate significant deviations from an ideal dipole model due to the presence of electronics around the head, the non-circular nature of our coils, and the arbitrary positioning of the coils for form-factor purposes. Instead, for a forward model, we train a separate neural network to estimate the magnetic field vector at any position, p around the headset. We again adopt a computationally-simple neural network with a single hidden layer of 32 nodes that maps the 3D position, $\hat{p}$, to three 3D vectors, $\hat{h}_1$, $\hat{h}_2$, and $\hat{h}_3$. The model is trained using a ground truth source of position and an estimate of $\hat{h}$ computed by rotating the measured $\hat{f}$ by a ground truth source of orientation. This training process is again intended to be a one-time calibration procedure.

To obtain an orientation estimate at runtime, Aura uses the magnetic field estimates and the well-known singular value decomposition (SVD) method to find the least-squared error between the two coordinate systems. For a complete review of this approach, see Sorkine-Hornung and Rabinovich [34].

6 DATA COLLECTION

6.1 Optical Motion Capture

A ground truth source for position and orientation is required to train and evaluate Aura's tracking models. To this end, we use a seven-camera Vicon motion capture system to record the real-time position and orientation of both the headset and the controller at 240 Hz. To enable tracking, we place retroreflective spheres on both devices in known locations, as shown in Figure 12.

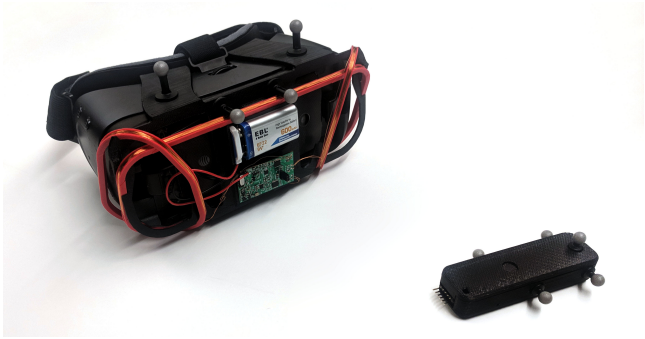

Figure 12: Retroreflective markers placed on the HMD and controller enable tracking with a ground truth optical motion capture system.

Importantly, we define the coordinate system of the controller to be the precise center of the magnetic coil so that a rotation about the origin does not change the position of the sensor. We run a one-time calibration step to find the rigid body transform between the object coordinate systems reported by Vicon and our desired coordinate system. We use the average position of the markers from Vicon and the expected marker positions from the known geometry of our system to derive this transformation, which we then apply to each frame that the Vicon system reports.

Because Aura estimates the relative pose of the controller with respect to the head, we use standard coordinate system transformations to compute this 6-DoF pose from the pose of each device in room coordinates.

6.2 Synchronization

Once data has been recorded from both the 91 Hz Aura system and the 240 Hz motion capture system, we must align and synchronize the two data streams. Without an electrical synchronization signal between the two, we rely on characteristics of each signal for alignment. For this, we compare the distance between the controller and headset as measured from motion capture with an approximation of the distance from Aura as specified by Equation 6. While this signal does not have a physically significant value, it was found to correlate strongly with the distance from the headset, as shown in Figure 13.

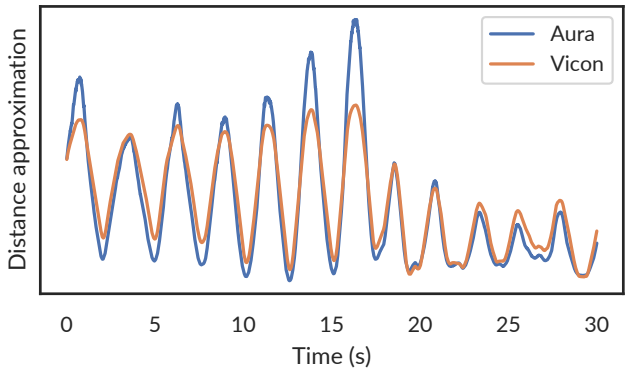

Figure 13: Signals used to compute alignment between Aura and the motion capture system.

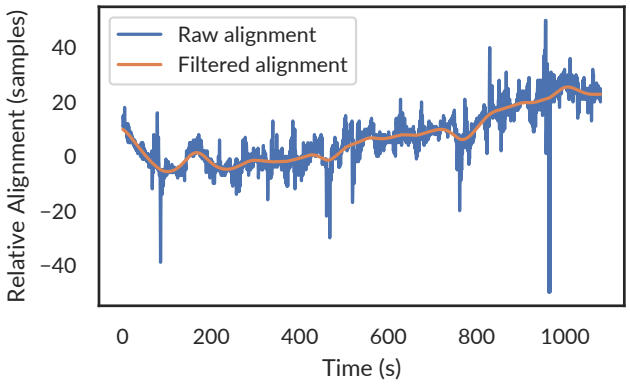

Figure 14: Cross-correlation between the sensor data and the distance from the headset.

$$\frac{1}{\sqrt{\sum_{i=1}^{9} f_i^2}} \tag{6}$$

We first use these signals to achieve alignment at the start and end of the recorded data streams. However, due to effects like temperature changes, it was observed that the two clocks drift relative to each other over time by as much as 1 part per thousand. At a typical hand speed of 10 cm/sec, each frame misalignment represents an additional error of over 1 mm, so mitigating the effects of this drift is essential.

We use sliding cross-correlation between the two signals to compute a dynamic estimate of frame shift and filter this signal for smoothness. Figure 14 shows how the two systems drift over time. We then resample the motion capture data to align with the Aura system.

7 EVALUATION

The main performance metrics for Aura are position and orientation accuracy, precision analysis, latency, and power consumption. We also evaluated the effects of magnetic interference on Aura's measurements.

7.1 Position Estimation Accuracy

7.1.1 2D Position Accuracy. As an initial verification of the tracking capabilities of Aura, we evaluated its positional tracking accuracy in a constrained 2D task. For this experiment, the headset was placed on a flat surface while the controller was manually moved along the surface at a mean distance of 0.5 m from headset. First, training data was collected by sweeping the controller across a 44 cm × 32 cm area. A test set was collected by randomly moving the controller about within the same tracking area.

Figure 15 shows the trace from the Vicon data and the reconstructed path from Aura. The Aura system is able to track the controller with a 2D median tracking error of 1.6 mm. A Kalman filter is used to smooth the estimate of position and reduces the

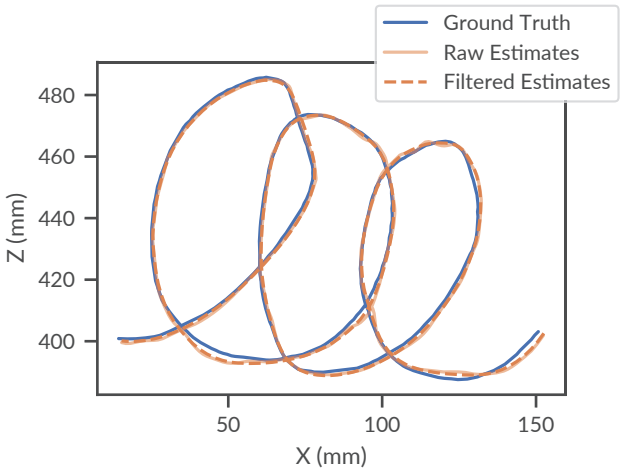

Figure 15: 2D positional tracking performance

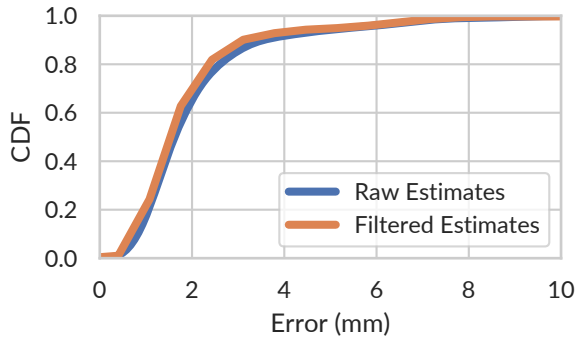

Figure 16: CDF of 2D positional error

median error to 1.5 mm. Figure 16 shows the CDF of 2D position accuracy for the raw and filtered estimated of position.

7.1.2 3D Position Accuracy. While the 2D test demonstrates the feasibility of this approach, Aura is intended to track handheld controllers in 3D space. For evaluating the 3D positional tracking capabilities, one of the authors wore the headset while holding the handheld controller and moving it about within arm's reach. Training data was collected for 15 min by systematically exploring the space in front of the user. An additional two minutes of random motion was then collected as the test set.

Figure 17 shows the ground truth and estimated position over a representative 30 s segment of the test set. Aura is able to track the controller with a 3D median error of 7.0 mm. After applying the Kalman filter, the median error drops to 5.5 mm. Figure 18 shows the CDF of 3D position accuracy for the raw and filtered estimates of position. We expect further performance gains by fusing the electromagnetic tracking with an onboard IMU.

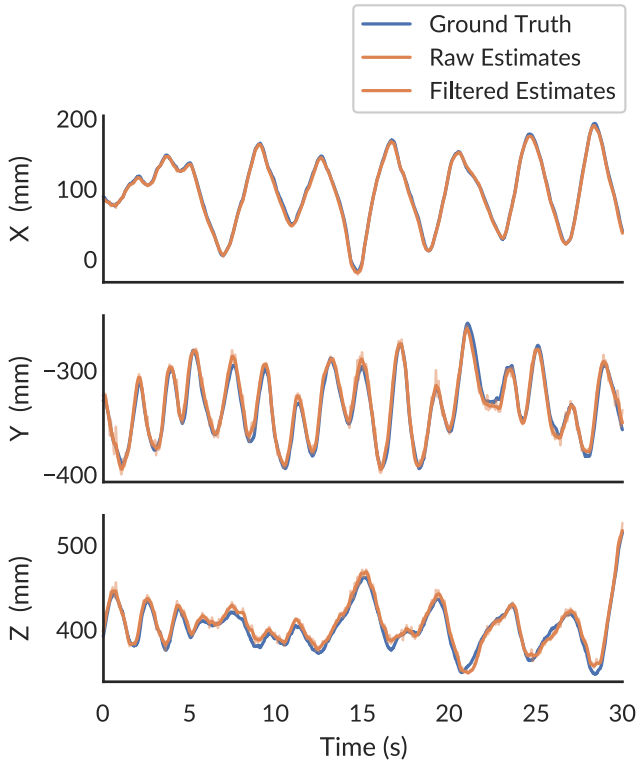

Figure 17: 3D positional tracking performance

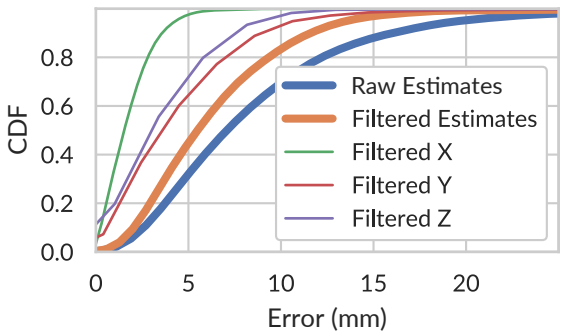

Figure 18: CDF of 3D positional error

7.2 Orientation Accuracy

We use the same dataset to train the magnetic field models for orientation estimation. Aura estimates the orientation of the handheld controller using the algorithm described in Section 5 and reached a median accuracy of 0.8°. Once again, we expect significant performance gains after fusing this data source with an IMU. However, as an initial approximation, we simulate the effects of leveraging an accelerometer to recover the gravity vector. We derive this estimate of the gravity vector from the ground truth motion capture system.

Adding this additional vector reference to the SVD calculation improves the median orientation accuracy to 0.5°. Figure 19 shows the CDF of orientation accuracy for both Aura and the simulated approach using the added gravity vector.

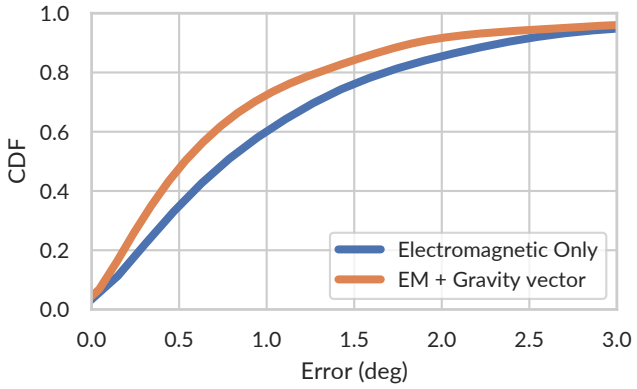

Figure 19: CDF of orientation tracking error

7.3 Precision Analysis

We also evaluated Aura's precision in its tracking estimates. To measure the precision, we used the same set up as in Section 7.1.1. However, instead of continuous motion, the device was placed in eight different locations and left motionless for a few seconds. We then compute the jitter in the estimated position while the device is motionless. The median jitter across all points is calculated to be 0.4 mm. This calculation includes the digital low-pass filter applied to the raw magnetic signals, but no additional Kalman filtering of the position. Figure 20 shows the CDF of jitter.

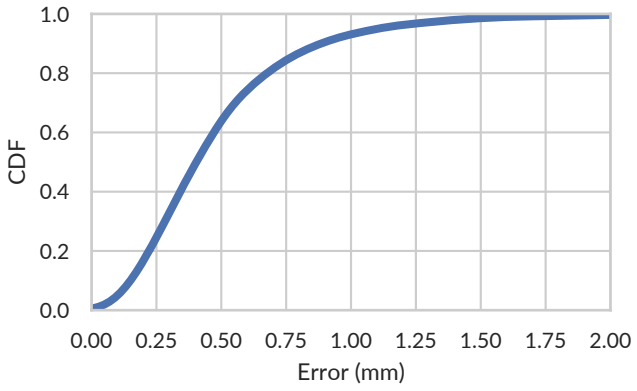

Figure 20: CDF of stationary measurement jitter

7.4 Speed and Latency

Although the tracking analysis was performed offline, it was designed with realtime operation in mind. Like any handheld controller, we envision Aura to be used alongside an onboard IMU to capture high-speed motions. Nonetheless, we characterize the latency and speed aspects of our system.

Aura's latency is impacted by delays introduced by the analog signal chain, digital signal processing, and tracking algorithm. To quantify the impact of the analog signal chain, we place the Aura controller within the Helmholtz coil and activate the field while measuring both the current through the Helmholtz coil and the input the to Aura ADC on a oscilloscope. The time between the field turning on and the signal stabilizing is 200 µs. This introduces negligible latency and validates our decision to use 3 ms "on" states for each coil. On the digital side, the use of interpolation to resample the magnetic signal reduces the impact of the 3 ms delay between channels. An additional digital filter is used to smooth data and can be set according to anticipated device usage. In our prototype, a second-order Butterworth filter with a 10 Hz cutoff frequency was used. Finally, we anticipate negligible latency from the tracking algorithm, which was designed with simple computation in mind. The computation consists primarily of a few small matrix multiplications, which can be performed in real time on most microprocessors.

7.5 Power

We measured the power consumption of both the Aura transmitter and receiver. Our measurement setup consists of a National Instruments (NI) USB 6003 data-acquisition (DAQ) unit configured for taking analog measurements in fully differential mode −10 V to 10 V at a sampling rate of 10 kHz. We measure the voltage drop across a shunt resistor of size 10 Ω for amount of 5 s and calculate the power consumed in Aura.

The handheld controller consumes an average of 13.5 mA (45 mW). Using a 700 mA h LiPo, which would comfortably fit in the controller housing, the Aura handheld controller could be tracked continuously for over two full days. Note that this does not include power consumption for a Bluetooth module or other wireless communication device.

The head-mounted transmitter consumes 29.8 mA average current (224 mW). Using a 9 V battery (600 mA h), the transmitter lasts for 24 h of continuous use. Although the current design of the transmitter uses a 9 V battery to operate, the Aura design is easily modified to operate at 3.3 V. Adding more windings to the transmitter coils and reducing the operating voltage has led to prototype designs that consume only 30 mW on the head-mounted transmitter with similar field strengths. With this design, the Aura transmitter would last nearly four days on a 700 mA h LiPo battery.

For reference, the electromagnetic Polhemus G4 tracking system uses 5 W for the transmitter and 2.5 W for the sensor hub [27]. The HTC Vive Lighthouse base stations, an optical head and controller tracking solution, use approximately 5 W each, as measured by an inline power monitor. The Magic Leap One controller contains an 8.4 W h battery and is rated for 7.5 h of continuous use [18], suggesting a power consumption on the order of 1 W. Aura uses at least an order of magnitude less power than these alternatives.

7.6 Interference

Fundamentally, electromagnetic tracking is prone to interference from nearby metallic objects. While this work does not claim any

Table 1: Effects of various devices on the Aura sensor readings. Each distance represents the closest distance at which the device changes the measured signal by less than 1%.

Device	Distance to transmitter	Distance to receiver
Smartwatch (Apple Watch Series 2)	2 cm	2 cm
Smartphone (iPhone 7)	10 cm	5 cm
Laptop (2018 MacBook Pro)	20 cm	20 cm

specific algorithmic contributions to account for interference, we note that the use of inside-out tracking, as opposed to outside-in tracking, significantly reduces the scope of possible interference sources. Based on our observations, significant distortion was only observed when metallic objects were placed close to either the transmitter or receiver. To quantify these effects on the Aura system, we investigated the interference from three common electronic devices that are likely to be in close proximity to the system: a smartwatch, a smartphone, and a laptop. The Aura controller and headset are placed on a flat surface as in 7.1.1. The authors then bring each interfering device from a far distance to close proximity to either the handheld controller or the HMD while recording the change on the received signals. The distance at which the signals changed by 1% was recorded. The results from this experiment are summarized in Table 1.

In summary, a small electronic device such as a smartwatch, has no effect as long as it is at least a few centimeters from the Aura system. Larger electronic devices must be kept further away before significant distortion is observed. As with any other electromagnetic tracking system, larger ferromagnetic materials, such as iron beams or vehicles, would have a much larger impact on the signal. While accounting for such distortions is out of scope for this work, we note that there is existing research on accounting for sources of interference [14, 15]. We also note the possibility of fusing electromagnetic tracking with inertial or optical tracking to dynamically calibrate in the presence of dynamic sources of interference.

Although not common today, it is important to consider the implications of using multiple devices within the same room. In theory, the signal from one Aura device could interfere with another if they are tuned to the same frequency. Fortunately, the strength of the generated magnetic field falls off with the cube of the distance to the transmitter. We measured the distance from the headset at which the signals fall below the noise floor of the sensor to be 1.5 m. This indicates that there will be negligible interference between separate systems as long as they remain more than 1.5 m apart. For optimal operation at closer distances, the two devices should be set to different frequencies to avoid interfering.

8 DISCUSSION

We demonstrate a 6-DoF tracking system capable of tracking position with a median error of 5.5 mm and a median orientation error of 0.8° within arm's reach around the head while using less than 50 mW on the controller. This approaches the performance of commercial electromagnetic tracking systems, such as the Polhemus G4, while using an order of magnitude less power and allowing optimization for form-factor. This performance does not include integration with an onboard IMU. In a production-grade system, one would use Kalman filtering techniques to fuse the electromagnetic pose estimate with inertial measurements to improve speed and precision. In our simulated result, we demonstrate that with the inclusion of a gravity vector estimate, the rotational error was reduced to 0.5°. We leave the IMU integration to future work.

While our approach demonstrates the feasibility of precise tracking, it does not fully account for all possible sensor positions and orientations. Additional robustness can be achieved by collecting data in all possible configurations, perhaps with the use of a robotic arm. For manufacturing purposes, a separate, externally calibrated sensor can be used to train the magnetic field models, as these depend only on the magnetic field, not on any specific measurement of the Aura device.

By modeling the magnetic fields empirically, Aura can operate near ferromagnetic materials, such as those found within a head-mounted display. However, this approach accounts only for static distortions to the field. Large metallic objects brought nearby the device will degrade tracking performance. Still, because Aura is an inside-out tracking system, metallic objects must be near the head or hand for distortions to occur. Future work can explore techniques to fuse electromagnetic tracking with optical or inertial tracking to maintain accuracy in the presence of nearby distortions. Dynamic distortions caused by particular electronics within the display will likely be localized to a particular frequency and can be eliminated by carefully choosing the frequency of operation.

While Aura was designed with low-power operation in mind, we expect that additional engineering improvements can further reduce power consumption. For example, additional coil windings on the transmitter and the use of a fixed oscillator instead of a programmable function generator can save significant power on the head-mounted transmitter.

One of the advantages of our system over other electromagnetic tracking systems is the ability to use arbitrary transmitter coil configurations. While we designed the Aura prototype as a snap-on device, the transmitter coils could also be placed directly onto a PCB behind the display of a VR system or embedded within the frames of a pair of glasses. By reducing the dependence on orthogonal dipole models, we widen the design space for head-mounted computing systems with tracked devices.

Aura is designed to track a handheld controller, but by eliminating the need for bulky tracking markers, it also opens the door to other kinds of tracked objects. With small engineering improvements, Aura could be used to track accessories to improve the virtual experience, limbs for precise motion capture and avatar reconstruction, or other handheld tools, such as a pen.

Additional performance can likely be achieved by optimizing the placement and shape of the transmitter coils. The Aura coils were designed for the form-factor of a particular HMD, but our simulation results reveal a performance difference between a dipole model and a model of our coils. Leveraging simulation and optimization tools, we expect one can optimize the design for a particular use case. Despite this, the Aura prototype demonstrates reasonable tracking performance without any iteration over transmit coil design.

Aura's tracking models were trained using data collected from an optical motion capture system. However, for researchers who wish to use our system for their own projects, we anticipate that

training can also be performed using commodity low-cost trackers, such as the HTC VIVE Tracker.

9 RELATED WORK

9.1 Inside-out Controller Tracking

Some commercial devices (Windows Mixed Reality, Oculus Quest) use head-mounted cameras that track a controller with an external LED ring. While this technique is precise, it requires the use of bulky controllers and line of sight to the headset, preventing use with arms at the side. The upcoming Vive Focus uses ultrasound and IMU tracking, though there are few details about how this works.

In the research space, a number of different techniques have been explored to eliminate the use of bulky tracking rings and markers. Pocket6 is a solution that uses ARKit on an iPhone X to localize the controller using the built-in SLAM and IMU fusion algorithms [1]. Pandey et al. demonstrate a technique to track markerless controllers using only the front-facing camera on the HMD [22]. These classes of solution tend to demand significant power and computation.

Though not explicitly a controller tracking technique, Shen et al. show a method [31] to track the position of a smartwatch using IMU sensors and a kinematic model. Ultrasound tracking systems have long been used for tracking the head [7, 8]. These rely on pairs of beacons and microphones that use time-of-flight measurements to estimate 6-DoF pose. Though they are usually small and light, they are sensitive to temperature, occlusion, and ultrasound noise. Nandakumar et al. use RF backscatter to track the position of sub-centimeter scale devices [21]. Although they can localize objects tens of meters away, their accuracy is not sufficient for VR controller trackers and requires instrumenting the environment. See Baillot et al. for a more complete review of VR tracking technologies [2].

9.2 Advances in Electromagnetic Tracking

The use of electromagnetic (EM) fields for 6-DoF tracking has a rich history that dates back to the 1980's [16, 29]. This technique is characterized by extremely precise position and orientation tracking. Since then, it has been used in contexts ranging from surgery [9, 17] to biomechanics [20, 24] to localization [4].

In the commercial space, electromagnetic tracking is performed by products from Polhemus, NDI, and Sixense. These products consist of a large transmitter that emits a consistent magnetic field that spans a volume on the order of a cubic meter. The base stations are coupled with small receiver sensors that plug into a large processing hub. In general, such devices offer incredible tracking precision and accuracy, but rely on large infrastructure and distortion-free fields that make integration into a mobile device difficult. For example, the Polhemus G4 [27] tracking system uses a 5 W transmitter and 2.5 W receiver hub. In contrast, the Aura system uses only 224 mW for the transmitter and 49 mW for the receiver.

Magic Leap has recently released a proprietary electromagnetic tracking solution. However, this system uses a multi-frequency transmitter in the controller while placing the receiver coil in the headset. This configuration makes it difficult to scale to multiple controllers, as nearby controllers may interfere with each other. In contrast, Aura uses only a single set of transmitter coils on the head and supports any number of peripheral devices. While Magic Leap's system is proprietary, the relatively large controller and placement of the sensor suggest this system also relies on distortion-free dipole fields. There remains a need for head-mounted inside-out electromagnetic tracking solutions that do not require a particular field structure or large transmit coils. In contrast, Aura consumes significantly less power, has a smaller controller that fits entirely within the hand, and relies on custom-designed transmit coils that make no assumptions about the nature of the resulting field.

Traditionally, EM tracking solutions rely on iterative algorithms, but there have been recent attempts to simplify the computation involved. Ge et al. use rotating transmitters [10] to continuously track the object. Kim et al. show a closed form solution when using a 3-axis dipole generator [13], but only evaluate their approach on a small 3D trajectory with no ground truth reference. These approaches rely on analytic solutions to the forward and reverse problems. Because Aura can operate with any transmitter coil shape, such analytic solutions are infeasible. Unlike these systems, our data-driven solution requires no feedback loop with the transmitter and makes no assumptions about the dipole nature of the transmitter or field distortions.

Unlike most EM systems, which rely on environmentally placed transmitter coils, Pirkl et al. developed a wearable low-power electromagnetic system [25]. However, this solution was only used for gesture recognition, not positional tracking. Roetenberg et al. created a wearable EM tracking system with 5 mm accuracy [30] using a pyramidal structure of transmitter coils.

Some electromagnetic tracking solutions rely on magnetometers instead of field coils. Dai et al. demonstrate an electromagnetic tracking technique using a single transmitter coil and a 3-axis magnetometer [6]. In Finexus, Chen et al. use four magnetometers to track the position of electromagnets placed on the fingertips [3]. Both of these approaches are limited to shorter distances (<20 cm).

Islam et al. show a technique using resonance coupling to improve the efficiency of electromagnetic tracking systems [12]. Our experimentation with this technique suggests that while range is increased, the crosstalk due to the mutual coupling of different coil axes makes field reconstruction difficult. For a more thorough review of magnetic positioning systems, see [23].

10 CONCLUSION

In this work, we present Aura, a head-mounted inside-out tracking system for handheld devices. We demonstrate a novel low-power architecture that enables precise tracking without the need for external infrastructure, line-of-sight, or bulky tracking markers. In an evaluation with an optical motion capture system, we demonstrate Aura's ability to track a controller with a median error of 5.5 mm and 0.8° within arm's reach. We hope that our system enables increased adoption of mobile spatial computing systems.

ACKNOWLEDGMENTS

The authors would like to thank Svet Kolev, the Movement Control Laboratory, and the Global Innovation Exchange for access to their motion capture systems as well as Manoj Gulati and Elliot Saba for useful technical discussions.

REFERENCES

[1] Teo Babic, Harald Reiterer, and Michael Haller. 2018. Pocket6: A 6DoF Controller Based On A Simple Smartphone Application. In *SUI'18: 6th ACM Symposium on Spatial User Interaction*. 2–10.

[2] Yohan Baillot, L Davis, and J Rolland. 2001. A survey of tracking technology for virtual environments. *Fundamentals of wearable computers and augumented reality* (2001), 67.

[3] Ke-Yu Chen, Shwetak N Patel, and Sean Keller. 2016. Finexus: Tracking precise motions of multiple fingertips using magnetic sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1504–1514.

[4] Wei-Tung Chen and Ling-Jyh Chen. 2017. Pokeball: A 3D Positioning System Using Magnetism. In *Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), 2017 IEEE International Conference on*. IEEE, 719–726.

[5] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. 2018. CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 654.

[6] Houde Dai, Shuang Song, Xianping Zeng, Shijian su, Mingqiang Lin, and Max Q.-H. Meng. 2017. 6D Electromagnetic Tracking Approach Using Uniaxial Transmitting Coil and Tri-Axial Magneto-Resistive Sensor. *IEEE Sensors Journal* PP (12 2017), 1–1. https://doi.org/10.1109/JSEN.2017.2779560

[7] Michael Deering. 1992. High Resolution Virtual Reality. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 195–202. https://doi.org/10.1145/142920.134039

[8] Eric Foxlin, Michael Harrington, and George Pfeifer. 1998. Constellation: A Wide-range Wireless Motion-tracking System for Augmented Reality and Virtual Set Applications. In *SIGGRAPH*.

[9] Marvin P Fried, Jonathan Kleefield, Harsha Gopal, Edward Reardon, Bryan T Ho, and Frederick A Kuhn. 1997. Image-guided endoscopic surgery: results of accuracy and performance in a multicenter clinical study using an electromagnetic tracking system. *The Laryngoscope* 107, 5 (1997), 594–601.

[10] X. Ge, D. Lai, X. Wu, and Z. Fang. 2009. A novel non-model-based 6-DOF electromagnetic tracking method using non-iterative algorithm. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 5144–5117. https://doi.org/10.1109/IEMBS.2009.5332723

[11] Peter Henry, Michael Krainin, Evan Herbst, Xiaofeng Ren, and Dieter Fox. 2012. RGB-D mapping: Using Kinect-style depth cameras for dense 3D modeling of indoor environments. *The International Journal of Robotics Research* 31, 5 (2012), 647–663.

[12] Mohd Noor Islam and Andrew J Fleming. 2018. Resonance-Enhanced Coupling for Range Extension of Electromagnetic Tracking Systems. *IEEE Transactions on Magnetics* 54, 4 (2018), 1–9.

[13] Wooyoung Kim, Jihoon Song, and Frank C Park. 2018. Closed-form position and orientation estimation for a three-axis electromagnetic tracking system. *IEEE Transactions on Industrial Electronics* 65, 5 (2018), 4331–4337.

[14] Volodymyr V Kindratenko. 2000. A survey of electromagnetic position tracker calibration techniques. *Virtual Reality* 5, 3 (2000), 169–182.

[15] Volodymyr V Kindratenko and William R Sherman. 2005. Neural network-based calibration of electromagnetic tracking systems. *Virtual Reality* 9, 1 (2005), 70–78.

[16] Jack B Kuipers. 1980. SPASYN-an electromagnetic relative position and orientation tracking system. *IEEE Transactions on Instrumentation and Measurement* 29, 4 (1980), 462–466.

[17] Katja M Langen, Twyla R Willoughby, Sanford L Meeks, Anand Santhanam, Alexis Cunningham, Lisa Levine, and Patrick A Kupelian. 2008. Observations on real-time prostate gland motion using electromagnetic tracking. *International Journal of Radiation Oncology* Biology* Physics* 71, 4 (2008), 1084–1090.

[18] Magic Leap. [n.d.]. Magic Leap Fact Sheet. https://www.magicleap.com/static/magic-leap-fact-sheet.pdf. Accessed: 2019-04-05.

[19] Michael Meehan, Brent Insko, Mary Whitton, and Frederick P Brooks Jr. 2002. Physiological measures of presence in stressful virtual environments. *ACM Transactions on Graphics (TOG)* 21, 3 (2002), 645–652.

[20] CGM Meskers, HM Vermeulen, JH De Groot, FCT Van der Helm, and PM Rozing. 1998. 3D shoulder position measurements using a six-degree-of-freedom electromagnetic tracking device. *Clinical biomechanics* 13, 4 (1998), 280–292.

[21] Rajalakshmi Nandakumar, Vikram Iyer, and Shyamnath Gollakota. 2018. 3D Localization for Sub-Centimeter Sized Devices. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. ACM, 108–119.

[22] Rohit Pandey, Pavel Pidlypenskyi, Shuoran Yang, and Christine Kaeser-Chen. 2018. Egocentric 6-DoF Tracking of Small Handheld Objects. *CoRR* abs/1804.05870 (2018). arXiv:1804.05870 http://arxiv.org/abs/1804.05870

[23] Valter Pasku, Alessio De Angelis, Guido De Angelis, Darmindra D Arumugam, Marco Dionigi, Paolo Carbone, Antonio Moschitta, and David S Ricketts. 2017. Magnetic field-based positioning systems. *IEEE Communications Surveys & Tutorials* 19, 3 (2017), 2003–2017.

[24] D Perie, AJ Tate, PL Cheng, and GA Dumas. 2002. Evaluation and calibration of an electromagnetic tracking device for biomechanical analysis of lifting tasks. *Journal of biomechanics* 35, 2 (2002), 293–297.

[25] G. Pirkl, K. Stockinger, K. Kunze, and P. Lukowicz. 2008. Adapting magnetic resonant coupling based relative positioning technology for wearable activitiy recogniton. In *2008 12th IEEE International Symposium on Wearable Computers*. 47–54. https://doi.org/10.1109/ISWC.2008.4911584

[26] Anton Plotkin and Eugene Paperno. 2003. 3-D magnetic tracking of a single subminiature coil with a large 2-D array of uniaxial transmitters. *IEEE Transactions on Magnetics* 39, 5 (2003), 3295–3297.

[27] Polhemus. [n.d.]. Polhemus G4. https://polhemus.com/motion-tracking/all-trackers/g4. Accessed: 2018-12-11.

[28] L. Quéval. 2015. *BSmag toolbox user manual*. Technical Report. Dept. Elect. Eng., University of Applied Sciences Düsseldorf, Düsseldorf, Germany. http://www.lqueval.com Accessed: 2018-12-11.

[29] Frederick H Raab, Ernest B Blood, Terry O Steiner, and Herbert R Jones. 1979. Magnetic position and orientation tracking system. *IEEE Transactions on Aerospace and Electronic systems* 5 (1979), 709–718.

[30] D. Roetenberg, P. J. Slycke, and P. H. Veltink. 2007. Ambulatory Position and Orientation Tracking Fusing Magnetic and Inertial Sensing. *IEEE Transactions on Biomedical Engineering* 54, 5 (May 2007), 883–890. https://doi.org/10.1109/TBME.2006.889184

[31] Sheng Shen, He Wang, and Romit Roy Choudhury. 2016. I Am a Smartwatch and I Can Track My User's Arm. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '16)*. ACM, New York, NY, USA, 85–96. https://doi.org/10.1145/2906388.2906407

[32] Shuang Song, Chao Hu, Baopu Li, Xiaoxiao Li, and Max Q.-H. Meng. 2013. An Electromagnetic Localization and Orientation Method Based on Rotating Magnetic Dipole. *IEEE Transactions on Magnetics* 49 (03 2013), 1274–1277. https://doi.org/10.1109/TMAG.2012.2211375

[33] Shuang Song, Hongliang Ren, and Haoyong Yu. 2014. An Improved Magnetic Tracking Method Using Rotating Uniaxial Coil With Sparse Points and Closed Form Analytic Solution. *IEEE Sensors Journal* 14 (2014), 3585–3592.

[34] Olga Sorkine-Hornung and Michael Rabinovich. 2017. Least-squares rigid motion using svd. (2017).

[35] Evan Strasnick, Christian Holz, Eyal Ofek, Mike Sinclair, and Hrvoje Benko. 2018. Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 644.

[36] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. 2018. Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 86.

[37] Andre Zenner and Antonio Krüger. 2017. Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 23, 4 (2017), 1285–1294.