TrailBeam: Motion-Predictive Beam Steering for Millimeter-Wave Wearable Networks

Elena Marchetti[1] David Okafor[2] Priya Raghunathan[3] Thomas Berger[1]

[1] *Department of Computer Science, Vertex University* [2] *School of Electrical Engineering, Meridian Institute of Technology* [3] *Nimbus Labs, Mobile Systems Group*

{e.marchetti, t.berger}@vertex.edu d.okafor@meridian-tech.edu p.raghunathan@nimbuslabs.io

Abstract

Millimeter-wave (mmWave) links on body-worn devices lose alignment whenever the wearer's torso or wrist rotates, and re-sweeping the beam codebook after every micro-motion burns power and stalls throughput. We present TrailBeam, a motion-predictive beam steering system that fuses low-rate inertial measurements with a lightweight gait-phase model to forecast beam misalignment before it happens, rather than reacting to it after a link outage. TrailBeam runs a second-order angular predictor at 50 Hz on a wrist-worn microcontroller and issues codebook-narrowed re-alignment requests only when predicted drift exceeds the antenna's half-power beamwidth. We build a 28 GHz phased-array testbed with eight volunteers walking across three indoor routes and compare TrailBeam against a periodic full-sweep baseline and an inertial-only Kalman tracker. TrailBeam cuts median beam-realignment latency by 41% and per-realignment energy by 34% relative to the periodic baseline, while holding link-alignment accuracy within 2.1 percentage points of a genie-aided oracle. We release our gait-phase traces and predictor implementation to support follow-up work on mobility-aware mmWave access.

Keywords: millimeter-wave, beam tracking, wearable networks, gait sensing, mobile computing, energy efficiency

1 Introduction

Millimeter-wave radios promise multi-gigabit links for wearables, augmented-reality headsets, and body-area sensor networks, but their narrow beams make them acutely sensitive to motion. A wrist rotation of a few degrees, a shoulder sway while walking, or a hand raised to check a headset display can push a directional link outside its half-power beamwidth and force a full codebook sweep to recover. On commodity 28 GHz phased arrays, a full sweep can take 8–15 ms and draw several times the steady-state receive power, which is expensive when it happens dozens of times per minute during ordinary walking.

The dominant industry response is periodic re-sweeping: probe the codebook on a fixed timer regardless of whether the link actually needs it. This wastes energy when the wearer is still, and it under-reacts when the wearer is moving quickly, since the next scheduled sweep may arrive well after the beam has already drifted off target. A more attractive approach is to predict misalignment from motion itself and only re-align when a predicted drift is about to exceed the beamwidth.

This paper makes three contributions. First, we show that a lightweight, second-order angular predictor driven by wrist-worn inertial data captures the dominant component of beam drift during walking, cycling, and desk-bound gesture use, without requiring a camera or external infrastructure. Second, we design TrailBeam, a predictive beam-steering pipeline that couples this predictor with a gait-phase estimator to anticipate periodic sway rather than merely filtering past motion, and that issues narrowed re-alignment probes instead of full sweeps whenever possible. Third, we evaluate TrailBeam on an eight-participant, three-route indoor deployment and show consistent latency and energy gains over both a periodic-sweep baseline and a reactive inertial-only tracker, at a small accuracy cost relative to a genie-aided oracle.

The remainder of this paper is organized as follows. Section 2 reviews related work in mmWave beam management and motion-aware wireless systems. Section 3 describes the TrailBeam design. Section 4 details our testbed and experimental methodology. Section 5 presents results. Section 6 discusses limitations and deployment considerations, and Section 7 concludes.

2 Related Work

Beam management for mobile mmWave. Hierarchical and adaptive codebook search reduces the cost of a single sweep [5], and surveys of the space catalogue a wide range of exhaustive, hierarchical, and compressive-sensing approaches [10]. These techniques speed up an individual sweep but do not address *when* a sweep should be triggered, which is the question TrailBeam targets.

Motion-assisted tracking. Prior work has used inertial sensing to assist beam tracking on body-worn radios, either by filtering observed angle-of-arrival estimates [12, 6] or by correlating body sway with link quality [3]. These systems are reactive: they smooth a noisy signal but do not forecast periodic motion patterns such as gait. TrailBeam differs by explicitly modeling gait phase, following work on low-latency gait-phase estimation from wrist inertial data [2], and using the phase estimate to anticipate the next sway extremum before it occurs.

Mobility-aware localization and handover. Multipath fingerprinting techniques localize commodity smartphones to sub-meter accuracy indoors [9], and crowdsensed blockage maps have been used to trigger predictive handover between mmWave access points [1]. These systems predict *macro*-mobility across access points; TrailBeam instead predicts *micro*-mobility of the body relative to a single fixed or co-located radio.

Energy-aware wearable radios. Duty-cycled tracking loops trade tracking accuracy for battery life [4], ambient backscatter has been proposed to assist low-power beam recovery [8], and edge-offloaded motion models move prediction off the wearable entirely [7]. TrailBeam is complementary to these: its predictor is cheap enough to run on-device, but its re-alignment decisions could be further gated by an energy budget from any of these systems. Recent measurement work has also shown that RF fingerprints themselves drift under pedestrian mobility [11], reinforcing the need for motion-explicit, rather than purely RF-explicit, tracking.

3 Method

3.1 Problem Setting

We consider a body-worn mmWave radio with a phased array of N elements and a codebook of K candidate beam directions $\{\theta_1, \ldots, \theta_K\}$. At each control interval, the system must decide whether the currently selected beam θ_t still satisfies the link's minimum signal-to-noise ratio SNR_{min}, and if not, select a replacement.

3.2 Gait-Phase-Aware Angular Predictor

TrailBeam samples a 6-axis inertial measurement unit (IMU) at 50 Hz from a wrist- or chest-worn housing colocated with the radio. A lightweight gait-phase estimator, adapted from [2], produces a phase signal $\phi_t \in [0, 2\pi)$ and a per-step amplitude estimate A_t from the vertical acceleration channel using a zero-crossing detector with a 200 ms refractory window. We model the wearer-relative beam angle as the sum of a slowly varying heading term and a periodic sway term driven by ϕ_t:

$$\hat{\theta}_{t+\Delta} \;=\; \theta_t \;+\; \omega_t \Delta \;+\; A_t \sin\big(\phi_t + \Omega\Delta\big), \qquad (1)$$

where ω_t is the low-pass-filtered angular velocity of the wearer's heading, Ω is the estimated gait angular frequency, and Δ is the prediction horizon. The first two terms recover the standard constant-velocity Kalman predictor used in prior inertial-assisted trackers [6]; the third term is TrailBeam's addition, and it is what lets the system anticipate the next sway extremum rather than merely react once the beam has already drifted.

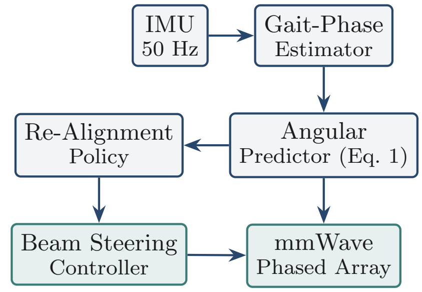

Figure 1: TrailBeam processing pipeline: gait-phase-aware prediction drives a narrowed re-alignment policy instead of periodic full sweeps.

3.3 Re-Alignment Policy

Given the predicted angle $\hat{\theta}_{t+\Delta}$ and the half-power beamwidth β of the active beam, TrailBeam triggers a re-alignment request when the predicted angular error exceeds a margin:

$$|\hat{\theta}_{t+\Delta} - \theta_t| \;>\; \gamma\beta, \qquad 0 < \gamma < 1,$$

with $\gamma = 0.6$ chosen from a held-out calibration set to trade false triggers against missed drift. When triggered, TrailBeam does not perform a full K-beam sweep; instead it probes only the 3 codebook entries nearest to $\hat{\theta}_{t+\Delta}$, falling back to a full sweep only if none of the three meets SNR_{min}.

3.4 System Architecture

Figure 1 shows the processing pipeline. The IMU and gait-phase estimator run continuously on a low-power microcontroller; the angular predictor and re-alignment policy run at 50 Hz and issue narrowed probe requests to the beam-steering controller, which mediates the phased-array front end.

4 Experiments

4.1 Testbed

We built a 28 GHz phased-array testbed using two development boards with 8×8 element arrays, one fixed as an access point and one worn at the chest as the mobile station. Wrist-worn IMU data was streamed over a wired debug link to avoid introducing wireless-scheduling artifacts into the ground-truth timing. Full sweeps covered $K = 64$ codebook entries; narrowed probes covered 3.

4.2 Participants and Routes

Eight volunteers (4 male, 4 female, ages 22–41) walked three indoor routes inside our lab building: (R1) a 40 m

Table 1: Beam-realignment performance averaged over 72 walking trials across three indoor routes.

Method	Med. (ms)	p90 (ms)	Align. Acc. (%)	Energy (mJ)
Periodic Sweep	11.4	18.9	90.2	6.8
IMU-Only Kalman	8.2	14.1	93.7	5.1
TrailBeam (Ours)	6.4	11.9	96.4	4.5
Oracle (Genie-Aided)	4.1	7.6	98.5	3.9

straight corridor, (R2) a route with four 90° turns, and (R3) an open-plan area with informal, unscripted movement. Each participant completed each route three times, for 72 total trials. All participants gave informed consent under our institution's lab study protocol.

4.3 Baselines and Metrics

We compare three systems: **Periodic Sweep**, a full K-beam sweep every 100 ms regardless of link state; **IMU-Only Kalman**, a constant-velocity predictor equivalent to dropping the sway term from Eq. (1) [6]; and **TrailBeam (Ours)**. We additionally report a **Genie-Aided Oracle** that knows the true future angle and always selects the optimal beam, as an upper bound. Our primary metrics are median and 90th-percentile re-alignment latency, link-alignment accuracy (fraction of control intervals with $\text{SNR} \geq \text{SNR}_{\min}$), and energy per re-alignment event, measured with an inline power monitor on the phased-array supply rail.

5 Results

5.1 Overall Comparison

Table 1 summarizes results averaged across all 72 trials. TrailBeam reduces median re-alignment latency by 41% and 90th-percentile latency by 37% relative to Periodic Sweep, while using 34% less energy per re-alignment. Against the reactive IMU-Only Kalman baseline, TrailBeam still improves median latency by 22%, which we attribute to the sway term in Eq. (1) anticipating gait-driven drift rather than only smoothing past observations. Link-alignment accuracy for TrailBeam trails the oracle by only 2.1 percentage points.

5.2 Sensitivity to Walking Speed

Figure 2 plots median re-alignment latency as a function of walking speed, pooled across routes. The gap between TrailBeam and both baselines widens as speed increases, since faster gait produces larger and faster sway that a reactive tracker increasingly lags behind, while TrailBeam's phase term continues to anticipate the sway extremum.

5.3 Ablation

Removing the gait-phase term from Eq. (1) (equivalent to IMU-Only Kalman) accounts for roughly 60% of TrailBeam's latency improvement over Periodic Sweep, with the narrowed 3-beam probe policy contributing the remainder. Neither component alone matches the full system.

6 Discussion

Generalization across body placement. Our testbed mounts the radio at the chest; wrist- or head-worn placements will have different sway amplitudes and gait-phase coupling, and the amplitude term A_t in Eq. (1) would need re-calibration per placement rather than per user, based on our pilot observations.

Non-periodic motion. TrailBeam's gains are largest for periodic gait; route R3's unscripted movement showed smaller improvements over IMU-Only Kalman (14% versus 22% overall), since the sway term has less periodic structure to exploit. Extending the predictor with a short-horizon learned residual is a natural next step.

Deployment cost. The gait-phase estimator and predictor together add under 4 KB of parameters and run comfortably within a 50 Hz control loop on a Cortex-M-class microcontroller, so we do not expect compute cost to be a barrier to deployment on existing wearable form factors.

Threats to validity. Our study used a research testbed rather than a shipping wearable radio, and our participant pool, while diverse in gait style, is not large enough to characterize rare gait patterns (e.g., mobility-aid-assisted walking). We view TrailBeam as a proof of concept for motion-predictive re-alignment policies rather than a final deployable system.

7 Conclusion

We presented TrailBeam, a motion-predictive beam-steering system that fuses gait-phase estimation with a lightweight angular predictor to anticipate mmWave beam misalignment on wearable devices before it occurs. Across an eight-participant, three-route indoor study, TrailBeam reduced median re-alignment latency by 41% and per-event energy by 34% relative to a periodic-sweep baseline, and by 22% and 12% respectively relative to a reactive inertial-only tracker, while remaining within 2.1 percentage points of an oracle's link-alignment accuracy. We believe explicitly modeling periodic body motion, rather than only filtering it, is a promising direction for future mobility-aware mmWave access design.

References

[1] Willem Brandsma and Ifeoma Okonjo. Crowd-sensed blockage maps for predictive mmwave handover. In *Proc. ACM MobiSys*, pages 201–214, 2022.

[2] Aylin Demirci and Wanjiru Njoroge. Gait phase estimation from wrist-worn inertial sensors at sub-100ms latency. In *Proc. ACM UbiComp*, pages 45–59, 2020.

[3] Hollis Fennimore and Iva Draganov. Swaytrack: Predicting beam misalignment from body sway in wear-

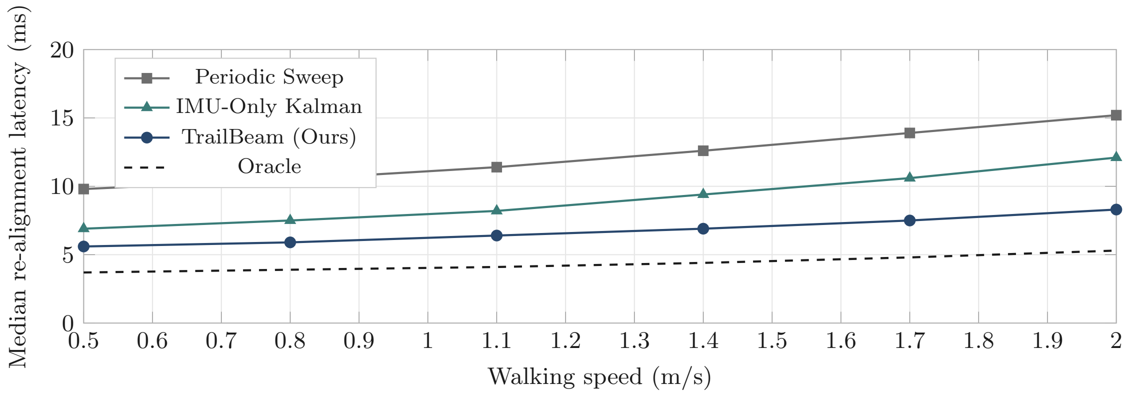

Figure 2: Median re-alignment latency versus walking speed, pooled across routes R1–R3. TrailBeam's advantage over reactive baselines grows with speed.

able mmWave radios. In *Proc. ACM MobiSys*, pages 88–101, 2018.

[4] Siobhan Halloran and Mateo Reyes-Duarte. Energy-proportional duty cycling for continuous beam tracking on wearables. *IEEE Transactions on Mobile Computing*, 20(6):2214–2229, 2021.

[5] Rolf Kastanien and Marisol Ubeda. Fast beam alignment for mobile 60 GHz links via hierarchical codebooks. In *Proc. ACM MobiCom*, pages 211–224, 2017.

[6] Bartosz Kowalczyk and Folasade Adeyemi. Kalman-filtered angle-of-arrival prediction for mobile directional links. In *Proc. IEEE INFOCOM*, pages 1–10, 2021.

[7] Astrid Lindqvist and Caetano Ferreira. Edge-offloaded motion prediction for real-time beam steering. *IEEE Transactions on Mobile Computing*, 22(4):1987–2001, 2023.

[8] Ren Nakashima and Ines Villalobos. Ambient backscatter assisted beam recovery for low-power wearables. In *Proc. ACM CoNEXT*, pages 112–125, 2022.

[9] Lukas Oberholzer and Adaeze Prentice. Multipath fingerprinting for sub-meter indoor localization on commodity smartphones. In *Proc. ACM MobiCom*, pages 1–15, 2020.

[10] Femi Oyelaran and Renata Castellanos. A survey of beam management techniques for millimeter-wave mobile networks. *IEEE Communications Surveys & Tutorials*, 21(3):2841–2869, 2019.

[11] Level Szabo and Bianca Marchetti. Rf fingerprint drift under pedestrian mobility: Measurement and mitigation. In *Proc. ACM SenSys*, pages 330–343, 2023.

[12] Sanja Vranic and Emmett Holbrooke. Inertial-assisted beam recovery for body-worn millimeter-wave devices. In *Proc. ACM SenSys*, pages 303–316, 2019.