mmRidge: Revealing Emergent Crowd Flow Structures with mmWave MIMO Radar

Anurag Pallaprolu, Winston Hurst and Yasamin Mostofi
{apallaprolu,winstonhurst,ymostofi}@ucsb.edu
University of California Santa Barbara, USA

Abstract

In this paper, we propose mmRidge, a novel framework that reveals emergent crowd flow structures using commodity mmWave MIMO radar. More specifically, mmRidge reveals flow boundaries that separate distinct motion behaviors in dense crowds. Our approach first converts radar point clouds into high-fidelity 2D flow fields that capture the collective motion of people. Building on concepts from dynamical systems theory, mmRidge computes a scalar Finite-Time Lyapunov Exponent (FTLE) field to uncover Lagrangian Coherent Structures (LCS) as "ridges" that demarcate boundaries between distinct crowd motion behaviors. Experimental validation with crowds of 18 or more individuals, across three real-world environments, demonstrates the effectiveness of mmRidge in detecting flow separators, highlighting the impact of advection, time horizons, and directionality on structure identification. Overall, the paper shows that mmWave radar can provide a low-power, privacy-preserving modality for identifying spatial crowd flow structures in dense and complex environments.

CCS Concepts

• **Human-centered computing** → **Ubiquitous and mobile computing**; **Visualization techniques**; • **Computer systems organization** → *Sensor networks.*

Keywords

mmWave Radar, Crowd Analytics, RF Sensing, Flow Fields, Finite Time Lyapunov Exponents, Crowd Semantics

ACM Reference Format:
Anurag Pallaprolu, Winston Hurst and Yasamin Mostofi. 2026. mmRidge: Revealing Emergent Crowd Flow Structures with mmWave MIMO Radar. In *The 27th International Workshop on Mobile Computing Systems and Applications (HotMobile '26), February 25–26, 2026, Atlanta, GA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3789514.3792044

1 Introduction

Crowd analytics provides important insights that drive decision making across a diverse set of domains. These insights are essential for a wide range of applications, such as real-time anomaly detection for public safety [7], design of urban spaces that minimize congestion and enable efficient evacuations [8], and optimization of commercial layouts based on customer behavior [3], to name a few. In these settings, frameworks that systematically identify regions of coherent flow can play critical roles in accurate, timely decision-making.

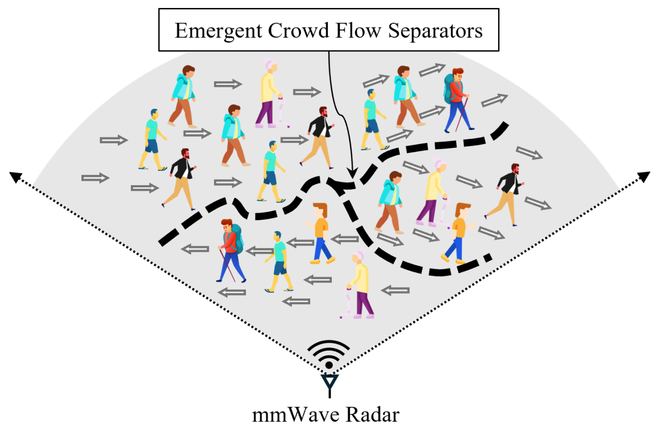

Figure 1: Illustrative schematic of coherent crowd flows showing opposing lanes and a diverging branch. mmRidge extracts the boundaries separating these regions.

Vision-based systems have emerged as the earliest mainstream approach for crowd analytics, driven by their high resolution and rapid development of machine learning techniques. However, mounting privacy concerns have prompted increased resistance towards large-scale video surveillance [4], shifting attention toward RF sensing, which naturally preserves privacy. Furthermore, the prospect of integrated sensing and communication (ISAC) in 6G systems [14], which would allow communication signals and networks to provide sensing capabilities, has spurred research into RF sensing at mmWave frequencies using commodity transceivers.

A growing body of work investigates the use of mmWave radar for crowd analytics, spanning tasks from occupancy estimation [12] to tracking a limited number of individuals simultaneously [1, 5]. In more recent work [11], we explored a flow field-based representation that summarizes aggregate crowd motion as 2D vector fields, enabling the analysis of collective movement patterns and the characterization of emergent structures (e.g., regions of attraction, dispersion, and rotation) from local vector field properties. [11], however, employs an Eulerian approach [2] that captures local motion patterns from instantaneous flow properties, and thus cannot reveal how these patterns evolve over time, or anticipate general, extended flow separations because of its pointwise nature.

In this paper, we present mmRidge, a novel mmWave radar-based framework for systematically identifying regions of coherent crowd movement. Unlike prior work, mmRidge adopts a new methodology based on a Lagrangian perspective, which can then trace the

ACM ISBN 979-8-4007-2471-8/2026/02
https://doi.org/10.1145/3789514.3792044

evolution of trajectories over time to delineate boundaries between regions exhibiting distinct behaviors (Fig. 1). This is accomplished by computing the Finite-Time Lyapunov Exponent (FTLE) field and extracting the associated Lagrangian Coherent Structures (LCS), a classical tool in the analysis of dynamical systems [15], as we shall see. We next summarize our key contributions.

Statement of Contributions:

a) We present mmRidge, a novel mmWave radar–based framework that identifies flow boundaries separating coherent regions of crowd movement (Fig. 1).

b) We introduce a Lagrangian crowd flow analysis framework (Fig. 2) that derives 2D flow fields from radar point clouds and computes Finite-Time Lyapunov Exponent (FTLE) fields to extract Lagrangian Coherent Structures (LCS).

c) Experimental validation across four real-world crowd scenarios with 18 or more individuals demonstrates reliable detection of flow separators, and further provides analysis of the impact of advection, time horizon, and directionality.

2 Radar Point Cloud Generation

We start with a primer on FMCW MIMO radar signal processing, followed by our methodology for point cloud generation. We consider a mmWave FMCW MIMO radar equipped with N_{TX} transmit antennas and N_{RX} receive antennas. Each transmitter periodically radiates a linear chirp whose frequency sweeps upward from the carrier f_0. Formally, the instantaneous transmit frequency is $f_{\mathrm{TX}}(t) = f_0 + St,\ 0 \le t < T_c$, where T_c is the chirp duration, B is the signal bandwidth, and the chirp slope is defined as $S = B/T_c$. The complex baseband waveform corresponding to the m^{th} chirp from the i^{th} transmitter can be written as $s_{\mathrm{TX}}^{(i)}(m,t) = \sqrt{P_{\mathrm{TX}}}\,\exp\big(j\{2\pi f_0 t + \pi S t^2\}\big)$, where P_{TX} is the transmit power.

Consider an object at distance $d_j[m]$ from the j^{th} RX. The signal received at RX j from TX i after reflection is $s_{\mathrm{RX}}^{(i,j)}(m,t) = \frac{s_{\mathrm{TX}}^{(i)}(m,t-\tau_j)}{d_j[m]^2}$, where $\tau_j = 2d_j[m]/c$ is the round-trip delay and c is the speed of light. In a TDM-MIMO configuration, TXs emit chirps sequentially, so each RX directly mixes its received signal with the corresponding TX chirp [13] to produce the intermediate frequency (IF) signal $s_{\mathrm{IF}}^{(i,j)}(m,t) = s_{\mathrm{RX}}^{(i,j)*}(m,t)\,s_{\mathrm{TX}}^{(i)}(m,t)$. To increase angular resolution, we adopt the standard virtual array construction in MIMO radar, which maps each TX-RX pair (i,j) to a virtual element (p,q), yielding an effective aperture of size $N_x N_y = N_{\mathrm{TX}} N_{\mathrm{RX}}$. Finally, let $s_{\mathrm{IF}}^{(p,q)}[m,n]$ denote the discretized IF signal, where n indexes ADC samples and $N_r = T_c/T_r$ is the number of range bins with an ADC sampling period T_r.

We next introduce a novel alternative to Constant False Alarm Rate (CFAR) based detection strategies that allows us to generate high quality point clouds for crowd flow scenarios. While traditional approaches employ a range-Doppler-angle FFT followed by CFAR thresholding [16], we instead adopt the methodology of [12], which begins with activity detection along the range dimension. More specifically, each IF stream $s_{\mathrm{IF}}^{(p,q)}[m,n]$ is first collapsed across the virtual array and passed through a range FFT, yielding per-chirp range profiles $R_{\mathrm{IF}}[m,r]$. Rather than applying a Doppler FFT along the chirp axis, temporal dynamics are captured by computing a Wrapped Phase Spectrum [12, Sec. 4.2] over sliding windows of chirps, from which the spectral bandwidth $\mathcal{H}[w,r]$ is extracted as a robust measure of local activity. This representation adapts naturally to varying target scales while suppressing background noise. The resulting bandwidth features are then binarized via a pretrained Gradient Boosted Machine (GBM) ensemble [12, Sec. 7], producing a Binary Trace Map $\mathcal{H}_{\mathrm{b,vis}}[w,r]$ that encodes which range bins are active. Range detections are thus defined as $\mathcal{D}_R[w] = \{r_i \mid \mathcal{H}_{\mathrm{b,vis}}[w,r_i] = 1\}$, and we refer the reader to [12] for a full derivation of this pipeline.

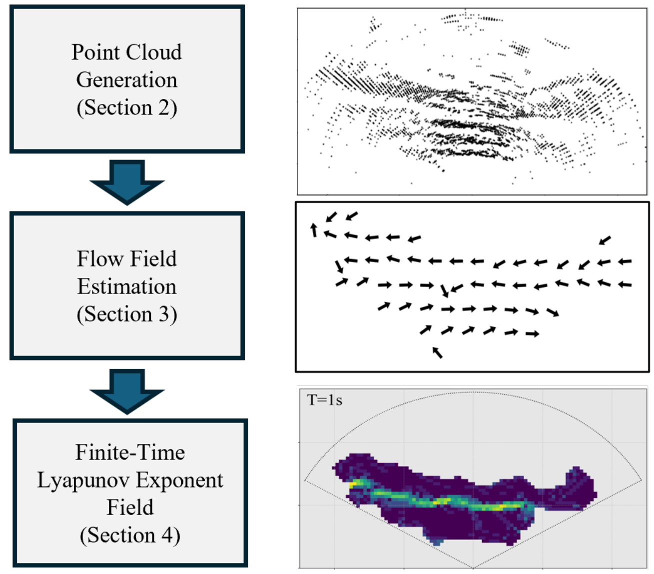

Figure 2: High-level view of the mmRidge pipeline.

Building on this, we extend the method with a virtual array to perform AoA estimation only at the detected ranges, as in [11]. The range–Doppler transform at array element (p,q) within window w of length N_c chirps is defined as

$$Q_{\mathrm{IF}}^{(p,q)}[w,v,r] = \mathcal{F}_{m,n}\Big\{s_{\mathrm{IF}}^{(p,q)}[m,n]\,\mathbb{1}_{w\le m<w+W}\Big\}, \tag{1}$$

and a 2D FFT over aperture indices (p,q) yields the azimuthal spectrum

$$\mathcal{A}[w,\theta,r] = \sum\nolimits_{\phi} \frac{1}{N_c} \sum\nolimits_{v=0}^{N_c-1} \mathcal{F}_{p,q}\{Q_{\mathrm{IF}}^{(p,q)}[w,v,r]\}. \tag{2}$$

In this formulation, Doppler measurements are marginalized by averaging across v, reflecting our focus on spatial localization rather than velocity estimation. Likewise, because commodity radar arrays exhibit limited elevation resolution, we collapse the elevation dimension after the angle FFT to obtain a cleaner 2D range–azimuth spectrum. Finally, the point cloud is assembled as

$$\mathcal{D}[w] = \{(r_i, \theta_i = \arg\max_{\theta} \mathcal{A}[w,\theta,r_i]) \mid r_i \in \mathcal{D}_R[w]\}. \tag{3}$$

Having established our method for generating radar point clouds, we next turn to constructing flow fields. A flow field provides a continuous representation of motion by interpolating from discrete radar detections to a spatial vector field that describes how the crowd is moving across the scene. This representation captures the collective dynamics needed for higher-level analysis and, crucially, provides the foundation for computing Finite Time Lyapunov Exponents (FTLE) and uncovering Lagrangian Coherent Structures (LCS) that partition the flow field into coherent sub-flows.

3 Crowd Flow Field Estimation

The sequence of point clouds $\mathcal{D}[w]$ offers a discrete snapshot of crowd dynamics, but for analyzing collective structure we require a continuous flow field $\mathbf{v}(\mathbf{x}) : \mathbb{R}^2 \to \mathbb{R}^2$. Our estimation process follows the pipeline detailed in [11, Sec. 3], and unfolds in three stages: (i) pairwise flow estimation, (ii) statistical denoising, and (iii) spatial pruning and smoothing. For completeness, we next briefly present each stage.

Pairwise Flow Estimation: Directly associating detections between successive point clouds is unreliable due to sparsity, missed detections, and false alarms. Instead, we adopt a Lucas–Kanade optical flow formulation [6], which infers local motion from intensity changes between frames. Each point $\mathbf{x}_{w,i} \in \mathcal{D}[w]$ is first embedded into a binary occupancy map $\mathcal{H}[w](\mathbf{x})$, and the displacement $\hat{\mathbf{v}}_{w,i}$ is then estimated by minimizing local intensity differences:

$$\hat{\mathbf{v}}_{w,i} = \arg\min_{\mathbf{h}} \sum_{\mathbf{y} \in \mathcal{N}_o(\mathbf{x}_{w,i})} \|\mathcal{H}[w+1](\mathbf{y}+\mathbf{h}) - \mathcal{H}[w](\mathbf{y})\|^2.$$

where $\mathcal{N}_o(\mathbf{x}_{\mathbf{w},\mathbf{i}})$ denotes a local neighborhood around $\mathbf{x}_{\mathbf{w},\mathbf{i}}$. This produces a set of candidate pairwise flows (PWFs) $\psi(\mathbf{x}) = \{\hat{\mathbf{v}}_{w,i} | \mathbf{x}_{w,i} = \mathbf{x}\}$, at each location $\mathbf{x}$ in the field of view. A natural first step is to average these candidates to obtain a time-averaged flow (TAF) estimate, $\mathbf{v}_{\text{TAF}}(\mathbf{x}) = \frac{1}{|\psi(\mathbf{x})|} \sum_{\hat{\mathbf{v}}_{w,i} \in \psi(\mathbf{x})} \hat{\mathbf{v}}_{w,i}$, with $\mathbf{v}_{\text{TAF}}(\mathbf{x}) = 0$ when $\psi(\mathbf{x})$ is empty. While this provides a straightforward flow estimate, it cannot by itself distinguish genuine crowd motion from spurious detections, motivating a statistical denoising step described next.

Kolmogorov–Smirnov Flow Denoising: While PWFs capture instantaneous crowd motion, many radar returns arise from persistent foliage noise or multipath, rather than genuine crowd movement. To filter these out, we analyze the angular distribution of flows at each location $\mathbf{x}$. Real motion produces consistent directional preference, whereas pure noise leads to angles that are effectively uniform on $[0, 2\pi)$. Formally, let $\Omega_{\mathbf{x}}(\alpha)$ denote the empirical CDF of PWF angles at location $\mathbf{x}$. We compare this against the uniform CDF $F_{\text{unif}}(\alpha) = \alpha/2\pi$ using the Kolmogorov-Smirnov statistic $D_{\mathbf{x}} = \sup_{\alpha \in [0,2\pi)} |\Omega_{\mathbf{x}}(\alpha) - F_{\text{unif}}(\alpha)|$, which quantifies the maximum deviation between the two. We then use the complementary CDF of the Kolmogorov distribution, κ, to compute the p-value, $p_{\mathbf{x}} = \kappa(D_{\mathbf{x}})$. If the corresponding p-value exceeds a threshold p_{th}, we cannot reject uniformity and treat the flows as unreliable. Otherwise, the flows are retained. We then take the Hadamard product of the resulting mask $M(\mathbf{x}) = \mathbb{1}_{p_{\mathbf{x}} \leq p_{\text{th}}}$ and the time-averaged flow field to produce the denoised flow field $\mathbf{v}_{\text{DFF}}(\mathbf{x}) = M(\mathbf{x}) \odot \mathbf{v}_{\text{TAF}}(\mathbf{x})$.

Flow Pruning and Smoothing: We next refine $\mathbf{v}_{\text{DFF}}(\mathbf{x})$ by exploiting the spatial coherence of natural crowd motion. Small, disjoint clusters below μ_S m^2 are removed via connected-component analysis, eliminating artifacts from multipath reflections or transient returns. Residual noise is suppressed with a two-stage median filter of window size μ_M m^2: the first pass interpolates across gaps, while the second enforces smoothness. This sequence mimics a morphological closing, yielding a cohesive flow field $\hat{\mathbf{v}}(\mathbf{x})$ when the above steps are applied on $\mathbf{v}_{\text{DFF}}(\mathbf{x})$.

As an additional step beyond the pipeline presented in [11, Sec. 3], we artificially pad the edges of the nonzero-flow region with vectors directed inward toward the main body of the flow. Conceptually, this enforces the idea that trajectories terminating at the boundary should remain within the observed region, consistent with the spatial distribution of radar returns. This boundary treatment ensures well-behaved advection dynamics and is especially important for reliable FTLE computation, which we explore next.

4 Finite Time Lyapunov Exponent Fields

In this section, we provide an introduction to Finite Time Lyapunov Exponents (FTLEs) [15] and explain our methodology for computing the FTLE field associated with the flow field $\hat{\mathbf{v}}(\mathbf{x})$. In the study of dynamical systems, an FTLE characterizes the extent to which a small perturbation in initial conditions affects a trajectory that is advected for a specific duration of time over an underlying flow field. Consider a particle with position $\mathbf{q}(t)$, initialized at a location $\mathbf{q}(t_0) = \mathbf{x}$ at time t_0, and let $\phi(\mathbf{x})_{t_0}^{t_0+T} = \mathbf{q}(t_0+T) = \mathbf{x} + \int_{t=t_0}^{t_0+T} \hat{\mathbf{v}}(\mathbf{q}(t))\, dt$ denote the position of the particle at time $t_0 + T$. Within the underlying flow field, an infinitesimal perturbation, $\delta\mathbf{x}$, of the initial location results in a deviation of

$$\|\delta\mathbf{x}(t_0+T)\| = \left\| \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}} \delta\mathbf{x} \right\|.$$

This value is maximized when $\delta\mathbf{x}$ is aligned with the eigenvector corresponding to the maximum eigenvalue of the finite-time Cauchy-Green deformation tensor [15], given by

$$\Delta = \left(\frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}} \right)^T \left(\frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}} \right).$$

The FTLE at $\mathbf{x}_0$ is then defined as $\sigma_{t_0}^T = (1/T) \ln \sqrt{\lambda_{\max}(\Delta)}$, where $\lambda_{\max}(\Delta)$ gives the largest eigenvalue of Δ. Importantly, the FTLE is a function of both the time horizon, T, as well as the initial position, $\mathbf{x}_0$.

Evaluation of FTLE Field: Calculating FTLEs requires evaluating the derivative $d\phi_{t_0}^{t_0+T}(\mathbf{x}_0)/d\mathbf{x}$. To achieve this, we initialize particles at each point in space associated with a non-zero flow and advect these particles forward using the flow field $\hat{\mathbf{v}}(\mathbf{x})$ for a fixed amount of time, T. We then evaluate the derivative by comparing the final locations of the particles with those of particles that began at nearby locations. Repeating this process across the domain produces a scalar field of FTLE values over the region of interest, where larger values indicate regions in which trajectories are highly sensitive to initial conditions. While standard (or forward-time) FTLE analysis emphasizes locations where trajectories diverge, regions of convergence can also be identified by instead advecting particles backward in time.

Lagrangian Coherent Structures: Analyzing the FTLE field uncovers Lagrangian Coherent Structures (LCS), which manifest as ridges delineating regions of distinct flow behavior. In the context of crowd dynamics, these ridges capture semantically meaningful boundaries, such as the division between self-organized pedestrian lanes, the merging of streams in a congested corridor, or the branching of flows at walkway intersections. **It is worth noting that such structures are not imposed by external annotation but**

Figure 3: Experimental setup: A mmWave MIMO radar transmits FMCW pulses, which reflect off of individuals moving in a crowd according to the white arrows, emulating the configurations C1-C4 in three different areas. The inset shows a front-facing view of the radar. See color PDF for best viewing.

arise directly from the underlying motion, making them robust indicators of collective behavior. The extraction of LCS is carried out using well-established image processing techniques applied to the scalar FTLE field, yielding a principled, data-driven method for identifying flow boundaries in dense crowds. In the next section, we demonstrate the performance of this analysis across several real-world scenarios and show that the resulting structures align closely with intuitive and observable crowd behaviors.

5 Experimental Validation

This section presents the results of several real-world experiments that validate our proposed approach. We first discuss our experimental setup, followed by crowd flow scenarios and system specifications. We then provide several examples of FTLE fields (and their corresponding LCS) extracted from flow fields, using the pipeline discussed in Sections 3 and 4.

Experimental Setup: We validate our pipeline through several real-world experiments with crowds of 18 or more individuals. As shown in Fig. 3, we utilize a TI AWR2243BOOST mmWave MIMO radar board, which operates at a base frequency of $f_0 = 76$ GHz and transmits FMCW pulses with a bandwidth of $B = 5$ GHz. The board includes $N_{\text{TX}} = 3$ transmitters and $N_{\text{RX}} = 4$ receivers, from which we construct a 1×8 virtual uniform linear array for azimuthal angle estimation. A TI DCA1000EVM FPGA streams the Channel State Information of the radar to a laptop via Ethernet.

Crowd Flow Configurations: We conduct our experiments in three distinct outdoor environments, with the high-level crowd flow configurations illustrated in Fig. 3. For each experiment, the radar was located at the midpoint of the area's lower periphery. Configurations C1, C2, and C4 involve groups of volunteers who are provided with general guidance regarding accessible areas and intended directions of movement, but without prescribed paths, lanes, or pacing. This allows participants to navigate freely while still maintaining the overall flow structure, thereby capturing natural variations in pedestrian behavior. In contrast, configuration C3 is based on naturally occurring crowd flows on campus during transitions between scheduled classes.

System Parameters: The pipeline described in Sections 3 and 4 operates with the following parameter settings. Pairwise flow field estimation is performed using OpenCV's implementation of the Lucas–Kanade algorithm [9]. For each point, the neighborhood $\mathcal{N}_o(\mathbf{x})$ is defined as a 2 m × 2 m window around $\mathbf{x}$, and we set the Kolmogorov-Smirnov threshold as $p_{\text{th}} = 0.15$ to denoise the pairwise flow fields. Flow components smaller than $\mu_S = 2$ m^2 in area are discarded, as such regions are unlikely to represent meaningful group motion, and the median filter is applied with a window size of $\mu_M = 0.25$ m^2, which roughly corresponds to the footprint of an individual. For trajectory advection in the FTLE computation, we integrate particle paths using an adaptive Runge–Kutta scheme (RK45). While there are several ways to implement LCS extraction, we use Otsu thresholding [10] on the FTLE field, followed by morphological dilation and opening, and we finally skeletonize the image using the classical approach developed in [17].

Experimental Results: To demonstrate our pipeline's ability to identify flow boundaries, we start by discussing an experiment with a crowd of 18 individuals adhering to configuration C1, as illustrated in Fig. 3. More specifically, we introduced a simple physical divider (a row of small training cones placed on the ground) in order to create a controlled reference for flow separation. Participants then freely walked on either side of the divider in opposite directions, producing two coherent counter-flows. We collected data for 35 seconds, and the corresponding FTLE field for this experiment is shown in Fig. 4 (a). The result is generated using forward advection (Sec. 4) with a time horizon of $T = 1$ s. As we can see, a pronounced LCS emerges with an average deviation of 19 cm from the reference line defined by the cones, confirming the pipeline's efficacy in accurately recovering flow boundaries.

To further test our approach, we next considered an unconstrained, real-world scenario: pedestrian traffic on a university campus during class transitions, which represents a setting characterized by dense and uncoordinated crowd motion (Fig. 3, C3). In this experiment, the radar was positioned alongside the sidewalk, and data were collected for 120 seconds. Despite the apparent complexity of the scene, the resulting FTLE field with the same forward advection horizon of $T = 1$ s reveals a clear dominant LCS that delineates the counter-flows (Fig. 4 (b)). It is worthwhile to note that this separation is not easily discernible by visual inspection of the point clouds over time, underscoring the method's ability to extract latent structure from complex crowd motion. Beyond the main separator, weaker LCSs reveal transient cross-traffic and nuanced pedestrian dynamics.

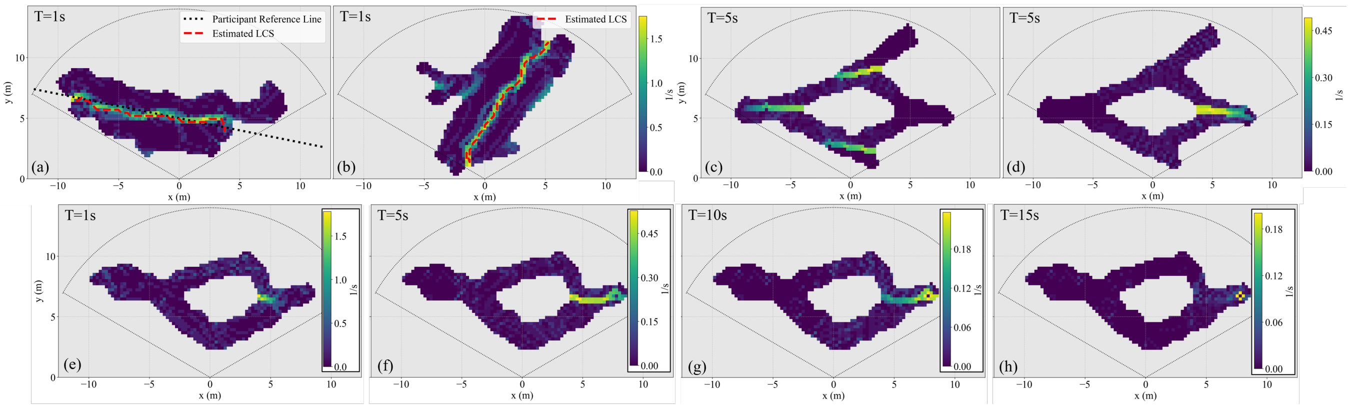

Figure 4: Results of FTLE fields from all experiments, with LCS appearing as green and yellow ridges highlighting dominant flow structures and regions of divergence or convergence. (a) Configuration C1 with participant reference line shown in black. (b) Spontaneous crowd flow in a campus environment (C3). (c) - (d) Configuration C4, showing forward-time (c) and backward-time (d) analysis, corresponding to flow separation and convergence, respectively. (e) - (h): Configuration C2 with different advection time horizons, T, demonstrating how increasing T reveals multi-scale flow behavior as trajectories that initially diverge later re-converge. See color PDF for best viewing.

We next illustrate our framework's ability to capture transitions within pedestrian flows through an experiment based on configuration C4 of Fig. 3. This setup emulates real-world scenarios such as intersections or bottlenecks, where lanes of foot traffic split or merge at distinct regions in space. The nodes of the flow graph were marked with training cones to indicate the intended branching and merging points, and participants were asked to follow the graph's edges while retaining freedom of movement, thereby preserving natural variations in pacing and spacing. We collected data for 38 seconds, and the resulting FTLE fields are evaluated with an advection time horizon of $T = 5$ s. More specifically, in the forward-time advection result of Fig. 4 (c), prominent LCS emerge at the branching location, effectively delineating the divergence of the two lanes. Conversely, when the FTLE is computed in reverse time (Fig. 4 (d)), the resulting LCS clearly highlights the location where the lanes reconverge, marking the merging behavior of the crowd.

Finally, we demonstrate the impact of the choice of the advection time horizon, T, on the FTLE field by presenting the results associated with experiment configuration C2 (see Fig. 3). In this configuration, participants enter from the right side of the area, split into two streams at a des- ignated point (one moving toward the radar and the other away from it) and subsequently merge at another point to exit as a single stream. We collected data for 38 seconds and we calculated the FTLE field with advection time horizons $T \in \{1\text{ s}, 5\text{ s}, 10\text{ s}, 15\text{ s}\}$, as shown in Fig. 4 (e) - (f). For short horizons ($T = 1$ s), a pronounced LCS appears on the right-hand side of the flow, reflecting the initial divergence of trajectories toward either the upper or lower path. As T increases, this LCS grows and extends, consistent with the continued separation of trajectories over intermediate timescales. However, for sufficiently long horizons ($T = 15$ s), the LCS fades as the two paths merge at the left-hand exit. This progression highlights how choosing T governs the multi-scale structure of crowd dynamics, balancing

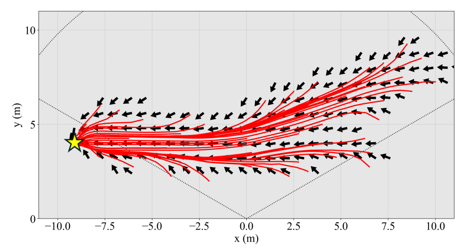

Figure 5: Identification of a sink using advection: trajectories (red) evolve along the flow field and accumulate at the exit (yellow star). See color PDF for best viewing.

transient separations and longer-term convergence, and motivating adaptive, context-aware horizon selection as future work.

Remark 1: The results presented thus far rely on advection on the flow field to estimate FTLEs. It is worth noting that advection by itself can also reveal flow structures. As an example, the real experiment of Fig. 5 shows advected trajectories correctly accumulating at the sink (exit), highlighting the generality of advection-based analysis for uncovering structures in crowd flows.

Remark 2: While [11] utilizes local curl and divergence to capture flow semantics, mmRidge complements such Eulerian methods by identifying (i) boundaries where trajectories diverge over extended regions rather than isolated points (Fig. 4(c)), (ii) multi-scale separators across advection horizons (Figs. 4(e)–(h)), and (iii) gradual separation zones not well captured by the local curl (Fig. 1, top-right).

6 Conclusion

In this paper, we have presented mmRidge, a framework that utilizes commodity mmWave MIMO radar to analyze emergent structures in crowd flows. By combining a novel crowd flow field estimation methodology with our proposed FTLE analysis, mmRidge is able to highlight persistent structures in complex and challenging crowd flow scenarios. Experimental results across diverse environments validate the robustness of our approach, highlighting the impact of advection time horizons and directionality on the extracted flow separators. These results demonstrate the potential of mmWave radar as a low-power, privacy-preserving sensing modality for capturing multi-scale structure in complex crowd behaviors.

7 Acknowledgments

This work is funded in part by ONR award N00014-23-1-2715, in part by NASA award 80NSSC25M7102, and in part by NSF CNS award 2226255.

References

[1] Valentín Barral, Tomás Domínguez-Bolaño, Carlos J Escudero, and José A García-Naya. 2024. An IoT system for smart building combining multiple mmWave FMCW radars applied to people counting. *arXiv preprint arXiv:2401.17949* (2024).

[2] George Keith Batchelor. 2000. *An Introduction to Fluid Dynamics*. Cambridge University Press.

[3] Antonio Greco, Alessia Saggese, and Bruno Vento. 2022. A robust and efficient overhead people counting system for retail applications. In *International Conference on Image Analysis and Processing*. Springer, 139–150.

[4] The Guardian. 2025. Is it time to ring the alarm on internet door cameras? *The Guardian* (2025). https://www.theguardian.com/world/2025/feb/09/is-it-time-to-ring-the-alarm-on-internet-door-cameras

[5] Hankai Liu, Xiulong Liu, Xin Xie, Xinyu Tong, and Keqiu Li. 2024. PmTrack: Enabling Personalized mmWave-based Human Tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 4 (2024), 1–30.

[6] Bruce D Lucas and Takeo Kanade. 1981. An iterative image registration technique with an application to stereo vision. In *IJCAI'81: 7th international joint conference on Artificial intelligence*, Vol. 2. 674–679.

[7] Rabia Nasir, Zakia Jalil, Muhammad Nasir, Tahani Alsubait, Maria Ashraf, and Sadia Saleem. 2025. An enhanced framework for real-time dense crowd abnormal behavior detection using YOLOv8. *Artificial Intelligence Review* 58, 7 (2025), 202.

[8] Hugo S Oliveira. 2025. Optimizing crowd evacuation: Evaluation of strategies for safety and efficiency. *Journal of Reliable Intelligent Environments* 11, 1 (2025), 2.

[9] OpenCV Project. 2024. OpenCV: Object Tracking. https://docs.opencv.org/3.4/dc/d6b/group__video__track.html Accessed: 2025-02-17.

[10] Nobuyuki Otsu. 1979. A Threshold Selection Method from Gray-Level Histograms. *IEEE Transactions on Systems, Man, and Cybernetics* 9, 1 (1979), 62–66. doi:10.1109/TSMC.1979.4310076

[11] Anurag Pallaprolu, Winston Hurst, and Yasamin Mostofi. 2025. mmFlux: Crowd Flow Analytics with Commodity mmWave MIMO Radar. *arXiv preprint arXiv:2507.07331* (2025). Accepted to appear in npj Wireless Technology.

[12] Anurag Pallaprolu, Phillip Peng, Shaan Sandhu, Winston Hurst, and Yasamin Mostofi. 2024. Crowd Analytics with a Single mmWave Radar. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 830–844.

[13] Sandeep Rao. 2017. Introduction to mmWave sensing: FMCW radars. *Texas Instruments (TI) mmWave Training Series* (2017), 1–11.

[14] Eliane Semaan, Erika Tejedor, Rajat Kumar Kochhar, Sverker Magnusson, and Stefan Parkvall. 2024. *6G spectrum - enabling the future mobile life beyond 2030*. Technical Report. Ericsson.

[15] Shawn C Shadden, Francois Lekien, and Jerrold E Marsden. 2005. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D: Nonlinear Phenomena* 212, 3-4 (2005), 271–304.

[16] Hongliu Yang and *et al.* 2025. From Spatial Domain to Temporal Domain: Unleashing the Capability of CFAR for mmWave Point Cloud Generation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 2 (2025), 1–29.

[17] T. Y. Zhang and C. Y. Suen. 1984. A fast parallel algorithm for thinning digital patterns. *Commun. ACM* 27, 3 (March 1984), 236–239. doi:10.1145/357994.358023