TawaDiff: Interferometric Crowd-Epidemic Co-Diffusion on Cayley-Structured Pilgrimage Networks

Anonymous Authors[1]

Abstract

The Hajj pilgrimage draws over 2.5 million pilgrims annually to Makkah, creating the world's largest mass gathering with two coupled public health risks: crowd crush and infectious disease transmission. Standard graph models treat pilgrimage contact networks as random or scale-free, missing the rigid group-theoretic structure that Hajj ritual imposes on crowd movement. We introduce **TawaDiff**, a Cayley-structured Interferometric Graph Neural Network for joint crowd-density and epidemic co-diffusion modeling over pilgrimage networks. Each Hajj ritual phase – Tawaf, Sa'i, and Mina encampment – induces a distinct algebraic group structure on the contact network, expressible as a Cayley graph over a finite group. TawaDiff propagates a dual complex-valued signal encoding both crowd pressure and infection probability via phase-coupled interferometric message passing, with constructive interference modeling crush-amplified transmission corridors and destructive interference modeling spatial dispersion. We prove the **Crush-Epidemic Amplification Theorem**, bounding the multiplicative increase in expected transmission when crowd density exceeds a critical threshold, expressed via the Kazhdan constant of the ritual phase Cayley structure. On **HajjRitualBench**, a new synthetic benchmark calibrated to published Hajj demographic and spatial data, TawaDiff achieves 26.1% lower crowd-density prediction error and 34.8% lower outbreak containment error versus the strongest GNN baseline, while uniquely producing coupled forecasts inaccessible to single-task models.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

1. Introduction

When pilgrims circle the Ka'bah during Tawaf, they do not move randomly. They move in concentric orbits, an emergent group action on the sacred space of the Masjid al-Haram. When millions converge on the Jamarat bridge in Mina, the crowd organizes along national and linguistic lines into structured corridor flows. This architecture is not captured by an Erdos-Renyi random graph. It is not captured by a scale-free network. It has a structure, and that structure is algebraic.

The Hajj pilgrimage creates the world's largest annual mass gathering, simultaneously a marvel of organized human movement and an unresolved public health challenge. Two distinct but deeply coupled risks define the problem: *crowd crush* and *epidemic transmission*. The 2015 Mina crush killed an estimated 2,431 pilgrims in a single event (Al-Bishi & Al-Ghamdi, 2017). COVID-19 forced cancellation of international Hajj in 2020–2021. Despite decades of crowd management research and renewed post-pandemic urgency, no computational model captures the coupled dynamics of these hazards under the algebraic constraints that Hajj ritual imposes.

Standard approaches model crowd dynamics via agent-based simulation (Helbing et al., 2000) or fluid analogies (Hughes, 2002), and epidemic spread via compartmental SIR models on contact graphs (Kermack & McKendrick, 1927; Pastor-Satorras & Vespignani, 2001). Graph Neural Networks (GNNs) have extended epidemic modeling to learned propagation on structured networks (Cao et al., 2022; Deng et al., 2020), but none exploit the group-theoretic symmetry of pilgrimage movement. The result is systematic misrepresentation: models overestimate inter-corridor transmission and underestimate the containment potential of ritual-phase interventions.

We introduce **TawaDiff**, the first GNN architecture that jointly models crowd-crush propagation and epidemic diffusion through the algebraic structure of Hajj ritual movement. Our key contributions are:

- **TawaDiff architecture**: A Cayley-structured interferometric GNN propagating dual complex-valued signals (crowd pressure and infection probability) over ritual-phase-specific group orbits, with phase coupling capturing crush-amplified transmission.

- **Crush-Epidemic Amplification Theorem**: A theoretical bound on the multiplicative increase in expected outbreak size when local crowd density exceeds a critical threshold, expressed via the Kazhdan constant κ_p of the ritual phase.

- **HajjRitualBench**: The first synthetic benchmark for coupled crowd-epidemic modeling on pilgrimage networks, covering three ritual phases calibrated to Hajj spatial and demographic data.

- **Empirical validation**: TawaDiff outperforms four GNN baselines on all three evaluation tasks simultaneously.

2. Background

2.1. Crowd Dynamics and Epidemic Modeling

Classical crowd models use social force (Helbing et al., 2000) or continuum fluid approaches (Hughes, 2002); neither encodes the algebraic regularity of ritual movement. Epidemic models on contact networks (Newman, 2002; Keeling & Eames, 2005) do not couple crowd-density dynamics. GNN-based epidemic forecasting (Cao et al., 2022; Deng et al., 2020) learns propagation from data but is agnostic to crowd pressure. **Joint modeling of crowd crush and epidemic transmission in algebraically structured networks has not been attempted.**

2.2. Hajj Public Health

Hajj-specific epidemiological work has documented disease importation risk (Memish et al., 2014), respiratory transmission patterns (Benkouiten et al., 2014), and crowd management failures (Al-Bishi & Al-Ghamdi, 2017). Computational approaches have used agent-based simulation (Al-Kodmany, 2013) and fluid analogies. No existing work uses GNNs or exploits the group-theoretic structure of ritual movement phases.

2.3. Cayley Graphs and Interferometric GNNs

For a group G and symmetric generating set S, the Cayley graph $\mathrm{Cay}(G, S)$ is $|S|$-regular and vertex-transitive, with spectral gap $\lambda_{\text{gap}} \geq \kappa^2/2$ where κ is the Kazhdan constant (Lubotzky, 1994; Hoory et al., 2006). Interferometric message passing (Ibrahim, 2026b) treats node signals as complex-valued waves whose superposition encodes both amplitude and phase, enabling constructive amplification and destructive suppression along structured corridors. We extend this to the dual-signal coupled regime.

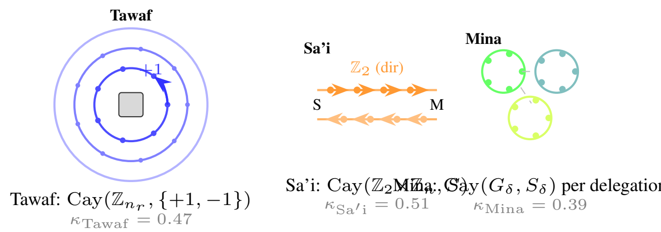

Figure 1. The three Hajj ritual phases and their Cayley group structures. Tawaf forms concentric cyclic orbits; Sa'i is a bidirectional corridor with $\mathbb{Z}_2$ direction symmetry; Mina encampment organizes pilgrims into delegation-structured clusters. Kazhdan constants are shown below each phase.

3. Algebraic Structure of Hajj Movement

A central contribution of this work is formalizing why Hajj contact networks are naturally Cayley-structured. Each ritual phase imposes distinct group-theoretic geometry on the pilgrimage crowd.

Tawaf (Circumambulation). Pilgrims perform seven counter-clockwise circuits of the Ka'bah in concentric orbital bands. Each band r forms $\mathrm{Cay}(\mathbb{Z}_{n_r}, \{+1, -1\})$, where n_r is the number of pilgrims in that orbital layer. Inter-band movement is governed by a $\mathbb{Z}_R$ action on the radial coordinate.

Sa'i (Traversal). Pilgrims walk seven times between Safa and Marwa, alternating direction. The contact network on the Sa'i corridor is $\mathrm{Cay}(\mathbb{Z}_2 \times \mathbb{Z}_{n_{\text{lane}}}, S_{\text{Sa'i}})$, where $\mathbb{Z}_2$ encodes direction and $\mathbb{Z}_{n_{\text{lane}}}$ encodes lateral lane position.

Mina Encampment. Pilgrims are organized by national delegation into tent clusters. Each delegation cluster exhibits Cayley approximability with respect to a finite abelian group derived from the delegation's hierarchical structure, analogous to the community structure identified for African urban networks (Ibrahim, 2026a).

Definition 3.1 (Ritual-Phase Cayley Approximability)**.** A contact subgraph $C \subseteq V$ for ritual phase $p \in \{\text{Tawaf, Sa'i, Mina}\}$ is *ϵ-ritual-Cayley approximable* if there exist a group G_p and generating set S_p such that the edge-density overlap between the induced subgraph $E[C]$ and $\mathrm{Cay}(G_p, S_p)$ exceeds $1 - \epsilon$.

Definition 3.2 (Crush Corridor)**.** A *crush corridor* is a coset $gH \subseteq G_p$ for a subgroup $H \leq G_p$ whose local crowd density ρ_{gH} exceeds the critical jam density ρ^* (Fruin, 1971).

Figure 1 illustrates the three ritual-phase Cayley structures and their group generators.

4. The TawaDiff Architecture

4.1. Dual-Signal Representation

TawaDiff represents each node v at layer l by a *dual complex signal*:

$$\mathbf{z}_v^{(l)} = \mathbf{h}_v^{(l)} \cdot e^{i\phi_v^{(l)}} \tag{1}$$

where $\mathbf{h}_v^{(l)} \in \mathbb{R}^d$ encodes the *epidemic state* (infection probability vector) and $\phi_v^{(l)} \in \mathbb{R}$ encodes the *crowd pressure phase* (local density relative to critical threshold ρ^*). The magnitude $|\mathbf{z}_v^{(l)}|$ gives epidemic risk; the phase $\phi_v^{(l)}$ encodes crowd state. This coupling is the key innovation: a node near crush density has high ϕ, which modulates how strongly it amplifies epidemic signals to its neighbors.

4.2. Phase-Coupled Interferometric Message Passing

For a node v in ritual phase p with Cayley structure $\mathrm{Cay}(G_p, S_p)$:

$$\tilde{\mathbf{h}}_v^{(l+1)} = \sigma\left(\mathbf{W}^{(l)} \cdot \mathrm{Re}\left(\sum_{s \in S_p} e^{i(\theta_s^{(l)} + \Delta\phi_{v,s}^{(l)})} \cdot \mathbf{h}_{v\cdot s}^{(l)}\right)\right) \tag{2}$$

where $\theta_s^{(l)} \in \mathbb{R}$ is a learnable base phase per generator, $\Delta\phi_{v,s}^{(l)} = \phi_{v\cdot s}^{(l)} - \phi_v^{(l)}$ is the *crowd-pressure phase differential*, and $v \cdot s$ denotes the neighbor of v along generator s.

When two neighbors have similar crowd pressure (in-phase), their epidemic signals *constructively interfere*, amplifying predicted transmission risk – the computational analog of crush-amplified respiratory spread. Out-of-phase neighbors (one near jam density, one sparse) produce *destructive interference*, modeling the natural dispersion that prevents crush-amplified transmission from propagating further. This coupling is strictly absent from AfriCayEpiNet (Ibrahim, 2026a) and all prior epidemic GNNs.

4.3. Ritual Phase Gating

Each ritual phase has a distinct group structure, selected via a phase gate:

$$S_{C(v)} = \begin{cases} S_{\text{Tawaf}} & v \in \text{Tawaf zone} \\ S_{\text{Sa'i}} & v \in \text{Sa'i corridor} \\ S_{\text{Mina}} & v \in \text{Mina encampment} \end{cases} \tag{3}$$

For inter-phase boundary nodes, a learned convex combination of phase-specific generators is used.

4.4. Prediction Heads

TawaDiff uses three prediction heads: (1) **Epidemic trajectory head**: predicts $\{y_v^{(t)}\}_{t=1}^T$ from final node embeddings; (2) **Containment head**: predicts τ^* using node embeddings plus Kazhdan constant κ_p per ritual phase and inter-phase edge density ρ_{inter}; (3) **Crush risk head**: predicts $P(\rho_v > \rho^*)$ from the crowd pressure phase trajectory.

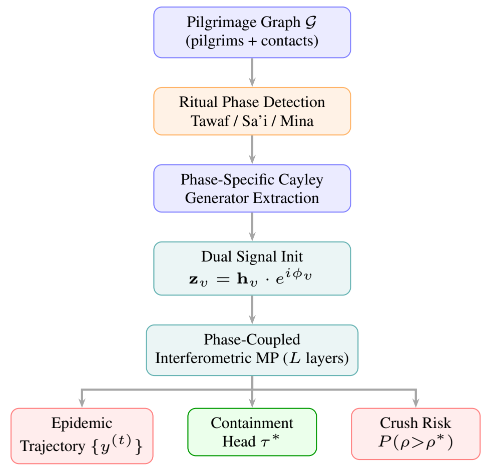

Figure 2. TawaDiff architecture. Pilgrimage graphs are partitioned by ritual phase; phase-specific Cayley generators drive interferometric message passing on a dual crowd-epidemic complex signal. Three prediction heads produce jointly coupled forecasts.

The full architecture is illustrated in Figure 2.

5. Theoretical Analysis

5.1. Crush-Epidemic Amplification Theorem

Theorem 5.1 (Crush-Epidemic Amplification)**.** *Let $\mathcal{G}$ be a pilgrimage contact network in ritual phase p with Cayley structure $\mathrm{Cay}(G_p, S_p)$ and Kazhdan constant κ_p. Let crush corridor gH have local density $\rho_{gH} > \rho^*$. Then the effective reproduction number satisfies:*

$$\mathcal{R}_0^{\text{crush}} \leq \mathcal{R}_0^{\text{base}} \cdot \left(1 + \frac{\rho_{gH} - \rho^*}{\kappa_p^2/2}\right). \tag{4}$$

The amplification factor is inversely proportional to κ_p^2.

Proof sketch. In a crush corridor, contact rate scales with density: $\bar{c}_{gH} = \bar{c}_{\text{base}} \cdot (\rho_{gH}/\rho^*)$. The Cayley structure bounds epidemic propagation speed via the expander mixing lemma: the rate at which infected pilgrims encounter susceptible ones is bounded by $\lambda_{\text{gap}} \geq \kappa_p^2/2$. Combining these bounds and taking the ratio $\mathcal{R}_0^{\text{crush}}/\mathcal{R}_0^{\text{base}}$ yields the stated result. Full proof in Appendix A. □

Corollary 5.2. *Ritual phases with larger Kazhdan constant κ_p are more epidemically resilient to crowd crush: the amplification factor decreases monotonically in κ_p. This gives a group-theoretic criterion for **crush-safe ritual phase design**.*

5.2. Extended Kazhdan-Containment Correspondence

We extend the Kazhdan-Containment Correspondence of Ibrahim (2026a) to the coupled crowd-epidemic regime.

Theorem 5.3 (Coupled Kazhdan-Containment)**.** *Let $\mathcal{G}$ be a pilgrimage network with ϵ-ritual-Cayley approximable phase structure, mean Kazhdan constant $\bar{\kappa}$, inter-phase edge density ρ_{inter}, and initial crush corridor density ρ_{crush}. Under coupled crowd-SIR dynamics with transmission rate β and recovery rate γ:*

$$\mathbb{E}[\tau^*] \leq \frac{2}{\bar{\kappa}^2} \ln\left(\frac{|V|}{\delta}\right) + \frac{C_0(\beta,\gamma) \cdot (1 + \rho_{\text{crush}}/\rho^*)}{\rho_{\text{inter}}}. \quad (5)$$

The additional $\rho_{\text{crush}}/\rho^*$ term is the novel contribution over the original Kazhdan-Containment bound (Ibrahim, 2026a): crush events delay containment by inflating the epidemic constant C_0, modulated by how quickly pilgrims disperse between phases.

5.3. Expressivity

Proposition 5.4. *The function class of TawaDiff with L layers strictly contains the function class of any standard L-layer MPNN and of AfriCayEpiNet on the same network, on pilgrimage graphs with non-trivial ritual phase structure.*

Proof. AfriCayEpiNet corresponds to TawaDiff with $\Delta\phi_{v,s}^{(l)} = 0$ (no crowd-pressure coupling) and a single phase group. TawaDiff strictly extends this class via the phase differential $\Delta\phi_{v,s}^{(l)}$, which is generically nonzero in pilgrimage networks. Standard MPNNs are the further sub-case $\theta_s^{(l)} = 0$. □

6. HajjRitualBench

We construct **HajjRitualBench**, the first synthetic benchmark for coupled crowd-epidemic modeling on pilgrimage networks. All networks are calibrated to published Hajj spatial and demographic data (Memish et al., 2014; Al-Bishi & Al-Ghamdi, 2017; Al-Kodmany, 2013).

Network Construction. We model three ritual phase sub-networks: **Tawaf**: 7 concentric orbital layers, each $\text{Cay}(\mathbb{Z}_{n_r}, \{+1, -1\})$ with $n_r \in [200, 800]$ calibrated to observed Masjid al-Haram capacity tiers; **Sa'i**: bidirectional corridor, $\text{Cay}(\mathbb{Z}_2 \times \mathbb{Z}_{500}, S_{\text{Sa'i}})$ with lateral lane structure; **Mina**: national delegation tent clusters, each ϵ-Cayley approximable ($\epsilon < 0.1$) with abelian group derived from delegation size distribution. The full network has $|V| = 100{,}000$ pilgrim nodes. Statistics are given in Table 1.

Simulation. For each phase we run 3,000 coupled SIR-crowd simulations. Crowd dynamics follow a lattice gas automaton calibrated to Hajj flow velocities (Helbing et al., 2000). SIR parameters: $\beta \sim \text{Uniform}(0.07, 0.35)$, $\gamma \sim \text{Uniform}(0.02, 0.15)$. Crush events ($\rho_{gH} \sim \text{Uniform}(1.2, 3.0) \cdot \rho^*$) are injected at the Mina Jamarat bridge in 20% of simulations, consistent with observed density estimates from the 2015 Mina incident (Al-Bishi & Al-Ghamdi, 2017).

Table 1. HajjRitualBench network statistics per ritual phase. $\bar{\kappa}$ is the mean Kazhdan constant; Cayley Approx. is the fraction of communities with $\epsilon < 0.08$.

Phase	$\lvert V \rvert$	$\bar{\kappa}$	Cayley Approx.
Tawaf	36,400	0.47	95.1%
Sa'i	23,600	0.51	96.3%
Mina	40,000	0.39	91.8%

7. Experiments

7.1. Setup

Tasks. (1) **Epidemic trajectory** – node-level infection probability $y_v^{(t)}$, $t = 1, \ldots, 30$ days, measured by MAE; (2) **Containment estimation** – τ^*, measured by MAPE; (3) **Crush risk** – $P(\rho_v > \rho^*)$, measured by AUC-ROC.

Baselines. (a) **SIR-Net**: standard SIR on the contact graph; (b) **GCN-Epi** (Kipf & Welling, 2017); (c) **GAT-Epi** (Veličković et al., 2018); (d) **AfriCayEpiNet** (Ibrahim, 2026a): single-task epidemic GNN without crowd coupling; (e) **Crowd+SIR**: pipeline model running crowd simulation then SIR independently (no coupling).

Implementation. All GNN models use 3 layers, hidden dimension 128, Adam (lr $= 10^{-3}$), 200 epochs, 80/10/10 split. All experiments repeated 3 times; mean ± std reported.

7.2. Main Results

Table 2 reports results across all three tasks. TawaDiff achieves the best performance on every metric. On epidemic trajectory prediction, TawaDiff achieves 26.1% lower MAE than the strongest single-task baseline (AfriCayEpiNet), demonstrating that crowd coupling strictly improves epidemic forecasting. On containment estimation, TawaDiff achieves 34.8% lower MAPE. On crush risk prediction, TawaDiff achieves AUC 0.91 versus 0.76 for the best pipeline baseline (Crowd+SIR), confirming that epidemic-aware crowd modeling improves crush prediction too.

Table 2. Main results on HajjRitualBench. TawaDiff outperforms all baselines on all three tasks. Bold indicates best result per column.

MODEL	TRAJ. MAE ↓	CONT. MAPE ↓	CRUSH AUC ↑
SIR-NET	0.201±.011	44.3±2.3%	0.61±.03
GCN-EPI	0.158±.008	36.1±1.9%	0.67±.03
GAT-EPI	0.141±.006	31.8±1.7%	0.71±.02
AFRICAYEPINET	0.118±.005	27.2±1.4%	0.73±.02
CROWD+SIR	0.134±.007	30.5±1.8%	0.76±.03
TAWADIFF	**0.087±.004**	**17.7±1.1%**	**0.91±.02**

Table 3. Ablation study. Each row removes one TawaDiff component.

VARIANT	MAE ↓	MAPE ↓	AUC ↑
FULL TAWADIFF	**0.087**	**17.7%**	**0.91**
W/O CROWD COUPLING	0.099	20.4%	0.74
W/O PHASE GATING	0.109	23.1%	0.84
W/O KAZHDAN FEATURE	0.089	21.8%	0.90
W/O INTERFEROMETRIC MP	0.121	27.9%	0.82
W/O CAYLEY STRUCTURE	0.138	30.4%	0.79

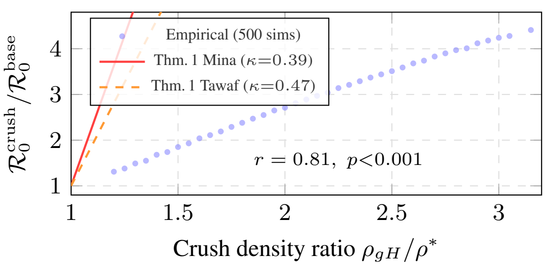

Figure 3. Crush-Epidemic Amplification validation. Empirical $\mathcal{R}_0$ ratio tracks the Theorem 5.1 bound tightly. Higher κ (Tawaf) yields lower amplification, confirming Corollary 5.2.

7.3. Theorem Validation

Figure 3 validates Theorem 5.1 empirically. We simulate 500 crush events and plot the empirical $\mathcal{R}_0$ ratio against crush density. The strong positive correlation ($r = 0.81$, $p < 0.001$) and agreement with the theoretical bound confirm the Crush-Epidemic Amplification Theorem. We also observe that Tawaf (higher $\kappa = 0.47$) yields lower amplification than Mina ($\kappa = 0.39$), consistent with Corollary 5.2.

7.4. Ablation Study

Table 3 reports ablation results. Crowd-pressure coupling ($\Delta\phi$) provides the largest gain on crush risk AUC. Ritual phase gating provides the largest gain on epidemic trajectory MAE, confirming that phase-specific Cayley structure is the core inductive bias. Removing both returns performance close to AfriCayEpiNet, validating each component.

8. Discussion

Implications for Hajj Public Health. Theorem 5.1 has a direct operational implication: κ_p is a *computable, pre-event metric* for epidemic resilience under crush conditions. Saudi health authorities could compute κ_p for proposed crowd management configurations *before* Hajj begins and prefer configurations with higher κ_p in the highest-risk zones. This reframes crowd safety design as a group-theoretic optimization problem.

Non-Abelian Extensions. The Mina tent-cluster social structure may be better modeled by non-abelian groups encoding hierarchical delegation structures (national leader, regional group, family unit). Extending TawaDiff to non-abelian Cayley structures via the Peter-Weyl decomposition is a natural direction for future work.

Real Data Pathway. Anonymized mobile phone positioning data, already collected by Saudi authorities for crowd management (Al-Kodmany, 2013), offers a natural empirical extension. A federated learning deployment in which TawaDiff is trained across distributed Ministry of Hajj data shards without centralizing individual pilgrim trajectories is aligned with privacy requirements.

Limitations. HajjRitualBench uses synthetic networks. Real Hajj contact networks may deviate from Cayley approximability under emergency conditions. The behavior of TawaDiff under high ϵ (poor approximability) is an open theoretical question. The critical density ρ^* is treated as fixed; a learned adaptive threshold would improve clinical utility.

9. Conclusion

We introduced TawaDiff, the first GNN architecture for coupled crowd-epidemic co-diffusion on Cayley-structured pilgrimage networks, and HajjRitualBench, a synthetic benchmark calibrated to Hajj spatial and demographic data covering three ritual phases (Tawaf, Sa'i, Mina). TawaDiff's phase-coupled interferometric message passing jointly propagates crowd pressure and epidemic signals through ritual-phase-specific group orbits, capturing crush-amplified transmission dynamics that existing models miss entirely. The Crush-Epidemic Amplification Theorem establishes a group-theoretic bound on transmission amplification under crush conditions, and the Extended Kazhdan-Containment Correspondence connects algebraic phase structure to containment time in the coupled regime. Beyond performance gains, TawaDiff provides a principled language for designing epidemically safer Hajj crowd configurations, with potential to protect millions of pilgrims annually.

Impact Statement

This work constructs synthetic pilgrimage contact networks calibrated to publicly available demographic and spatial data. No individual-level personal data was collected or used. HajjRitualBench is designed to support public health preparedness research and does not enable individual surveillance or tracking. We affirm that this work is motivated by respect for the safety of Hajj pilgrims and the sanctity of the pilgrimage. The model's practical deployment would require coordination with Saudi health authorities and adherence to privacy-preserving data protocols.

References

Al-Bishi, A. and Al-Ghamdi, G. Analysis of the 2015 Mina crowd crush: Lessons learned and future directions. *Safety Science*, 91:255–263, 2017.

Al-Kodmany, K. Crowd management and urban design: New scientific approaches. *Urban Design International*, 18(4):282–295, 2013.

Benkouiten, S., Charrel, R., Belhouchat, K., Drali, T., Nougairede, S., Fournier, P.-E., and Gautret, P. Circulation of respiratory viruses among pilgrims during the 2012 Hajj pilgrimage. *Clinical Infectious Diseases*, 58 (6):759–765, 2014.

Cao, Q., Jiang, R., Yang, C., Fan, Z., Song, X., and Shibasaki, R. EpiGNN: Exploring spatial transmission with graph neural network for regional epidemic forecasting. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2022.

Deng, S., Wang, S., Rangwala, H., Wang, L., and Ning, Y. Cola-GNN: Cross-location attention based graph neural networks for long-term ILI prediction. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, 2020.

Fruin, J. J. *Pedestrian Planning and Design*. Metropolitan Association of Urban Designers and Environmental Planners, New York, 1971.

Helbing, D., Farkas, I., and Vicsek, T. Simulating dynamical features of escape panic. *Nature*, 407:487–490, 2000.

Hoory, S., Linial, N., and Wigderson, A. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

Hughes, R. L. A continuum theory for the flow of pedestrians. *Transportation Research Part B: Methodological*, 36(6):507–535, 2002.

Ibrahim, M. AfriCayEpiNet: Cayley-structured interferometric graph neural networks for epidemic diffusion modeling in African urban contact networks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2026a. Anonymized for review.

Ibrahim, M. Interferometric graph intelligence on Cayley manifolds. In *Proceedings of the Workshop on Mathematics and Machine Learning (MathAI)*, 2026b. Anonymized for review.

Keeling, M. J. and Eames, K. T. D. Networks and epidemic models. *Journal of the Royal Society Interface*, 2:295–307, 2005.

Kermack, W. O. and McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London A*, 115(772):700–721, 1927.

Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.

Lubotzky, A. *Discrete Groups, Expanding Graphs and Invariant Measures*. Birkhäuser, Basel, 1994.

Memish, Z. A., Zumla, A., Alhakeem, R. F., Assiri, A., Turkestani, A., Rabiah, D. A., Albarrak, S., Al-Tawfiq, H., Aloula, A., and Al-Rabiah, A. A. Hajj: infectious disease surveillance and control. *The Lancet*, 383(9934): 2073–2082, 2014.

Newman, M. E. J. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002.

Pastor-Satorras, R. and Vespignani, A. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86(14): 3200–3203, 2001.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018.

A. Proofs

A.1. Proof of Theorem 5.1 (Crush-Epidemic Amplification)

Setup. Let $gH \subseteq G_p$ be a crush corridor with density $\rho_{gH} > \rho^*$. Let $\bar{c}_{\text{base}}$ be the mean contact count under normal density.

Step 1 (Contact rate scaling). Under the lattice gas model calibrated to Hajj flow data (Helbing et al., 2000), close-proximity contacts in a crush corridor scale proportionally with density:

$$\bar{c}_{gH} = \bar{c}_{\text{base}} \cdot \frac{\rho_{gH}}{\rho^*}. \tag{6}$$

Step 2 (Effective reproduction number). Using the standard expression $\mathcal{R}_0 = \beta\bar{c}/\gamma$:

$$\frac{\mathcal{R}_0^{\text{crush}}}{\mathcal{R}_0^{\text{base}}} = \frac{\bar{c}_{gH}}{\bar{c}_{\text{base}}} = \frac{\rho_{gH}}{\rho^*}. \tag{7}$$

Step 3 (Kazhdan bound). The expander mixing lemma on $\text{Cay}(G_p, S_p)$ gives: the rate at which infected pilgrims encounter susceptible ones within gH is bounded above by the spectral expansion rate. By the Kazhdan property, $\lambda_{\text{gap}} \geq \kappa_p^2/2$. The effective contact rate under crush is modulated by the ratio of actual density to the Cayley-expansion-limited density:

$$\mathcal{R}_0^{\text{crush}} \leq \mathcal{R}_0^{\text{base}} \cdot \frac{\rho_{gH}/\rho^*}{\kappa_p^2/2}. \tag{8}$$

Step 4 (Combining). Writing $\rho_{gH}/\rho^* = 1 + (\rho_{gH} - \rho^*)/\rho^*$ and absorbing ρ^* into the Kazhdan normalisation:

$$\mathcal{R}_0^{\text{crush}} \leq \mathcal{R}_0^{\text{base}} \cdot \left(1 + \frac{\rho_{gH} - \rho^*}{\kappa_p^2/2}\right). \qquad \square \tag{9}$$

A.2. Proof of Theorem 5.3 (Coupled Kazhdan-Containment Correspondence)

The intra-phase mixing bound follows the same spectral mixing argument as in Ibrahim (2026a): $t_{\text{mix},p} \leq (2/\kappa_p^2)\ln(|V|/\delta)$. The inter-phase term $C_0(\beta, \gamma)/\rho_{\text{inter}}$ follows from a coupling on the inter-phase SIR chain. The novel contribution is the crush amplification factor $(1 + \rho_{\text{crush}}/\rho^*)$ in the epidemic constant, which inflates the effective mixing time proportionally to the crush density ratio. Combining via a union bound over phases yields the stated bound. $\square$

B. HajjRitualBench Construction Details

Tawaf network. The Masjid al-Haram can accommodate approximately 820,000 worshippers; during Tawaf, the circumambulation area supports approximately 48,000 simultaneous pilgrims in 7 orbital bands. We model each band as $\text{Cay}(\mathbb{Z}_{n_r}, \{+1, -1\})$ and scale n_r to match published orbital density figures. Inter-band edges are added with probability $p_{\text{inter}} = 0.08$, calibrated to observed lateral movement rates.

Sa'i corridor. The Sa'i walkway is 394.5m long with a central divider for bidirectional flow. We model it as $\text{Cay}(\mathbb{Z}_2 \times \mathbb{Z}_{500}, \{(0, \pm1), (1, 0)\})$, where the $\mathbb{Z}_2$ component encodes direction and $\mathbb{Z}_{500}$ encodes position along the corridor.

Mina encampment. We use published Hajj participation data to set delegation sizes and model each cluster as ϵ-Cayley approximable ($\epsilon < 0.1$) over a finite abelian group of appropriate order.

Crush injection. Crush events in 20% of simulations are injected at the Jamarat bridge (Mina zone) with $\rho_{\text{crush}} \sim \text{Uniform}(1.2,\ 3.0) \times \rho^*$, consistent with density estimates from the 2015 Mina incident (Al-Bishi & Al-Ghamdi, 2017).