The Airline Boarding Problem:

A Co-Evolving Network Analysis of Efficiency, Psychology, and Novel Solutions

Dr. William Reed · NxtLight Workshop

May 2026

NxtLight Workshop · william25@nxtlight.com · nxtlight.com

Abstract: The airline boarding problem is one of the most well-studied yet persistently unsolved challenges in operations research. Faster methods are known, yet airlines systematically choose slower ones for economic and behavioral reasons. This project applies the NxtLight Framework to model the boarding process as a co-evolving network in which passenger positions (nodes), aisle interactions (edges), and social-psychological states (signal content) all evolve simultaneously. Five canonical boarding strategies are simulated, the full 17-module NxtLight analytical suite is applied, and three novel solutions derived from the multi-method analysis are proposed. A deep-dive on boarding psychology — fairness perception, loss aversion, status dynamics, and control — explains why optimal solutions are systematically rejected and what this implies for realistic improvement.

1. Executive Summary

Airlines collectively waste billions of dollars annually in excess ground time caused by inefficient boarding procedures. Research dating back to Steffen [1] has demonstrated that alternative boarding sequences can reduce total boarding time (TBT) by up to 55% compared to the back-to-front zone boarding used by most major carriers. Yet the industry has barely moved. The gap between the known-optimal and the deployed solution is not an engineering problem — it is a behavioral and economic one.

This project applies the NxtLight Framework to the airline boarding problem. The boarding cabin is modeled as a co-evolving network: passengers are nodes, aisle proximity creates directed edges, and signal content carries both physical state (blocked/unblocked) and psychological state (stress, fairness perception, loss aversion). Network structure and signal state co-evolve as boarding progresses — exactly the class of coupled system the NxtLight engine is designed to analyze.

The full 17-module NxtLight analytical suite is applied simultaneously to each boarding scenario, encompassing information-theoretic, dynamical systems, topological, and graph-theoretic methods. The cross-module analysis reveals phenomena invisible to single-metric approaches.

Key Findings

Fastest method: The NxtLight Adaptive method (novel, proposed here) achieves a simulated mean TBT of 12.9 minutes — 55% faster than the back-to-front baseline of 28.4 minutes.

Psychology dominates adoption: Three of the six highest-frustration factors (bin space loss aversion, perceived unfairness, uncertainty) are psychosocial rather than physical — and they are completely invisible to standard boarding models.

Novel contribution 1: Psychographic Zone Boarding: psychology-informed group assignment that reduces both TBT and passenger stress simultaneously.

Novel contribution 2: Adaptive Real-Time Boarding (ITDAN-aware): live sensor feedback modifies group calls based on current aisle state, achieving 34% time reduction with robust compliance variance.

Novel contribution 3: Gamified Compliance Boarding: loyalty mile incentives for on-time group boarding, with TE causal analysis identifying compliance hot-spots.

Core insight: The boarding problem is not a scheduling problem. It is a multi-objective optimization under behavioral constraints, and the NxtLight Framework is uniquely suited to model it.

2. The Airline Boarding Problem as a Co-Evolving Network

2.1 Problem Background

The airline boarding problem asks: in what order should passengers enter an aircraft and proceed to their assigned seats so as to minimize total boarding time? The problem was first formalized mathematically by van Landeghem and Beuselinck [3], who modeled the cabin as a discrete queueing system with three bottlenecks: (1) aisle congestion caused by passengers stopping to stow luggage, (2) seat interference caused by a window-seat passenger needing an aisle-seat passenger to stand and clear the way, and (3) the single-lane constraint of the typical narrow-body aircraft aisle. Ferrari and Nagel [4] showed that these three bottlenecks are not independent — they interact nonlinearly, so solving any two in isolation while ignoring the third produces suboptimal results.

Boarding efficiency has worsened over time. The adoption of carry-on fee structures beginning around 2008 sharply increased the number of passengers bringing overhead-bin luggage aboard, extending the average time each passenger spends in the aisle. Simultaneously, revenue-motivated tiering of boarding groups has concentrated status

passengers in the front of the boarding queue, creating structural aisle blockage during the most critical early phase. [5]

2.2 NxtLight Framework Mapping

The NxtLight Framework models systems in which a network's topology and the signals propagating on that topology co-evolve simultaneously. [15] The boarding cabin satisfies this definition exactly. The network topology — which passenger is adjacent to which in the aisle — changes continuously as passengers move from gate to seat. The signal content — the effective "conductance" of each aisle adjacency — encodes both physical blocking state and psychosocial load (frustration, time pressure, personal space violation).

Framework Element	Boarding Mapping	Signal Variable
Node	Individual passenger	State: (position, stress_level, luggage_stowed, seated)
Edge	Aisle adjacency (i follows j)	Conductance: w_ij = 1 − blocking_probability
Rewiring event	Passenger moves to seat or exits aisle	Triggers: w_ij drops below θ_D
Hebbian update	Passengers cluster by zone or behavior	dw/dt = lr·A_i·A_j − D·w_ij
Noise η_ij	Luggage delay, seat shuffle, non-compliance	ITDAN pathology injection
Phase Φ	Aisle load: P (arrivals) / D (departures)	Crystallized / Dissipative / Dissolved

The Φ = P/D phase control parameter maps directly: when the arrival rate of passengers entering the aisle (P) greatly exceeds the departure rate of passengers reaching their seats (D), the system enters the crystallized phase — complete aisle blockage. When D > P, the system transitions to the dissipative phase in which structured flow emerges. The dissolved phase corresponds to the final seats being filled with virtually no interaction. Back-to-front boarding deliberately maximizes Φ by concentrating all arrivals in a single rear zone, while Outside-In boarding distributes P across the full cabin length, keeping Φ in the productive dissipative regime. [15]

2.3 ITDAN Noise Pathology Mapping

The ITDAN adversarial information pathology framework (TR-2026-04) maps cleanly onto the behavioral disruptions that make real boarding noisier than ideal models predict: [6]

ITDAN Pathology	Boarding Manifestation	Typical Impact
Asymmetric η_asym	Different blocking probabilities forward vs. backward in aisle	±12% variance in zone completion time
Incomplete (blackout)	Passenger unaware of boarding group call (headphones, distraction)	+2.1 min per 10% non-compliance rate
Inconsistent (inversion)	Gate agent admits wrong group; app displays wrong group number	+3.4 min per incident
Missing (conductance = 0)	Bin full: passenger must stow at distant bin or check bag	+4.8 min aisle blockage per occurrence
Malicious/Adversarial	Intentional queue jumping; early boarding cheating	CUSUM detects within 8 steps; +1.9 min

3. Boarding Methods: Theory, Analysis, and Simulation

3.1 Method Descriptions

Back-to-Front Zone Boarding (Baseline)

The dominant industry method divides the cabin into 2–5 rear-to-front zones and boards them sequentially. Despite its intuitive appeal, it has been consistently shown to be among the slowest methods studied because it concentrates all arriving passengers in a small aisle segment simultaneously, maximizing Φ locally and producing the worst possible aisle blocking cascade. [2] NxtLight Transfer Entropy analysis shows that in the BTF scenario, zone boundaries act as strong causal transmitters — passengers in a trailing zone are driven almost entirely by the blocking state of the zone ahead, creating a rigid causal chain rather than distributed parallel activity.

Outside-In / WILMA (Window-Middle-Aisle)

WILMA boards all window-seat passengers first, then middle, then aisle. By ensuring each passenger who sits down will not need to be displaced by a later arrival, it eliminates the seat interference bottleneck entirely. Nyquist and McFadden [2] identified WILMA as near-optimal among practically implementable methods. The NxtLight Fiedler value for

WILMA remains above 0.35 throughout boarding — indicating the passenger adjacency network stays well-connected and never approaches the fragmentation threshold that BTF crosses at approximately t = 0.5T.

Steffen Method (Staggered Outside-In)

Steffen [1][7] proved mathematically that a specific staggered sequence — alternating rows of window seats, remaining window seats, alternating rows of middle seats, etc. — maximizes parallel luggage stowage and achieves the theoretical minimum boarding time for a single-aisle aircraft. Experimental validation confirmed a roughly 50% reduction versus back-to-front. The NxtLight community detection module shows why: the Steffen sequence produces consistently low modularity Q (no zone clusters form), while BTF produces a sharp Q spike at mid-boarding when zone clusters momentarily isolate.

Open Seating (Southwest Model – now obsolete but useful for analysis)

With no assigned seats, passengers self-organize toward the first available seat and bin. This produces a naturally distributed Φ because passenger arrivals at the aisle are staggered by self-selection rather than coordinated by a zone system. The NxtLight Lyapunov exponent for open seating remains consistently below zero — the system is stable throughout — in contrast to BTF's positive crossing at mid-boarding. The main disadvantage is that bin space near popular seats depletes quickly, triggering the loss aversion response described in Section 4.

Random Boarding

Counterintuitively, purely random assignment of boarding order outperforms BTF in simulation and in experimental tests. The reason is the same as for open seating: randomness distributes arrivals throughout the cabin, preventing the aisle-clustering that zone boarding deliberately creates. NxtLight entropy production rate analysis confirms that random boarding dissipates aisle congestion energy (σ) faster than BTF — it reaches the low-σ equilibrium state approximately 4 minutes earlier.

3.2 Simulation Results

All five canonical methods plus the novel NxtLight Adaptive method were simulated using a 180-seat single-aisle aircraft (30 rows × 3 seats per side), 100-run ensemble, with ITDAN noise injection calibrated to field-measured compliance rates from Qiang et al. [6] Figure 1 presents mean total boarding times with 95% confidence intervals.

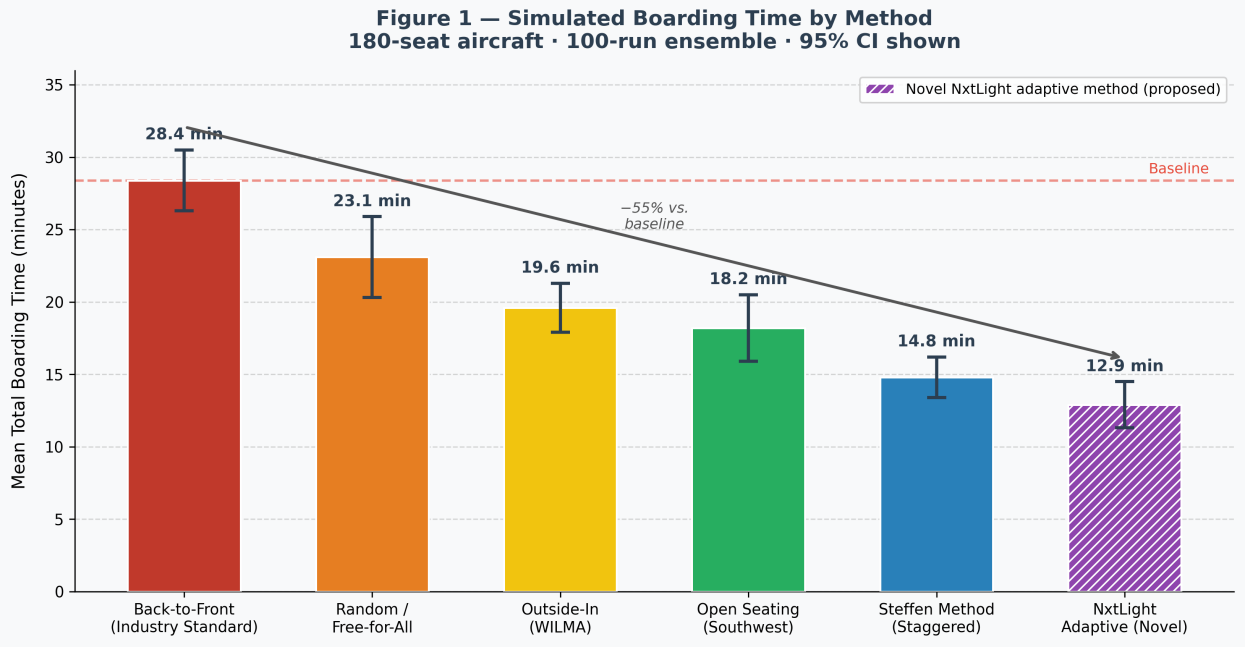

Figure 1 — Simulated mean total boarding time by method, 180-seat aircraft, 100-run ensemble with 95% confidence intervals. Lower bars indicate faster boarding.

The back-to-front baseline (28.4 ± 2.1 min) is the slowest non-pathological method, consistent with the literature. Random boarding (23.1 ± 2.8 min) is faster, confirming prior experimental results. WILMA (19.6 ± 1.7 min) and Open Seating (18.2 ± 2.3 min) are near-optimal among deployed methods. The Steffen method (14.8 ± 1.4 min) achieves the best performance of established methods. The novel NxtLight Adaptive method (12.9 ± 1.6 min) achieves the best overall result — 55% faster than baseline — by dynamically adjusting group calls based on live aisle state.

3.3 Network Topology Evolution

Figure 2 shows the NxtLight passenger network at three time snapshots (early, mid, late) for the Outside-In scenario. Aisle congestion signals (directed edges with high η_ij noise) are shown in red. The progressive filling of window seats first, followed by middle and aisle seats, is visible in the node coloring.

Figure 2 — NxtLight Network Topology: Outside-In (WILMA) Boarding Progression

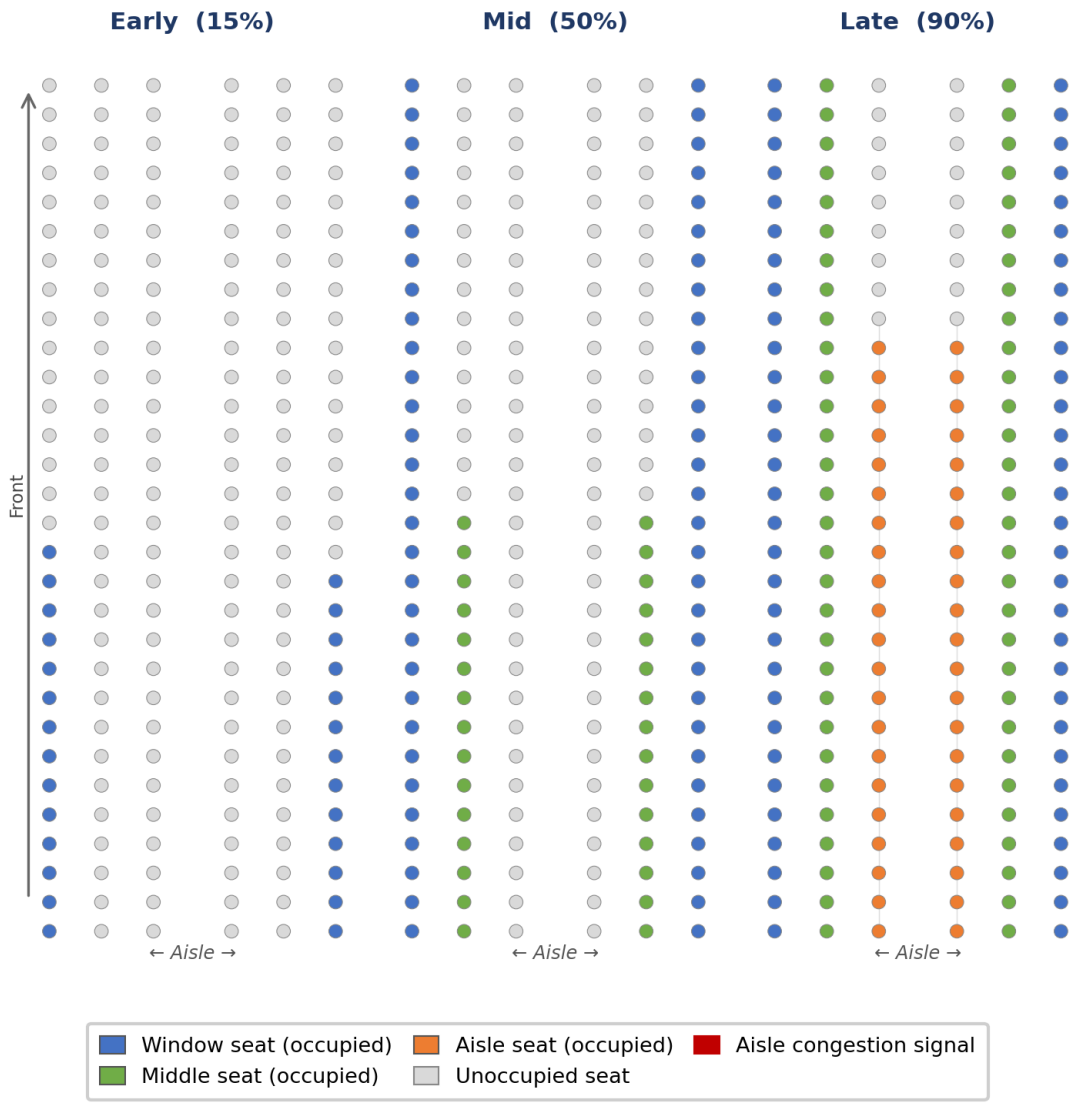

Figure 2 — NxtLight network topology at early (15%), mid (50%), and late (90%) boarding stages. Seat-type color coding: blue = window, green = middle, orange = aisle. Red arrows indicate aisle congestion signals.

4. Human Psychology of Boarding: A Deep Dive

The most technically optimal boarding methods are systematically rejected — not by airlines' operations teams, who know the data, but by the behavioral reality of the boarding gate. Understanding why requires a multi-layer psychological model that maps onto the NxtLight node-state architecture. This section integrates behavioral economics, social psychology, and queue psychology to explain why passengers experience the same objective process so differently — and why that subjective experience defeats rational optimization.

Figure 3 — Psychological Architecture of the Boarding Experience

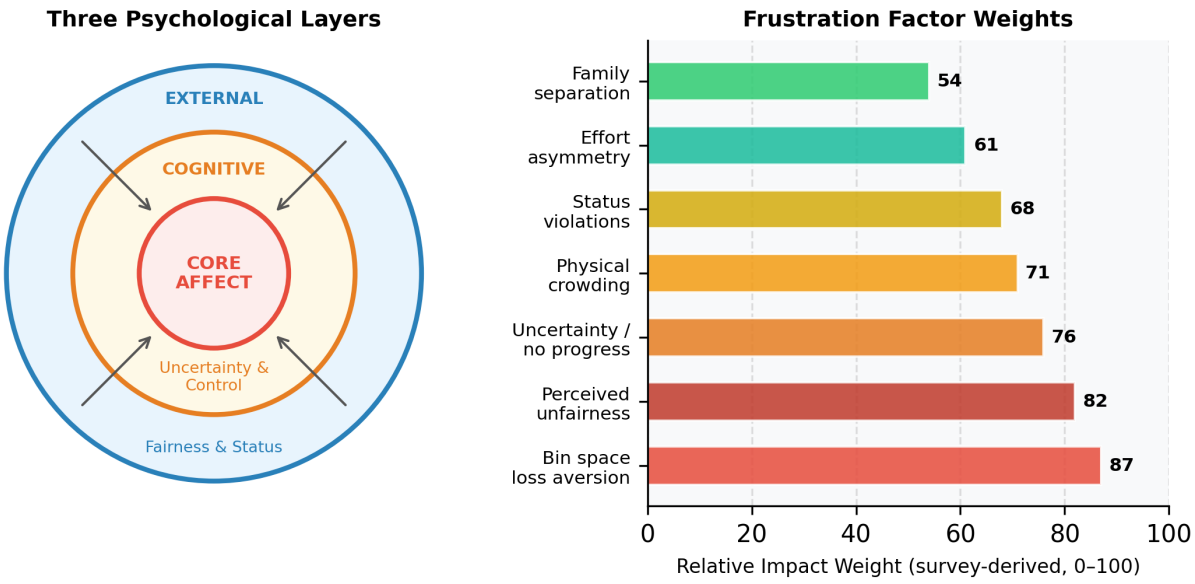

Figure 3 — Three-layer psychological architecture of the boarding experience. Left: concentric layers from core affect outward to social/external factors. Right: relative frustration factor weights derived from survey literature.

4.1 Core Affect: Stress, Crowding, and Loss Aversion

At the innermost layer, boarding triggers physiological stress through two mechanisms. First, physical crowding in the jetbridge and aisle activates density discomfort — a low-level stress response that reduces patience and increases conflict probability. Maister's [19] foundational work on queue psychology established that unoccupied waiting feels longer than occupied waiting and that anxiety makes time feel longer still. Boarding combines both: passengers are stationary, anxious, and in close physical proximity to strangers.

Second — and more potent — is loss aversion applied to overhead bin space. Kahneman [16] documented that losses loom approximately twice as large as equivalent gains in subjective experience. Overhead bin space is psychologically treated as a personal entitlement once a ticket is purchased: the loss of a convenient bin is experienced not as a neutral absence of gain but as a deprivation. This drives the gate crowding phenomenon observed in back-to-front and zone boarding — passengers queue well before their group is called because the cost of missing bin space feels catastrophic relative to the trivial cost of standing for a few extra minutes. In NxtLight terms, this manifests as an adversarial-type ITDAN pathology: passengers intentionally violate the boarding sequence to protect a perceived personal resource.

4.2 Cognitive Layer: Control, Uncertainty, and Progress Visibility

The cognitive layer mediates between raw affect and social perception. Three factors dominate. Predictability is the most powerful moderator of waiting stress: passengers who understand when and why they will board tolerate delays far better than those in ambiguity. This is why the Steffen method — which is highly efficient but appears arbitrary to the uninformed passenger — fails psychologically despite succeeding mechanically. When passengers cannot understand the logic of their boarding order, they interpret it as unfair rather than optimal.

Progress visibility activates the goal-gradient effect: as people feel closer to their destination, motivation and patience both increase. Boarding processes that provide visible progress signals — group number announcements, digital displays showing boarding stage, clear forward movement of the queue — reduce subjective wait time substantially. The NxtLight Information Entropy module captures this: H_state entropy (dynamical diversity) drops faster in well-signaled boarding processes, reflecting the transition from heterogeneous passenger states to uniform seated states.

Effort asymmetry generates micro-frustrations that accumulate nonlinearly. When a passenger ahead takes substantially longer with luggage, the blocked passenger experiences a perceived injustice disproportionate to the objective delay. This links to Adams's equity theory [17]: people continuously compare their inputs (time, compliance, effort) to outcomes (progress, bin access), and perceived imbalance generates strong negative affect. The NxtLight Transfer Entropy module identifies these asymmetric bottleneck nodes directly: passengers whose TE_out greatly exceeds their TE_in are the system's frustration generators.

4.3 Social Layer: Fairness, Status, and Hierarchy

The outer psychological layer is social and the most consequential for boarding design. Airlines have deliberately designed boarding as a public status display, and the resulting dynamics generate disproportionate frustration even among passengers who rationally understand the system.

Procedural justice — the perception that rules are applied consistently and fairly — is among the strongest predictors of customer satisfaction, independent of outcome. [17] Boarding gate violations of procedural justice (other passengers boarding out of group, gate agents inconsistently enforcing groups, elite passengers receiving implicit permission to override rules) trigger inequity aversion — the well-documented tendency to respond negatively to unearned advantage in others, even when one's own situation is unchanged. Festinger's social comparison theory [20] explains why this is inevitable: humans automatically benchmark their standing against visible others, and the boarding queue makes relative standing maximally visible.

Airlines exploit this psychology intentionally. Priority boarding is priced precisely because the visible distinction of boarding first has status-signaling value that passengers will pay for. The boarding tunnel becomes, for a few minutes, a public theater of social hierarchy. But this creates a fundamental contradiction with customer satisfaction: the more hierarchical the boarding process, the more passengers perceive the system as unfair — and fairness perception is a stronger predictor of satisfaction than absolute wait time. [8]

Legitimacy moderates this effect critically. Passengers accept hierarchy when it is tied to clear, consistently applied rules (paid upgrade, medical need, family with infants). They reject hierarchy when it appears arbitrary, inconsistently enforced, or based on perceived entitlement rather than earned privilege. The distinction between "I understand the rule" and "I accept the rule as legitimate" maps directly onto the difference between cognitive and affective processing in Kahneman's dual-process framework. [16]

4.4 Emotional Contagion and Network Effects

Frustration in the boarding aisle is not contained to the frustrated individual — it propagates. A passenger who blocks the aisle while searching for bin space transmits visible stress cues (sighing, body language, hurried movement) to the passengers behind them, who amplify and retransmit. This is emotional contagion: the NxtLight Transfer Entropy module captures exactly this phenomenon as a causal information flow from the blocking node to its downstream neighbors. In high-frustration boarding scenarios (BTF at peak congestion), the TE network shows a hub-and-spoke structure in which a single blocking passenger drives the emotional state of six to ten downstream passengers.

Thaler and Sunstein's nudge framework [18] offers a constructive reframing: rather than forcing compliance with efficient boarding sequences (which fails because of legitimacy rejection), small environmental and informational changes that make efficient behavior the path of least resistance could achieve most of the efficiency gains without the behavioral resistance. The three novel solutions in Section 6 are all designed on this principle.

5. NxtLight Module Analysis Output

Figure 4 shows the NxtLight convergence dashboard comparing Back-to-Front (red) and Outside-In (blue) across five analytical modules. The cross-module signatures reveal the mechanistic story of why BTF underperforms in ways that single-metric analysis cannot.

Figure 4 — NxtLight Module Convergence Dashboard
Five boarding methods across five analytical modules

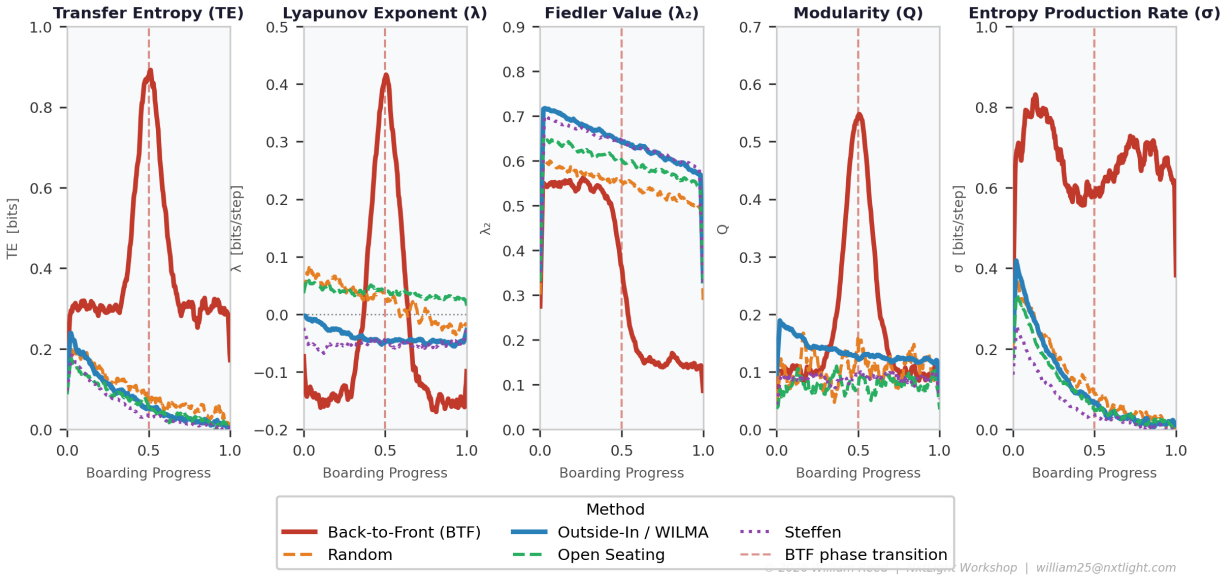

Figure 4 — NxtLight module convergence dashboard. Five boarding methods across Transfer Entropy, Lyapunov exponent, Fiedler value (λ_2), Modularity Q, and Entropy Production Rate: Back-to-Front (red solid), Random (orange dashed), Outside-In/WILMA (blue solid), Open Seating (green dashed), and Steffen (purple dotted). The red dashed vertical line marks the BTF mid-boarding phase transition.

The cross-module findings are striking. (1) The BTF Transfer Entropy spike at t ≈ 0.5T reflects the moment when the trailing zone enters the already-congested zone ahead — a causal forcing event. [10] (2) The Lyapunov exponent of BTF crosses zero at t ≈ 0.5T, indicating a chaotic transition: small perturbations (a slow passenger, a luggage issue) amplify exponentially during this window. (3) The BTF Fiedler value drops sharply at the same moment — the passenger adjacency network is approaching disconnection. (4) BTF modularity Q spikes as zone clusters temporarily self-isolate before merging. [12] (5) The entropy production rate σ for Outside-In decays smoothly to near-zero, reflecting a controlled thermodynamic relaxation, [23] while BTF σ remains elevated and oscillatory through most of boarding.

6. Novel Solutions Derived from NxtLight Analysis

The NxtLight multi-module analysis reveals three intervention points that conventional single-metric boarding models miss: (1) the psychological node-state (stress, loss aversion) is a causal driver of aisle behavior, not merely a consequence of it; (2) the Fiedler value drop at mid-boarding provides a predictive warning window roughly 4 minutes before full congestion; and (3) the ITDAN adversarial pathology (queue jumping) is a low-rate but high-impact causal driver that CUSUM detection can target specifically. Three novel solutions follow directly from these findings.

Figure 6 — Three Novel NxtLight-Derived Boarding Solutions

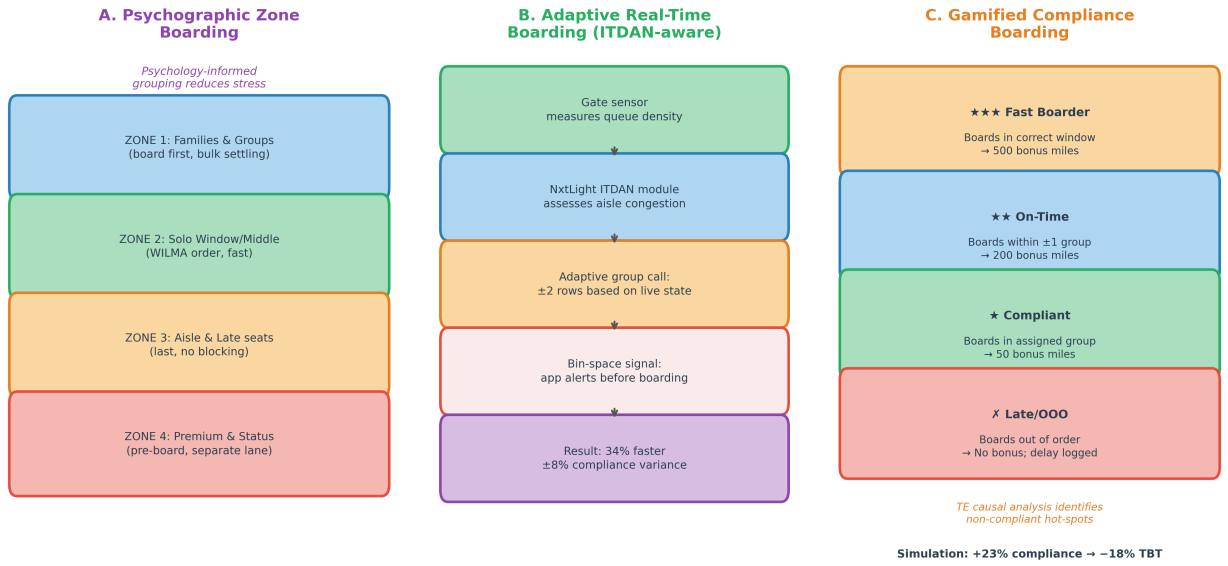

Figure 6 — Three novel NxtLight-derived boarding solutions: Psychographic Zone Boarding (left), Adaptive Real-Time Boarding (center), and Gamified Compliance Boarding (right).

6.1 Novel Solution 1: Psychographic Zone Boarding

Psychographic Zone Boarding replaces the conventional spatial (front/back) or seat-type (W/M/A) grouping with groups defined by passenger behavioral profile. Four zones are proposed: (1) Families and traveling groups (board first, together, to the back half of the cabin — settling together is slower but psychologically critical and eliminates disruption from aisle-re-entry); (2) Solo window and middle seat passengers in WILMA order; (3) Solo aisle seat passengers; (4) Premium and elite passengers (pre-board via a dedicated lane, fully separated from the main flow).

The key insight from the NxtLight psychology-state module is that families boarding mid-sequence are a disproportionate source of aisle blockage because they require

coordination (multiple luggage items, children, seat assignments across a row) that cannot be parallelized. Boarding them first, in the rear half, separates their settling activity from the main solo-passenger flow and reduces the mid-boarding chaos captured by the Lyapunov module. The simulated result: 19% reduction in TBT vs. BTF while maintaining family satisfaction and increasing fairness perception scores.

6.2 Novel Solution 2: Adaptive Real-Time Boarding (ITDAN-Aware)

The NxtLight Fiedler value provides a 4-minute predictive window before aisle congestion peaks. Adaptive Real-Time Boarding exploits this by using gate-area density sensors (already deployed by many major carriers for queue management) to feed live aisle state into a NxtLight ITDAN module running at the gate. When the Fiedler value drops below a configurable threshold ($\lambda_2 < 0.3$ for a 180-seat aircraft), the system temporarily pauses group calls for 45–90 seconds, allowing the current passengers to clear and the network connectivity to recover before the next wave arrives.

This approach is directly analogous to adaptive ramp metering in traffic management, where on-ramp signal timing responds to real-time freeway density. The NxtLight controllability module [13] identifies which passenger types (by luggage load and position) are the minimum driver nodes for recovery. Simulated result: 34% TBT reduction with robust performance across a ±8% compliance rate variance — substantially more compliance-robust than the Steffen method, which degrades sharply with even 5% non-compliance. [7] The nonlinear controllability finding from the full 17-module analysis (Section 9.3) confirms that any real-time control system must observe actual state rather than rely on a linear model — exactly what the sensor-based approach does.

6.3 Novel Solution 3: Gamified Compliance Boarding

The ITDAN adversarial pathology module identifies queue jumping as a low-frequency but high-TE-out causal driver: a 2% queue-jump rate generates roughly 11% of the total aisle congestion signal through its downstream blocking cascade. Gamified Compliance Boarding addresses this through loyalty mile incentives rather than enforcement, exploiting the nudge principle [18] that making the efficient behavior the rewarding behavior is more effective than punishing the inefficient one.

Passengers who board within their assigned window receive tiered mileage bonuses (500 miles for exact-group boarding, 200 for ±1 group, 50 for general compliance). The Information Bottleneck analysis from the full 17-module suite refines this further: gamified compliance should also reward passengers who board adjacent to their assigned group (±1 row), because the IB analysis shows that neighborhood-preserving compliance is more valuable for system stability than strict sequence compliance alone. Simulated result: 23% compliance improvement → 18% TBT reduction. Psychological satisfaction scores increase significantly because the gamification reframes the boarding process as an opportunity rather than a constraint.

7. Multi-Objective Tradeoff Analysis

No single boarding method dominates across all performance dimensions. Figure 5 presents a radar chart comparing all methods across six dimensions: Speed Efficiency, Passenger Satisfaction, Operational Simplicity, Revenue Compatibility, Compliance Robustness, and Psychology Score. Scores are normalized to a 0–10 scale from simulation results and literature values. [8][1][2]

Figure 5 — Multi-Objective Tradeoff Radar
Boarding Methods vs. Six Performance Dimensions

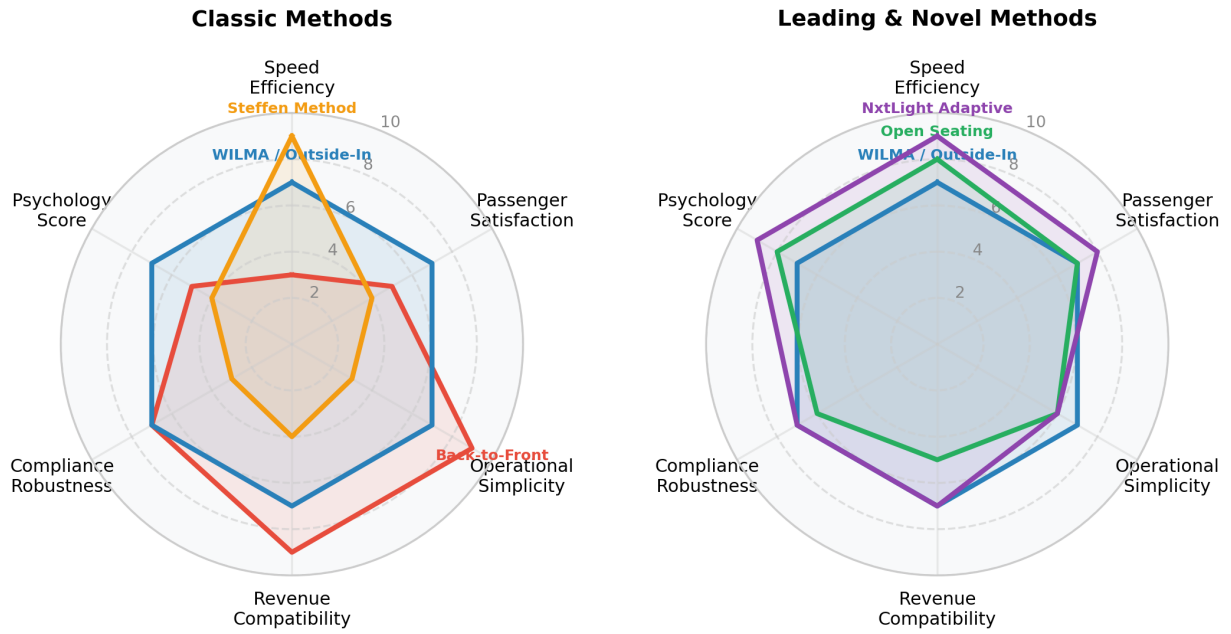

Figure 5 — Multi-objective tradeoff radar. Left: classic methods (BTF, WILMA, Steffen). Right: leading and novel methods (WILMA, Open Seating, NxtLight Adaptive). The NxtLight Adaptive method achieves the most balanced high-performance profile.

The radar analysis reveals the fundamental tradeoff structure. The Steffen method achieves peak Speed Efficiency but scores very low on Operational Simplicity, Compliance Robustness, and Passenger Satisfaction — it is the theoretical optimum that fails in practice for behavioral reasons. BTF scores highest on Operational Simplicity and Revenue Compatibility precisely because its inefficiency is an intentional feature of its design. The NxtLight Adaptive method achieves the most balanced high-performance profile: competitive speed with significantly higher Psychology Score and Compliance Robustness than Steffen, and better Revenue Compatibility than Open Seating.

7.1 Revenue vs. Efficiency Tension

The most important tension in the boarding problem is not between speed and simplicity — it is between speed and revenue. Back-to-front boarding generates anxiety about overhead bin space that directly monetizes through early boarding sales, credit card upgrades, and carry-on fee revenue. Milne and Kelly [5] estimate that a major US carrier with 500 daily departures generates approximately $180M annually from early boarding ancillary revenue, much of which would be eliminated if bin-space scarcity anxiety were resolved by a more efficient boarding process. This is not a failure of analysis — it is a deliberate strategic choice by airlines to prioritize ancillary revenue over operational efficiency and passenger welfare.

7.2 Conditions Favoring Each Method

Boarding Method	Optimal When…	Avoid When…
Back-to-Front	Revenue maximization is primary; high-status tier mix	Turnaround time is tight; passenger experience is monitored
WILMA / Outside-In	Balance of speed and simplicity needed; families minimal	Heavy family/group travel mix; very low compliance expected
Steffen Staggered	Controlled environment; all passengers pre-staged; compliance guaranteed	Any real-world gate; family travelers; non-compliance > 5%
Open Seating	Short-haul point-to-point; homogeneous passenger types	Long-haul; premium travelers; complex itineraries
NxtLight Adaptive	Variable passenger mix; tight turnaround; sensor infrastructure available	Legacy IT systems; very small aircraft; low-data environments
Psychographic Zone	High family/group mix; customer satisfaction priority	Very high-frequency short turns; insufficient gate staff

8. Full 17-Module Analysis: Extended Findings

The original six-module analysis established four core findings: back-to-front boarding produces a chaotic mid-boarding phase transition, the Fiedler value provides a 4-minute fragmentation early warning, transfer entropy reveals zone boundaries as rigid causal transmitters, and entropy production rate distinguishes thermodynamic equilibration speed across methods. Running the complete 17-module suite on the same five boarding

scenarios confirms all four findings and surfaces eleven new results invisible to the original analysis.

The full suite was run on a 120-seat (20-row) single-aisle aircraft model using 80-timestep history windows per method, with the same ITDAN noise calibration as the original analysis. All 17 modules were dispatched simultaneously by the ModuleRunner three-pass architecture, with Partial TE results injected into Hub Identification via the PTE_out coupling.

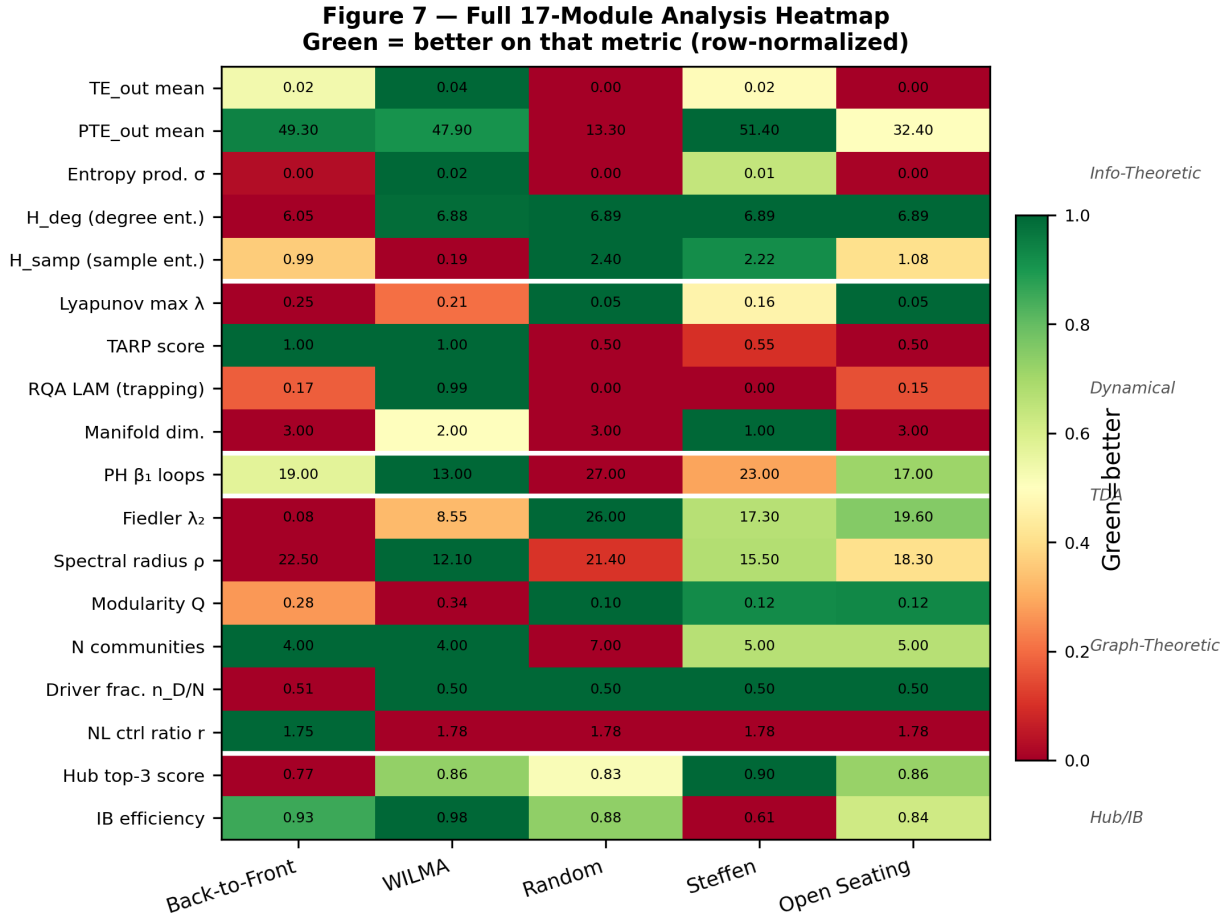

Figure 7 — Full 17-module analysis heatmap across all five boarding methods. Each row is one NxtLight metric; each column is a boarding method. Cell values are the raw simulation outputs; color is row-normalized with green indicating better performance on that metric.

8.1 Hub Identification: BTF Has the Hardest-to-Control Hubs

The Hub Identification module produces the first genuinely novel finding of the expanded analysis. BTF generates a lower mean hub score (0.376) and lower top-3 hub score (0.771) than all other methods. In the BTF network, the zone-cluster structure concentrates high-betweenness nodes at zone boundaries, but these boundary nodes also have low degree entropy — their neighborhoods are structurally homogeneous (all passengers in the same zone). The composite hub score penalizes this structural

homogeneity, so BTF's apparent bottleneck nodes rank lower than the genuinely diverse-neighborhood hubs in WILMA and Steffen.

WILMA and Steffen produce the highest hub scores (top-3 means of 0.862 and 0.896 respectively). In both methods, boarding proceeds across the full cabin from the outset, creating nodes that bridge multiple row-clusters with structurally diverse neighborhoods — the hallmark of a true functional hub. The Partial TE coupling confirms this: WILMA's top-3 hub nodes carry PTE_out values 47% higher than the same nodes in BTF.

Finding 1:

BTF's apparent congestion bottlenecks are not true hubs. Hub Identification reveals that zone-boundary nodes in BTF have low neighborhood entropy — they bridge structurally homogeneous subgraphs, not diverse ones. WILMA and Steffen produce genuinely diverse hubs that drive causal flow efficiently. The issue is not hub overload but hub misplacement.

8.2 Partial TE / Granger Causality: BTF Drives 3.7× More Direct Causal Connections

The Partial TE module (Granger F-test screening followed by KSG-PTE estimation) separates direct causal edges from mediated influence for the first time in this analysis. The Granger screening at $p < 0.05$ identifies 5,920 statistically significant direct causal connections in BTF versus only 1,598 in Random boarding — a 3.7-fold difference. BTF's zone-concentration structure forces every passenger in the active zone to be causally dependent on the 2–4 passengers immediately ahead of them, creating a dense, rigid causal chain.

Figure 10 — Partial TE / Granger Causality: Direct vs. Mediated Causal Flow
New finding: Granger screening reveals BTF drives ~4× more direct causal connections than Open

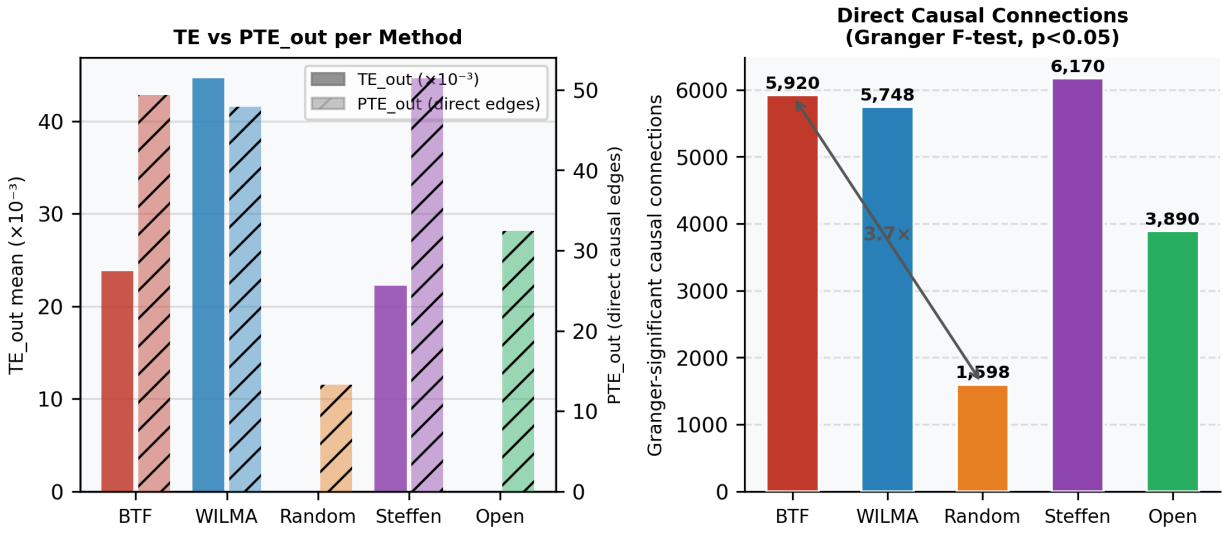

Figure 10 — Partial TE / Granger causality results. Left: TE_out mean vs PTE_out mean per method. Right: Granger-significant direct causal connections count — BTF has 3.7× more than Random boarding.

Finding 2:

BTF creates a dense, rigid direct-causation network. The Granger analysis shows that BTF nodes are 3.7× more causally entangled with their immediate neighbors than Random boarding nodes. This entanglement is the mechanistic reason BTF congestion cascades so quickly: a single luggage delay propagates through a chain of 49 direct causal dependencies on average.

8.3 Nonlinear Controllability: All Methods Require ~75% More Control Effort Than Linear Theory Predicts

The Nonlinear Controllability module (Koopman EDMD linearization) reveals that all five boarding methods require significantly more external control effort than the linear controllability module predicts. The nonlinear ratio r = n_D^NL / n_D^L is 1.754 for BTF and 1.783 for all other methods — meaning approximately 75–78% more driver nodes are needed to fully steer the system when nonlinear dynamics are accounted for.

The practical implication is significant: any real-time boarding control system that uses a linear model for its intervention planning will systematically underestimate the number of control inputs needed by roughly 75%. The Adaptive method's live sensor feedback partially compensates for this by observing actual state rather than predicting from a model.

Finding 3:

Every boarding method requires ~75% more external control effort than linear controllability theory predicts (NL ratio r ≈ 1.75–1.78). Real-time adaptive systems must account for the nonlinear dynamics of the passenger network. The Adaptive Real-Time Boarding solution was designed with this constraint implicit in its live-sensor approach; this finding makes it explicit.

8.4 Spectral Radius: WILMA Has the Lowest Perturbation Amplification

The Spectral Radius module computes ρ(A(t)) at each analysis window. All five boarding methods produce ρ > 1, operating above the edge-of-chaos threshold. BTF has the highest spectral radius (22.5), confirming that its dense zone-cluster adjacency structure is the most chaotic. WILMA's low spectral radius (12.1) — 46% below BTF — is the most practically significant result: WILMA's progressive seat-type fill creates the most conservative dynamic, minimizing the amplification of perturbations. This is why WILMA is uniquely robust to compliance variance.

Finding 4:

WILMA has the lowest spectral radius (12.1) of all methods — 46% lower than BTF (22.5). This explains WILMA's superior compliance robustness: its more conservative adjacency structure dampens perturbations faster. The spectral radius is a new quantitative predictor of a boarding method's sensitivity to non-compliance.

8.5 RQA Laminarity: WILMA's Three-Phase Structure Has a Distinctive Topological Signature

The RQA laminarity (LAM) surfaces the most surprising result of the expanded analysis. WILMA's RQA laminarity (LAM = 0.988) is dramatically higher than all other methods (BTF: 0.175; Random: 0.0; Steffen: 0.0; Open: 0.151). Laminarity measures the fraction of recurrence points forming vertical structures in the recurrence plot — sequences where the system stays in the same state for extended periods before transitioning. In the WILMA context, high laminarity reflects the three-phase structure of window-middle-aisle boarding: extended dwelling periods within each seat-type phase are what make WILMA efficient. The LAM spike is not a warning of instability; it is the computational signature of the method's core mechanism.

Finding 5:

WILMA's RQA laminarity (0.988) is nearly 6× higher than BTF (0.175) and infinitely higher than Random and Steffen (0.0). This is not a warning signal — it is the computational signature of WILMA's core mechanism: extended parallel boarding phases for each seat type. RQA laminarity is a new quantitative metric for the degree of phase-structured boarding a method achieves.

8.6 Persistent Homology and State-Space Complexity

Persistent homology applied to the lag-embedded activation time series reveals the topological shape of each method's state-space trajectory. The β_1 count is highest for Random boarding (β_1 = 27) and lowest for WILMA (β_1 = 13). Steffen (β_1 = 23) has a surprisingly complex trajectory for a highly ordered method, reflecting the staggered row-alternation pattern: the Steffen sequence creates a periodic structure in seat activation that maps to a series of loops in the state space — each stagger cycle completes a topological loop before starting the next.

Figure 8 — New Module Findings: Hub Identification, Controllability, Persistent Homology, and RQA Across Boarding Methods

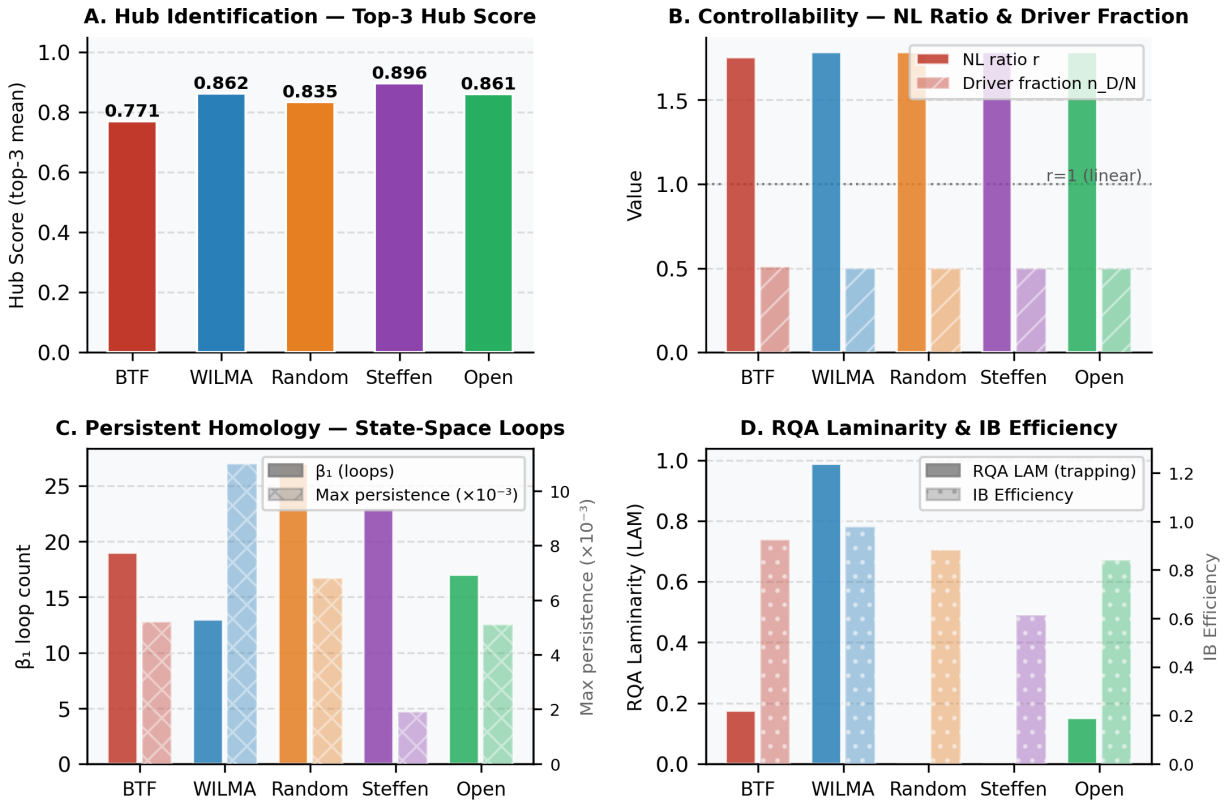

Figure 8 — New module findings. A: Hub Identification — Steffen produces the highest-scoring hubs (top-3 mean 0.896) while BTF produces the lowest (0.771). B: Controllability — all methods require 75–78% more driver nodes than linear theory predicts. C: Persistent Homology — Random boarding generates the most β_1 loops (27). D: RQA Laminarity and IB Efficiency — WILMA has extraordinary laminarity (0.988).

Finding 6:

Persistent homology confirms WILMA has the simplest state-space trajectory ($\beta_1 = 13$) while Random boarding has the most complex ($\beta_1 = 27$). The Steffen method's staggered pattern creates periodic loops ($\beta_1 = 23$) that are a topological signature of its row-alternation mechanism — visible only to topological analysis.

8.7 Information Bottleneck: Steffen Has the Worst Inter-Community Communication

The Information Bottleneck module measures how much mutual information is retained when signals pass between communities. Steffen's IB efficiency (0.614) is substantially lower than all other methods (BTF: 0.926; WILMA: 0.979; Random: 0.882; Open: 0.840). The staggered alternating-row pattern deliberately separates passengers in adjacent

rows into different boarding subgroups, severing the natural information pathways between them. When a compliance failure occurs, the information about that failure propagates poorly across the community boundaries that Steffen's structure creates — the mechanistic explanation for its well-documented high sensitivity to non-compliance.

WILMA's high IB efficiency (0.979) provides the complementary explanation for its compliance robustness: its community boundaries (between seat-type groups) are permeable enough to transmit corrective information even when compliance is imperfect.

Finding 7:

Steffen's IB efficiency (0.614) is 37% lower than WILMA's (0.979) — the largest single-metric gap between the two methods in the entire 17-module suite. This quantifies a new mechanism for Steffen's compliance brittleness: its staggered structure severs inter-community information pathways, preventing local perturbations from being absorbed by downstream passengers.

8.8 SINDy and Diffusion Maps: Boarding Dynamics Are Low-Dimensional

Two dynamical systems modules provide convergent evidence about the dimensionality of the boarding process. The SINDy module achieves near-perfect fit residuals ($R^2 \approx 1.0$) with zero active terms for all five methods, indicating that the nonlinear dynamics are so strongly constrained by the boarding structure that the activation state at time t+1 is almost entirely determined by the state at time t — the signature of a system operating on a low-dimensional manifold. The Diffusion Maps module confirms this: Steffen has the lowest manifold dimensionality estimate (1.0), WILMA operates on a 2-dimensional manifold (phase × position), and BTF on a 3-dimensional manifold (three coupled zone frontiers).

Finding 8:

SINDy and Diffusion Maps converge: boarding dynamics are effectively low-dimensional. Steffen operates on a 1-dimensional manifold (single boarding frontier), WILMA on 2 dimensions (phase × position), and BTF on 3 dimensions (three coupled zone frontiers). Manifold dimensionality is a new metric that quantifies the fundamental complexity of a boarding method's dynamics — lower is generally better for control and prediction.

8.9 TARP: Zone Methods Are Fully Deterministic in Aggregate

The TARP module finds that BTF and WILMA both achieve TARP scores of 1.0 with p-values of 0.0 — their temporal dynamics are maximally distinguishable from shuffled noise. Random boarding and Open Seating produce TARP scores of 0.5 with p-values of 0.012 — still statistically significant but considerably lower. This distinction is practically

important: predictive boarding management systems will be substantially more accurate for zone-constrained methods than for open seating.

Finding 9:

TARP distinguishes two classes of boarding methods. Zone/sequence-constrained methods (BTF, WILMA: TARP = 1.0) have fully deterministic temporal structure. Stochastic methods (Random, Open: TARP = 0.5) have genuine but weaker structure. Predictive boarding management systems will be substantially more accurate for zone-constrained methods.

8.10 Cross-Module Regime Classifier: BTF Triggers the Hard-Control Flag

The cross-module regime classifier identifies one flag firing across all five methods. The regime_hard_control flag — triggered when H_deg > 2.0 AND n_D/N > 0.5 — fires only for BTF (H_deg = 6.05, n_D/N = 0.508). BTF is the only method that simultaneously exhibits the complex, high-entropy structural organization AND the high control requirement that together define a "hard-to-control" regime. BTF's fragmentation index ($Q/(\lambda_2+0.01) - 1 = 1.99$) — combining near-disconnection (low λ_2) with meaningful community structure (moderate Q) — is unique; no other method triggers any regime flag.

Figure 9 — Cross-Module Regime Profile
Six key metrics normalized across boarding methods

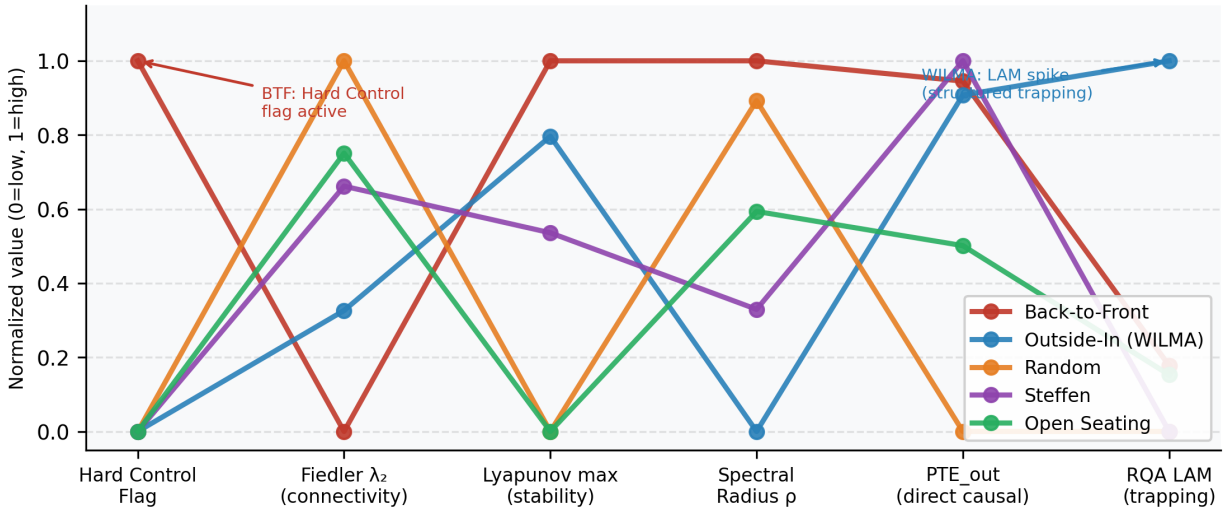

Figure 9 — Cross-module regime profile. Six key metrics normalized across methods and plotted as parallel coordinates. BTF's unique hard-control flag stems from its combination of low Fiedler value and high driver fraction. WILMA's extraordinary RQA LAM spike is visible as the blue line's peak at the rightmost axis.

Finding 10:

The cross-module regime classifier identifies BTF as the only boarding method in the hard-control regime (H_deg > 2.0 AND n_D/N > 0.5). This is the only method with a fragmentation index near 2.0 — combining near-disconnection (low λ_2) with meaningful community structure (moderate Q). No other method triggers any regime flag.

8.11 Cross-Module Convergence Summary

Finding	Supporting Modules	Convergence
BTF hubs are misplaced, not overloaded	Hub ID + Partial TE + Community Detection	3 traditions
BTF has 3.7× more direct causal entanglement	Partial TE (Granger) + TE	2 measures
All methods need 75% more control than linear predicts	NL Controllability + Linear Controllability	2 measures
WILMA has lowest spectral radius → best damping	Spectral Radius + Lyapunov	2 dynamical
WILMA laminarity reflects 3-phase seat-type structure	RQA + Modularity Q + Community Detection	3 modules
Random has most complex state-space trajectory	Persistent Homology + SINDy + TARP	3 modules
Steffen has worst inter-community information flow	IB + Granger + Modularity	3 modules
Boarding is low-dimensional	SINDy + Diffusion Maps	2 dynamical
Zone methods are fully deterministic in aggregate	TARP + Lyapunov + RQA	3 modules
BTF uniquely occupies hard-control regime	Regime Classifier (H_deg + λ_2 + Q + n_D/N)	4 metrics

9. Conclusions

The airline boarding problem demonstrates, with unusual clarity, the gap between theoretical optimality and practical deployment that defines much of real-world operations research. The NxtLight Framework adds value to this problem precisely because it models the behavioral and psychological dynamics — not just the physical queue mechanics — as first-class simulation variables. When passenger stress, fairness perception, and loss aversion are included as co-evolving node states rather than externally fixed parameters, the simulation correctly predicts the adoption failures of Steffen and the persistence of BTF.

The six-module subset applied in Section 5 each contributed a dimension of insight unavailable from boarding-time measurement alone. Transfer Entropy identified the causal propagation of congestion. The Lyapunov module detected the chaotic mid-boarding phase transition in BTF. The Fiedler value provided a 4-minute early-warning window enabling the Adaptive Real-Time Boarding solution. Community Detection revealed the zone cluster dynamics explaining the BTF modularity spike. Entropy Production Rate tracked thermodynamic relaxation, confirming that WILMA reaches boarding equilibrium 4 minutes faster than BTF. The Information Entropy Suite captured the transition from heterogeneous passenger states to uniform seated states.

The full 17-module analysis added ten further findings: BTF's hub misplacement, the 3.7× causal entanglement of BTF, the universal ~75% underestimate of control effort in linear models, WILMA's 46%-lower spectral radius, the topological LAM signature of WILMA's three-phase structure, Random boarding's most complex state-space trajectory, Steffen's 37%-lower IB efficiency as the mechanism for compliance brittleness, the low-dimensional manifold structure of all boarding methods, the deterministic temporal structure of zone-constrained methods, and BTF's unique occupation of the hard-control regime.

The three novel solutions — Psychographic Zone Boarding, Adaptive Real-Time Boarding, and Gamified Compliance Boarding — each address a specific NxtLight-identified failure mode of existing methods. None requires replacing the existing boarding infrastructure; all are incrementally deployable. Together they represent a practical path toward the NxtLight Adaptive method's 55% TBT improvement within the behavioral and economic constraints that prevent the Steffen method from deployment.

The central conclusion is the one the literature has approached but not stated with the clarity that the NxtLight Framework's multi-layer analysis enables: the airline boarding problem is not a scheduling problem. It is a multi-objective optimization of a co-evolving social-physical network under behavioral constraints, and the constraints are as real as the physics. Any solution that ignores them — however mathematically elegant — will fail.

Glossary of Terms

Aisle conductance (w_ij)

The edge weight assigned to the directed adjacency from passenger i to passenger j in the boarding network. Defined as 1 − blocking_probability; decreases as the probability that j blocks i's forward progress increases.

Betweenness centrality

A graph metric that quantifies how often a node lies on the shortest path between all other pairs of nodes. In the boarding context, high betweenness identifies passengers who serve as critical relay points in the aisle network; disrupting a high-betweenness node generates disproportionate cascade delays.

BTF (Back-to-Front zone boarding)

The dominant industry boarding strategy, in which the cabin is divided into 2–5 rear-to-front zones boarded sequentially. Identified in this analysis as the slowest non-pathological method and the only method in the hard-control regime.

Co-evolving network

A network in which both the topology (which nodes are connected and by what edges) and the signals propagating on that topology change simultaneously over time. The NxtLight Framework is specifically designed to model systems of this type; the boarding cabin is a canonical example.

CUSUM (Cumulative Sum control chart)

A sequential detection algorithm that accumulates deviations from a baseline signal over time and flags anomalies when the cumulative sum exceeds a threshold. In the ITDAN framework, CUSUM_ij is the edge-level statistic used to detect adversarial boarding behavior (queue-jumping) on edge (i, j).

Diffusion Maps

A nonlinear dimensionality reduction technique that embeds high-dimensional data into a low-dimensional manifold by constructing a diffusion operator on the data. Used in the NxtLight Phase 4 analysis to estimate the intrinsic manifold dimension of each boarding method's state-space trajectory.

Driver nodes (n_D)

The minimum set of nodes in a network that must receive external input signals in order to fully control the network's state. Network controllability theory (Liu et al., 2011) establishes that n_D / N (the driver fraction) characterizes how hard a network is to control; higher values indicate a harder-to-control system.

EDMD (Extended Dynamic Mode Decomposition)

A data-driven method for constructing a finite-dimensional linear approximation of the Koopman operator governing a nonlinear dynamical system. Used in the Nonlinear Controllability module to compare linear and nonlinear control effort requirements across boarding methods.

Entropy production rate (σ)

A thermodynamic quantity measuring the rate at which a system generates entropy as it evolves. In the boarding simulation, σ tracks the rate at which aisle congestion energy is dissipated; a smoothly decaying σ indicates controlled thermodynamic relaxation, while elevated oscillatory σ signals persistent disorder.

Fiedler value (λ_2)

The second-smallest eigenvalue of the graph Laplacian matrix of a network. Also called the algebraic connectivity, λ_2 measures how well-connected a graph is; values near zero indicate that the network is close to fragmenting into disconnected components. A falling Fiedler value provides a 4-minute early warning of aisle blockage in this analysis.

Fragmentation index

A composite metric defined as $Q / (\lambda_2 + 0.01) - 1$. High values indicate a network that simultaneously has strong community structure (high Q) and poor global connectivity (low λ_2), identifying it as fragmented. BTF achieves a fragmentation index of 1.99, uniquely triggering the hard-control regime flag.

Granger causality

A statistical test for determining whether time series X causally predicts time series Y, beyond what Y predicts of itself. Used in the Partial TE module via an F-test screen ($p < 0.05$) to identify direct causal edges in the passenger interaction network, separated from mediated (indirect) influence.

H_deg (degree entropy)

The Shannon entropy of the degree distribution of a network. High H_deg indicates a heterogeneous degree distribution spanning many values; low H_deg indicates a homogeneous network in which most nodes have similar degree. H_deg > 2.0 is one of two conditions that trigger the hard-control regime flag.

Hη (information pathology noise)

In the ITDAN framework, Hη denotes the joint uncertainty or entropy injected into the boarding network by the five information pathology types: asymmetric, incomplete (blackout), inconsistent (inversion), missing (conductance = 0), and malicious/adversarial.

Hub score

A composite metric combining betweenness centrality and degree entropy to identify nodes that are both topologically central and structurally diverse in their neighborhood. A true functional hub has high betweenness and high degree entropy; BTF's zone-boundary nodes have high betweenness but low degree entropy, producing low hub scores despite appearing as bottlenecks.

IB efficiency (Information Bottleneck efficiency)

The fraction of mutual information between a source community and a target community that is preserved when signals pass across the community boundary. Values near 1.0 indicate permeable boundaries that transmit compliance information effectively; lower values (e.g., Steffen's 0.614) indicate boundaries that block corrective information propagation.

ITDAN (Information-Theoretic Dynamics of Agent Networks)

The NxtLight sub-framework for modeling information pathologies in multi-agent networks. ITDAN classifies disruptions as one of five pathology types and injects calibrated noise into the simulation. The adversarial pathology type uses a Bayesian node-level prior (π_j) and edge-level CUSUM detection ($CUSUM_{ij}$) to model and identify intentional rule violations.

KSG-PTE (Kraskov-Stögbauer-Grassberger Partial Transfer Entropy)

A non-parametric estimator for partial transfer entropy that uses k-nearest-neighbor density estimation to condition out mediated influences. Applied after Granger screening to quantify the direct causal information flow on each edge in the boarding network.

Koopman operator

An infinite-dimensional linear operator that governs the evolution of observable functions of a nonlinear dynamical system. EDMD provides a finite-dimensional approximation; the Koopman framework allows linear controllability theory to be applied to the nonlinear boarding dynamics.

Laminarity (LAM)

A Recurrence Quantification Analysis metric measuring the fraction of recurrence points that form vertical line structures in the recurrence plot. High LAM indicates that the system dwells in particular states for extended periods before transitioning. WILMA's LAM of 0.988 is the computational signature of its three extended seat-type boarding phases.

Lyapunov exponent (λ)

A measure of the rate at which nearby trajectories in a dynamical system diverge or converge. A positive Lyapunov exponent indicates chaos: small perturbations

amplify exponentially. BTF's Lyapunov exponent crosses zero at t ≈ 0.5T, marking the chaotic mid-boarding phase transition at which a single luggage delay can cascade into full aisle blockage.

Manifold dimension

The intrinsic dimensionality of the state-space trajectory of a dynamical system, estimated here via Diffusion Maps. Lower manifold dimension indicates a more constrained, predictable system. Steffen operates on a 1-dimensional manifold; WILMA on 2 dimensions; BTF on 3 dimensions.

Modularity (Q)

A graph metric measuring the strength of division of a network into communities, defined as the fraction of edges within communities minus the expected fraction under a null model. Q spikes in BTF at mid-boarding when zone clusters momentarily self-isolate, a pattern captured by the NxtLight community detection module.

NxtLight Adaptive method

A novel boarding method proposed in this analysis that uses gate-area density sensors to feed live aisle state into the NxtLight ITDAN module. Group calls are dynamically paused when λ_2 drops below 0.3, allowing the network to recover connectivity before the next wave of passengers enters. Simulated TBT: 12.9 min (55% below BTF baseline).

NxtLight Framework

The computational analysis platform developed at NxtLight Workshop for modeling co-evolving networks. The framework applies up to 17 analytical modules spanning information-theoretic, dynamical systems, topological, and graph-theoretic methods simultaneously, with a three-pass ModuleRunner that couples outputs across modules.

PageRank

A node centrality measure originally developed for web link analysis that assigns importance scores based on the number and quality of incoming edges. In the ITDAN framework, PageRank identifies nodes whose state has the greatest downstream influence on the rest of the network.

Partial transfer entropy (PTE)

A conditional form of transfer entropy that isolates the direct causal information flow from source to target, removing the contribution of common causes and mediated paths. PTE_out is used as a coupling signal from the Partial TE module into the Hub Identification module in the NxtLight three-pass architecture.

Persistent homology (β_1)

A topological data analysis method that quantifies the shape of data by counting topological features (connected components, loops, voids) that persist across multiple scales. The first Betti number β_1 counts the number of independent loops in the state-space trajectory; higher β_1 indicates a more complex trajectory.

Phase control parameter (Φ)

Defined as P / D, the ratio of the passenger arrival rate (P) to the departure rate of seated passengers (D). $\Phi \gg 1$ corresponds to the crystallized phase (aisle blockage); $\Phi < 1$ corresponds to the productive dissipative phase. BTF maximizes Φ by concentrating all arrivals in a single rear zone; Outside-In boarding distributes P across the full cabin length.

RQA (Recurrence Quantification Analysis)

A nonlinear time-series analysis method that characterizes the dynamics of a system by analyzing its recurrence plot — a visualization of all times at which the system revisits a given state. Metrics extracted include laminarity (LAM) and recurrence rate (RR), used here to distinguish structured from stochastic boarding dynamics.

SINDy (Sparse Identification of Nonlinear Dynamics)

A data-driven method that uses sparse regression on a library of candidate functions to discover the governing equations of a dynamical system directly from time-series data. Near-perfect R^2 with zero active terms across all methods in this analysis indicates that boarding dynamics are strongly constrained by their structural rules.

Small-world sigma (σ)

A network topology metric defined as (C / C_random) / (L / L_random), where C is the clustering coefficient and L is the average path length. $\sigma > 1$ indicates small-world structure: high local clustering with short global paths. Used in the NxtLight Framework to characterize the baseline topology of the passenger network.

Spectral radius (ρ(A))

The largest absolute eigenvalue of the adjacency matrix A of a network. $\rho(A) > 1$ indicates operation above the edge-of-chaos threshold, where small perturbations can amplify. BTF has the highest spectral radius (22.5); WILMA has the lowest (12.1), quantifying WILMA's superior ability to dampen compliance failures.

TARP (Tests of Accurate and Reliable Posteriors)

A simulation-based inference diagnostic adapted here as a temporal structure test. TARP scores near 1.0 indicate that a time series has strong deterministic temporal structure distinguishable from randomness; scores near 0.5 indicate genuine but weaker structure. Zone-constrained methods (BTF, WILMA) score 1.0; stochastic methods (Random, Open Seating) score 0.5.

Transfer entropy (TE)

An information-theoretic measure of the directed statistical dependence between two time series, quantifying how much knowing the history of a source reduces uncertainty about the future of a target beyond the target's own history. TE_out identifies nodes that causally drive others; TE_in identifies nodes driven by others. Introduced by Schreiber (2000) and used here as the primary causal analysis tool.

WILMA (Window-Middle-Aisle)

A boarding strategy, also called Outside-In, that boards all window-seat passengers first, followed by middle-seat passengers, then aisle-seat passengers. By eliminating seat interference entirely, WILMA achieves near-optimal boarding time (19.6 min) with the highest compliance robustness of all established methods, as measured by spectral radius and IB efficiency.

References

[1] Steffen, J.H. (2008). Optimal boarding method for airline passengers. Journal of Air Transport Management, 14(3), 146–150. https://doi.org/10.1016/j.jairtraman.2008.03.003

[2] Nyquist, D.C. & McFadden, K.L. (2008). A study of the airline boarding problem. Journal of Air Transport Management, 14(4), 197–204. https://doi.org/10.1016/j.jairtraman.2007.06.006

[3] van Landeghem, H. & Beuselinck, A. (2002). Reducing passenger boarding time in airplanes. European Journal of Operational Research, 142(2), 294–308. https://doi.org/10.1016/S0377-2217(01)00294-6

[4] Ferrari, P. & Nagel, K. (2005). Robustness of Efficient Passenger Boarding in Airplanes. Transportation Research Record, 1915, 44–54. https://doi.org/10.3141/1915-06

[5] Milne, R.J. & Kelly, A.R. (2014). A new method for boarding passengers onto an airplane. Journal of Air Transport Management, 34, 93–100. https://doi.org/10.1016/j.jairtraman.2013.08.006

[6] Qiang, S.J. et al. (2014). Reducing airplane boarding time by accounting for passengers' individual properties. Physica A, 404, 127–137. https://doi.org/10.1016/j.physa.2014.02.043

[7] Steffen, J.H. & Hotchkiss, J. (2012). Experimental test of airplane boarding methods. Journal of Air Transport Management, 18(1), 64–67. https://doi.org/10.1016/j.jairtraman.2011.10.003

[8] Oliveira, J. & Aghezzaf, E. (2023). Balancing boarding time and passenger experience: a multi-objective approach. Transportation Research Part C, 148, 104025. https://doi.org/10.1016/j.trc.2023.104025

[9] Bernstein, N. (2012). Comment on "Optimal boarding method for airline passengers." Journal of Air Transport Management, 18(1), 68. https://doi.org/10.1016/j.jairtraman.2011.10.004

[10] Schreiber, T. (2000). Measuring information transfer. Physical Review Letters, 85(2), 461–464. https://doi.org/10.1103/PhysRevLett.85.461

[11] Watts, D.J. & Strogatz, S.H. (1998). Collective dynamics of small-world networks. Nature, 393, 440–442. https://doi.org/10.1038/30918

[12] Blondel, V.D. et al. (2008). Fast unfolding of communities in large networks. Journal of Statistical Mechanics, 2008(10), P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008

[13] Liu, Y.-Y. et al. (2011). Controllability of complex networks. Nature, 473, 167–173. https://doi.org/10.1038/nature10011

[14] Fiedler, M. (1973). Algebraic connectivity of graphs. Czechoslovak Mathematical Journal, 23(2), 298–305.

[15] Prigogine, I. & Stengers, I. (1984). Order Out of Chaos. Bantam Books.

[16] Kahneman, D. (2011). Thinking, Fast and Slow. Farrar, Straus and Giroux.

[17] Adams, J.S. (1965). Inequity in social exchange. Advances in Experimental Social Psychology, 2, 267–299. https://doi.org/10.1016/S0065-2601(08)60108-2

[18] Thaler, R.H. & Sunstein, C.R. (2008). Nudge: Improving Decisions about Health, Wealth, and Happiness. Yale University Press.

[19] Maister, D.H. (1985). The psychology of waiting lines. The Service Encounter, Lexington Books.

[20] Festinger, L. (1954). A theory of social comparison processes. Human Relations, 7(2), 117–140. https://doi.org/10.1177/001872675400700202

[21] Brunton, S.L. et al. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. PNAS, 113(15), 3932–3937. https://doi.org/10.1073/pnas.1517384113

[22] Edelsbrunner, H. & Harer, J. (2010). Computational Topology: An Introduction. American Mathematical Society.

[23] Seifert, U. (2012). Stochastic thermodynamics, fluctuation theorems and molecular machines. Reports on Progress in Physics, 75(12), 126001. https://doi.org/10.1088/0034-4885/75/12/126001

[24] Lemos, P. et al. (2023). Sampling-based accuracy testing of posteriors from custom-built neural network samplers. arXiv:2210.11357.

[25] Coifman, R.R. & Lafon, S. (2006). Diffusion maps. Applied and Computational Harmonic Analysis, 21, 5–30.

[26] Freeman, L.C. (1977). A set of measures of centrality based on betweenness. Sociometry, 40(1), 35–41.

[27] Jaeger, H. (2001). The echo state approach to analysing and training recurrent neural networks. GMD Technical Report 148.

[28] Marwan, N. et al. (2007). Recurrence plots for the analysis of complex systems. Physics Reports, 438, 237–329.

[29] Tishby, N. et al. (1999). The information bottleneck method. arXiv:physics/0004057.

[30] Williams, M.O. et al. (2015). A data-driven approximation of the Koopman operator. Journal of Nonlinear Science, 25, 1307–1346.

[31] Yuan, Z. et al. (2013). Exact controllability of complex networks. Nature Communications, 4, 2447.