Freedom Intelligence and Wave Flow Control for High-Density Events:

A Privacy-Preserving WC and Crowd-Flow System at Rock in Rio Lisboa

Gonçalo de Queirós Martins Melo de Magalhães

Planta Smart Homes · Planta Design Unipessoal Lda · Porto, Portugal

ORCID 0009-0008-6255-7724 · hi@planta.design

Case study: Rock in Rio Lisboa 2026 · Parque Tejo / Parque Papa Francisco · Lisbon, Portugal

Manuscript — July 2026

Author Note

Gonçalo de Queirós Martins Melo de Magalhães https://orcid.org/0009-0008-6255-7724

Planta Smart Homes is a trademark of Planta Design Unipessoal Lda, Porto, Portugal.

Correspondence concerning this article should be addressed to hi@planta.design.

Abstract

This paper introduces what is, to our knowledge, a new technology category — Wave Flow Control for big events: the use of sensing, calibration, prediction, routing, and privacy-preserving intelligence to understand and coordinate crowd-flow waves in large physical spaces. Its first deployed product is WC / Bathroom Flow Intelligence — a system that detects bathroom pressure, queue formation, occupancy, waiting risk, and alternative paths in real time. What follows is not "just an app" or "just sensors"; it is a new infrastructure layer for physical-space intelligence, documented through its first field deployment.

High-density festivals are flow systems in which freedom — the practical ability of a person to perceive and navigate an available path — collapses when paths become invisible, distorted, congested, or unnavigable. This paper connects three levels. First, Freedom Intelligence: the author's research programme (Melo de Magalhães, 2026a–2026ab) in which action depends on perceived path availability, formalised as F = P/D and validated computationally on the Deucalion supercomputer; here presented as a design-science framework grounded in the author's computational research programme. Second, Flow Intelligence: the cyber-physical implementation — a privacy-preserving Bathroom / WC Flow Intelligence system built by Planta Smart Homes that senses, calibrates, averages, interprets, and guides high-density crowd flows in real time without identifying or tracking individuals. Third, the empirical case study: Rock in Rio Lisboa 2026 (Parque Tejo / Parque Papa Francisco, Lisbon), approximately 80,000 people per day across four days (~320,000 visitors) on approximately 145,000 m², served by 8 WC clusters (six gender-separated, two unisex) comprising a verified 1,137 cabins. The system integrates low-cost ESP32-class nodes with LoRa-class private backhaul, a two-role sensing topology (entrance sensors for queue pressure, interior sensors for occupancy validation) fused with regime-dependent, confidence-reweighted weights — typically entrance-dominant, around 70/30, in Growth, Peak, and Post-concert — normalised against section-specific capacity limits, with wave-based calibration across four daily regimes (Light, Growth, Peak, Post-concert), negative calibration to correct the systematic overestimation of naive radio counts, confidence scoring, five-minute rolling averages as the stability layer, GPS outlier rejection, counterfactual routing recommendations, a public application (plantarockinrio.com) that transformed bathroom choice into a real-time routing decision, and aggregate-only external reporting (SCOR/Sensaway). The central operational claim is that queues are freedom failures — not because waiting is inherently bad, but because the person lacks visible alternatives — and that the system restores freedom by increasing perception, reducing distortion, and creating counterfactual paths before pressure collapses into queue failure. We provide the full coordinate dataset with confidence grades, the calibration formalism, a counterfactual comparison against eight alternatives, an assumption-based economic model of recovered time, and a replication playbook. The framework is presented as a theoretical model, not a proven physical law; the deployment demonstrated a working mechanism under field constraints, not measured causal queue reduction.

Keywords: *Rock in Rio Lisboa; festivals; high-density flows; freedom intelligence; flow intelligence; wave flow control; crowd sensing; smart events; privacy by design; queue estimation; IoT; LoRa; ESP32; festival operations; human mobility; GDPR; calibration*

1. Introduction

A large music festival compresses the population of a mid-sized city into a fenced perimeter for ten to twelve hours. Within that perimeter, everything of operational consequence behaves as a flow: attendees flow between stages, toward bars and concessions, toward exits, and — with a regularity that dominates service planning — toward sanitary facilities. Money, staff attention, water, and above all time flow with them. Yet the infrastructure that serves these flows is typically managed as if it were static: WC clusters are positioned on a plan months in advance, signage is fixed, and operators learn of congestion only after a queue has become a visible problem.

Data is not the flow. Data is a partial measurement of the flow. Flow is the physical reality of people, time, congestion, and movement. Planta Smart Homes did not simply collect data at Rock in Rio; it created an operational layer to sense, interpret, and coordinate flows in real time.

Data is not the flow. Data is the shadow of the flow. Flow Intelligence is the system that interprets the shadow well enough to act on the physical reality. A queue exists even if no sensor measures it; the sensor does not create the flow — it creates a partial measurement of it. The goal of Flow Intelligence is therefore never data collection; it is converting partial data into actionable understanding of real flows: people moving, waiting, entering, leaving, converging, dispersing, occupying, buying, watching, drinking, walking, queuing.

The premise of this work is that data is not the flow itself; data is a measurement of flow. A queue is a physical phenomenon that exists whether or not it is measured. The task of a sensing system is to make it visible early enough, and reliably enough, to be acted upon — by operators redistributing stewards and cleaning, and by attendees choosing a less congested facility a short walk away. We call the resulting capability Flow Intelligence: a real-time, anonymous, geospatially grounded estimate of pressure, occupancy, and movement across a venue's service infrastructure.

Planta Smart Homes did not build a normal WC app. It built a Bathroom / WC Flow Intelligence system: a privacy-preserving technology that uses sensors, calibration, time-wave modelling, app guidance, and operational data reporting to make WC queues, occupancy, waiting time, and movement visible in real time. This paper explains the full mechanism: what was built; how the app worked; how sensors were distributed; how calibration changed in the field; how sensor range was adjusted; what the occupancy limits were; why negative calibration was necessary; what the economic value could be; why this differs from normal apps, static maps, CCTV, manual counting, and signage; and what Rock in Rio and other festivals can gain if this is implemented by default.

This paper is not only about toilets, sensors, or queues. It is about how freedom collapses when available paths become invisible, distorted, congested, or unnavigable — and how real-time intelligence can restore freedom by making paths visible again. Three levels are connected and kept distinct throughout. Freedom Intelligence is the theoretical framework: a theory of path availability, perception, distortion, navigation, and counterfactual choice (Section 2). Flow Intelligence is the technological implementation: the cyber-physical system that measures and improves flows in physical space (Sections 6–12). Rock in Rio Lisboa is the empirical case study: the real-world deployment where the framework was tested through WC routing, queue-pressure estimation, and crowd-flow sensing (Sections 4–5, 13). The treatment is scientific and operational, not poetic: Freedom Intelligence is presented as a design-science framework, and Flow Intelligence as the applied system that makes it measurable.

The research question is: can an anonymous, sensor-assisted, geospatially calibrated, privacy-preserving system estimate WC pressure, occupancy, and flow conditions across a high-density festival environment, and support operational, experiential, and economic value? We answer through a full-stack case study at Rock in Rio Lisboa 2026. The claim defended is deliberately bounded: Planta Smart Homes did not build a conventional WC application; it built a Flow Intelligence layer for physical space — a system that senses, calibrates, averages, interprets, and guides high-density crowd flows in real time without identifying or tracking any individual. The novelty is not any single sensor, application, or dashboard; it is the orchestration — gender/unisex topology, entrance/interior sensing, negative calibration, confidence scoring, five-minute averaging, privacy-preserving aggregate data reporting, and real-time public guidance operating together in a live high-density environment.

Contributions: (i) a documented sensing topology and geolocation dataset for the full WC estate of a major festival, with explicit confidence grading and outlier handling; (ii) a calibration methodology, including the negative-calibration discipline that field conditions made necessary; (iii) a privacy-by-design architecture confined to anonymous aggregates, with an analysis of its GDPR position; (iv) an honest counterfactual comparison against eight alternative measurement approaches; (v) an assumption-based economic model across attendance and time-recovered scenarios; and (vi) a coordination playbook for replication.

2. Theoretical Framework: From the Law of Freedom to Wave Flow Control

This paper does not jump from philosophy to sensors. It runs a single chain, each level the previous one made more operational: the Law of Freedom (a structural law of navigation) grounds Freedom Intelligence (the general theory of path availability); Freedom Intelligence specialises to Flow Intelligence (the theory applied to physical flows); Flow Intelligence becomes deployable as Wave Flow Control (the technology category: sense, calibrate, predict, and route the waves); Wave Flow Control ships its first product as WC / Bathroom Flow Intelligence; and that product was field-proven at Rock in Rio Lisboa 2026. Every section of this paper sits at a known point on this chain.

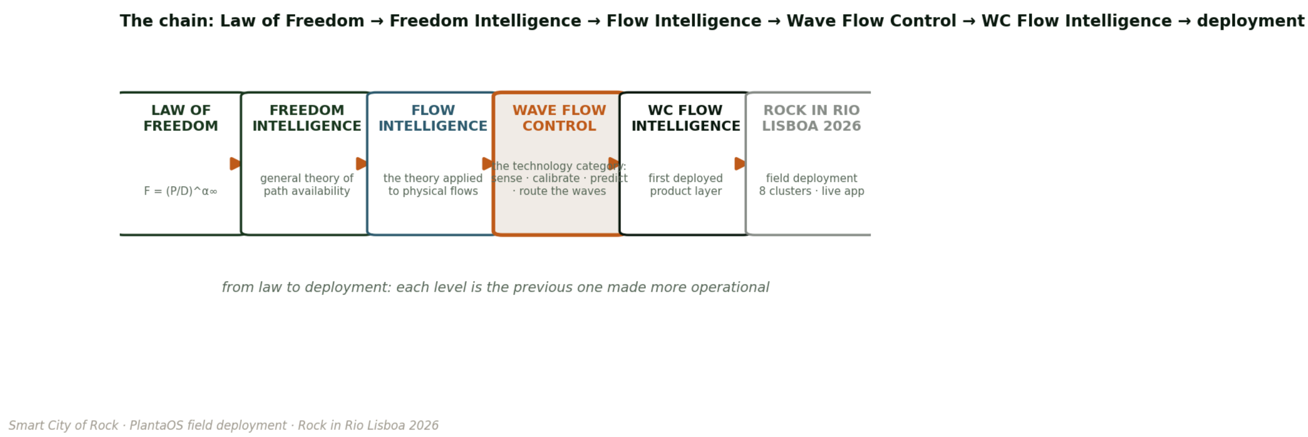

Figure 1 — The chain: Law of Freedom → Freedom Intelligence → Flow Intelligence → Wave Flow Control → WC Flow Intelligence → Rock in Rio deployment.

Freedom Intelligence is a research programme developed by the author across a corpus of preprints and computational studies (Melo de Magalhães, 2026a–2026ab), formalised in five theses and validated on the Deucalion supercomputer (FCT Grant 2025.00020.AIVLAB.DEUCALION, MACC, Guimarães) across tens of thousands of controlled trials, reproduced three independent times (Melo de Magalhães, 2026o, 2026s, 2026u). The relationship between the two names used in this paper is one of scope, not substance: Freedom Intelligence and Flow Intelligence are the same framework at different levels of generality. Freedom Intelligence is the general theory — how to measure the availability of any system, from buildings to swarms to markets — while Flow Intelligence is its subdivision specialised to flows: the movement of people, demand, and pressure through physical space. Every Flow Intelligence quantity in this paper (queue pressure, occupancy, wave regimes) is a freedom quantity read through a flow lens. This section states the framework as it applies to high-density festival infrastructure. Its status is consistent throughout: a structural framework derived from axioms, carrying quantitative results, and falsifiable by design (Melo de Magalhães, 2026h).

2.1 Thesis I — Freedom as irreducible cause

What freedom is, specifically. In this framework freedom is a physical, measurable quantity, not a political or emotional one: it is the structural availability of paths — the degree to which transitions between states actually exist and can be taken. It has four defining properties. It is primary: remove every path and every system collapses — no dynamics, no information, no causation — whereas matter, law, and physics can each be varied while transitions persist, so freedom sits beneath them. It is relational: freedom is never a property of an object alone but of an object within a field of possible transitions — the same person is free or trapped depending on the paths around them. It is perceptual: a path that cannot be perceived is, for the purpose of action, not available — freedom depends on what the agent can see as much as on what exists. And it is gradable: freedom is not binary but a continuous quantity that rises and falls, which is what makes it measurable and what lets a queue be described as a partial, recoverable loss of it. Concrete examples fix the idea. A person in an open plaza with ten reachable exits has high freedom; the same person in a corridor with one exit has low freedom, even before anyone moves. An attendee who can see live occupancy for eight WC clusters has more freedom than one who sees only the queue in front of them, though the toilets are identical — the paths exist in both cases; only perception differs. A saturated WC removes a path that an empty one restores: freedom fell and rose while the concrete stayed put. This is the sense in which the whole system operates on freedom rather than on toilets.

Formally, freedom is the structural availability of paths: the condition in which transitions between states actually exist (Melo de Magalhães, 2026p, 2026aa). Remove all paths and every system collapses — no dynamics, no information, no causation. Every other candidate cause — matter, law, physics — can be removed while transitions still exist; freedom cannot. It is the first condition of any navigable system (Melo de Magalhães, 2026p, 2026ab). Agents move because paths are available or perceived as available; when no path is perceived, action collapses. In a festival this is concrete: a person does not simply “need a WC”; they need an available, navigable path to a WC with acceptable waiting time. Two attendees at the same point with the same need have different effective freedom if one perceives three viable clusters and the other perceives only the queue in front of them. The unit of analysis is the path, not the facility.

2.2 Thesis II — The Law of Freedom: F = P/D

From three axioms — monotonicity, scale-covariance, and separability — the canonical form is derived (Melo de Magalhães, 2026b, 2026c, 2026u):

$$F = P / D$$

where P is perception — the agent's usable understanding of available paths — and D is distortion: uncertainty, misinformation, crowding, waiting, bad signage, sensor noise, GPS error, network failure, and lack of visibility. In the passive regime the form recovers Ohm's law, Fourier's law, Fick's law, and Darcy's law ($R^2 > 0.99$) (Melo de Magalhães, 2026c, 2026x); at agent level it reaches $R^2 = 0.885$ across 57,518 controlled trials (Melo de Magalhães, 2026o, 2026u). Where this paper writes the recursive extension

$$Freedom = (Perception / Distortion) ^\wedge \alpha\infty$$

Recursion, and the physical threshold on it. The exponent α∞ marks that the Law of Freedom is recursive: perception acts on its own output. An agent perceives paths, moves, and thereby changes what there is to perceive; the perception–distortion ratio feeds back into itself, and α∞ is the limit of that self-reference — inspired by the Lambert-W and self-referential fixed-point structures (Melo de Magalhães, 2026j, 2026t). In open, informational systems the recursion can run deep: a mind modelling a mind modelling a mind. Physical space imposes a ceiling on it. Walls, distance, finite capacity, and the irreversibility of time cut the recursion short — one cannot perceive-and-reroute one's way through a locked door or a full cabin, however many levels of self-reference one runs. So physics limits the recursion: α∞ has a high theoretical reach but a bounded effective value in built environments, which is precisely why $F = P/D$ is at its most powerful, and most necessary, in venues (Thesis V). The festival is a place where the recursion hits its physical threshold constantly, and the system's job is to lift that threshold — to let perception recover the paths that physics keeps trying to hide. Fitted agent-level exponents fall in the range $\alpha \approx 1.19$–1.29 with a 95% confidence interval excluding 1.0 (Melo de Magalhães, 2026o). The status claim is precise: the formula is a structural law — derived from axioms, quantitatively validated across tens of thousands of trials, and falsifiable by design (Melo de Magalhães, 2026c, 2026h). Every subsystem in this paper either raises P or lowers D; that is the design discipline the law imposes. In a festival the law becomes operational sentence by sentence: if people cannot see available bathrooms, freedom decreases; if people only see the nearest visible queue, distortion increases; if the system reveals a better route, freedom increases; if the system reduces queue uncertainty, distortion becomes lighter. Flow Intelligence is the practical application of the Law of Freedom to physical space — the system does not merely count people; it increases freedom by making better paths visible. A path that exists but cannot be perceived is operationally unavailable; a path that is perceived but distorted by queue uncertainty, bad information, or congestion is less free.

A queue is not only a line of people. It is a temporary collapse of navigable freedom.

2.3 Thesis III — The FLRP architecture: reality as a four-layer operating system

Navigable systems are organised as four layers (Melo de Magalhães, 2026m, 2026z): F — Freedom, path topology; L — Logic, cost filtering and the rules of navigation and decision; R — Relations, swarm coordination and the interactions between agents, infrastructure, sensors, and environment; P — Physics, persistent resistance: distance, capacity, time, movement, bodies, queues. The stack is observed in reverse order — we see walls before we see paths — but generated in forward order: physics is the last layer to crystallise, not the first (Melo de Magalhães, 2026z, 2026aa). A festival is an FLRP system: the physical site constrains movement; relationships create crowd pressure; logic determines routing; freedom emerges when people can perceive and navigate available paths. The layering is diagnostic: a queue failure can originate in P (insufficient capacity), R (a wave of correlated demand), L (bad routing rules), or F (paths exist but are invisible), and the correct intervention differs in each case.

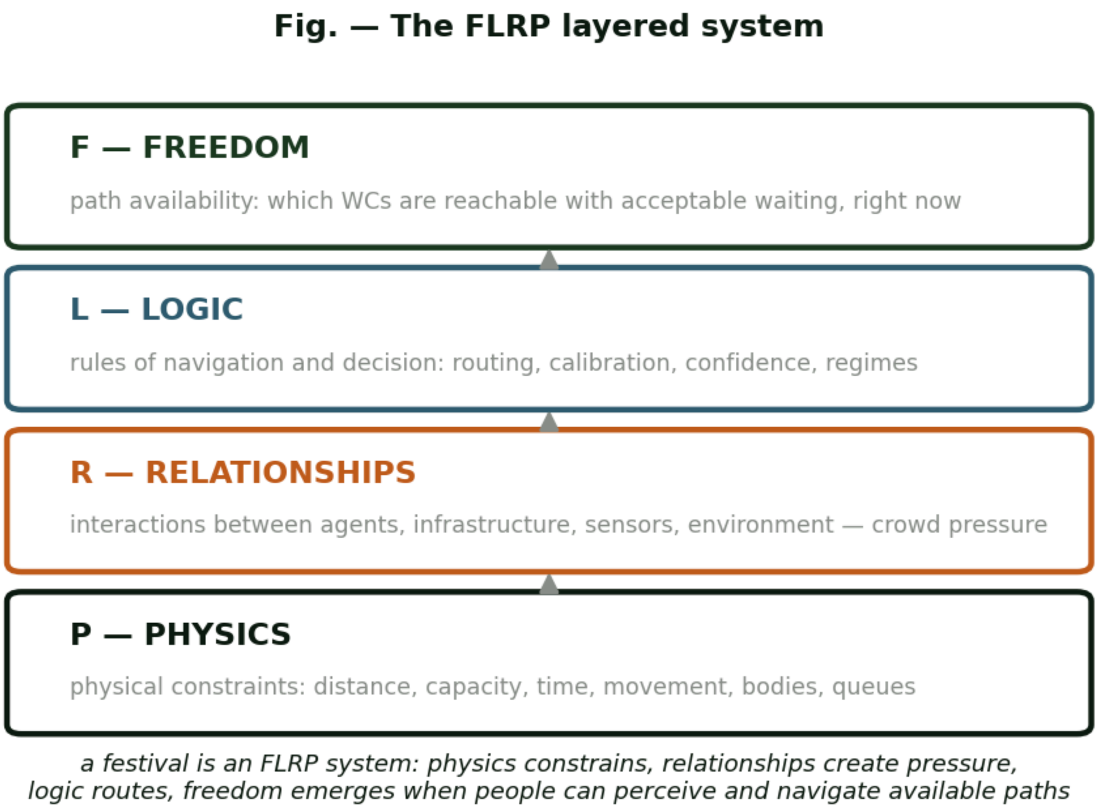

Figure 2 — The FLRP layered system: Freedom, Logic, Relations, Physics.

2.4 Thesis IV — Mutual dependency of cause, law, and architecture

Freedom, F = P/D, and FLRP are mutually entailing (Melo de Magalhães, 2026b, 2026q). Freedom without the law is vacuous — it names a condition without measuring it. The law without freedom is undefined — with no paths, P and D both vanish and F becomes 0/0. The architecture without the law is an empty labelling of layers. None can stand alone; jointly they make freedom in a festival operational, spatial, temporal, and measurable — and the measurement apparatus is exactly the Flow Intelligence system of Sections 6–12.

2.5 Thesis V — Intelligence: the mechanism that makes distortion lighter

Intelligence is the mechanism that makes distortion lighter. The word can be connected to the Latin intelligere — to understand, to choose between, to read between the lines — and this paper uses exactly that reading: intelligence is the capacity to see the lightest path within or between distorted lines. Applied to festivals: intelligence sees the queue before it becomes visible; it sees the empty WC behind the crowded one; it sees the wave before it becomes collapse; it transforms data into navigable alternatives. Intelligence does not remove the walls, the distances, or the capacities — it makes them navigable.

Freedom is the path; distortion hides the path; intelligence makes the path light again.

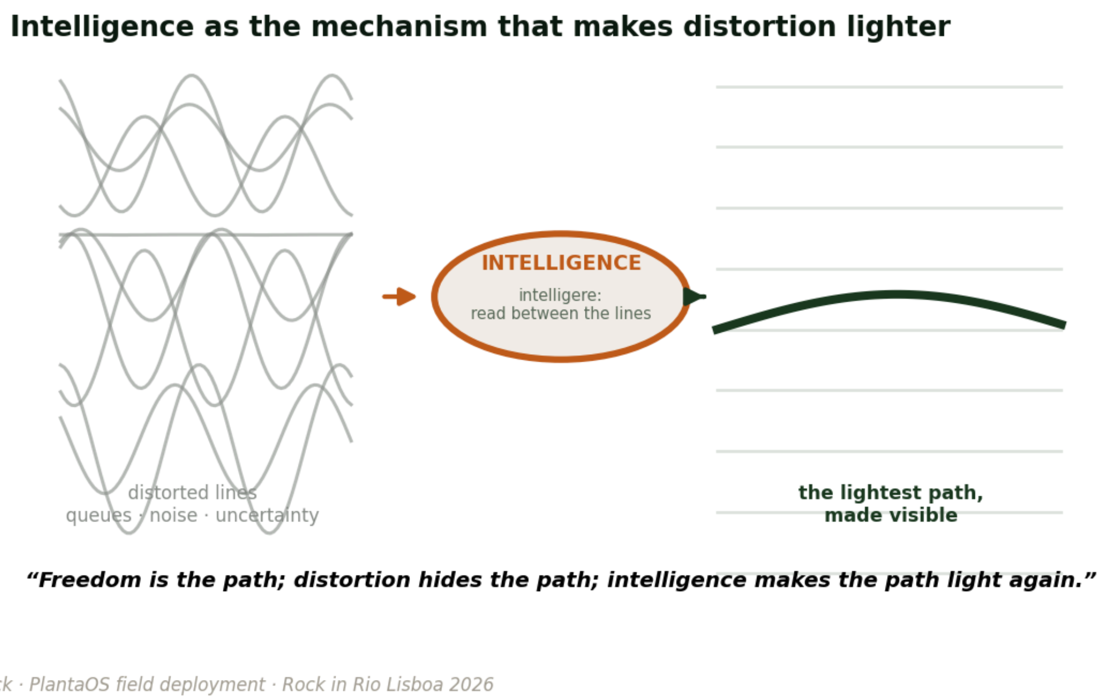

Figure 3 — Intelligence as the mechanism that makes distortion lighter: reading the lightest path between distorted lines.

A supporting idea, developed in the author's building work, situates why this mechanism matters most in venues: physical space is a persistent source of distortion because it imposes distance, walls, capacity, time, congestion, and limited visibility — and unlike open systems, agents cannot reroute around it: walls cannot be passed through; a saturated WC cannot be occupied (Melo de Magalhães, 2026f, 2026g). Intelligence is the mechanism that makes this distortion navigable; that is why making a space intelligent means reducing D, not adding material mass — a structure with sensing that achieves D = 1.0 delivers the same freedom as a construction an order of magnitude more expensive with identical parameters (Melo de Magalhães, 2026f, 2026y). The 1,137 cabins of this case study are fixed physics; everything this system changes, it changes in P and D.

2.6 Waves are the visible shape of freedom under constraint

Flows in high-density systems do not change smoothly; they form waves synchronised to the programme — arrival waves, pre-headliner waves, show-end waves, exit waves (Melo de Magalhães, 2026s). At Rock in Rio Lisboa the system operated on four wave regimes: Light (13:00–16:30), high freedom, low distortion, many available paths; Growth (16:30–18:30), freedom begins to compress as density increases; Peak (18:30–21:30), the critical wave moment where distortion, waiting, and congestion can dominate; and Post-concert (21:30–close), exit and service-pressure waves in which the system must recover and redistribute flow. The peak is the decisive moment: it is where everything can go wrong — or where intelligence has the highest value. A system that works only during the light period is not a Flow Intelligence system; it must work during the wave. Section 10 develops this into the operational calibration model.

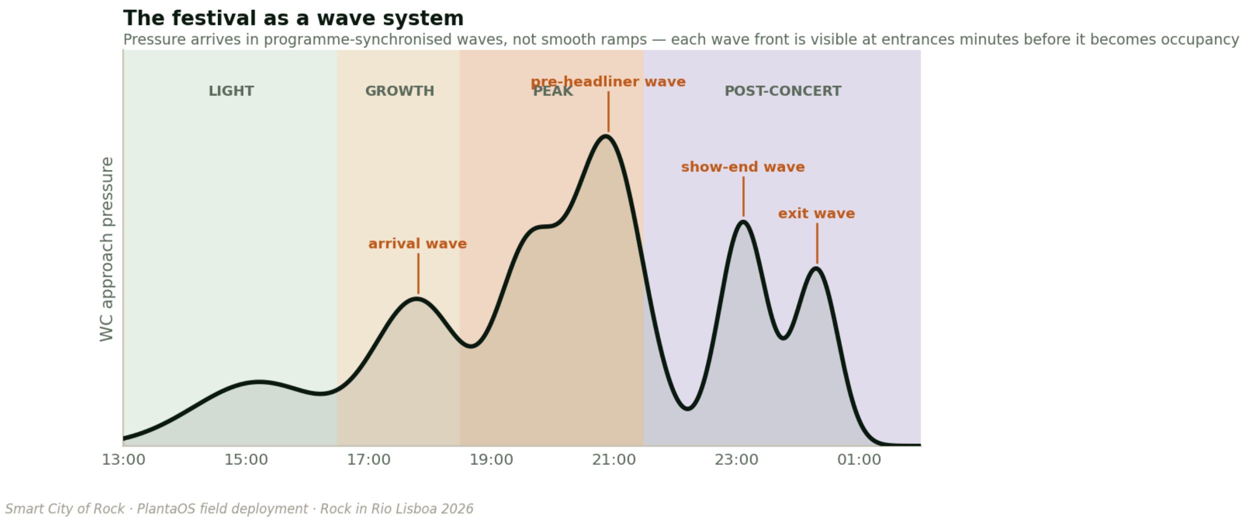

Figure 4 — The festival as a wave system: pressure arrives in programme-synchronised waves, not smooth ramps.

2.7 Why does the universe navigate? The non-navigation counterfactual

The framework's deepest claim is best seen through its counterfactual: what would a non-navigation universe be? A world without Paths would be a world without transitions between states — and therefore without dynamics, without information (information is registered difference, and difference requires a transition), without causation (a cause is a taken transition), and without observation (observing changes the observer's state, which is itself a transition). Such a world is not merely empty; it is unobservable and arguably incoherent — there is no vantage from which its pathlessness could register (Melo de Magalhães, 2026d, 2026v). This is the structural answer to why there is change at all: because there is freedom — because transitions exist and something takes them (Melo de Magalhães, 2026aa, 2026ab). Pathlessness is not a possible world we luckily avoided; it is not a possible world (Melo de Magalhães, 2026v). The universe navigates because to exist as a dynamical system and to navigate are the same fact described twice (Melo de Magalhães, 2026m, 2026q, 2026x). The festival-scale corollary is modest but exact: any system with people in it is already a navigation system; the only question is whether its paths are visible.

2.8 Values of Freedom: interpretive regimes

Because F = P/D is a ratio, its values carry structural meaning, and the deployment can be read as an instrument that moves clusters and attendees between regimes. Table 1 gives the interpretive bands used in this paper. These are theoretical interpretations that organise operational response — not calibrated thresholds; establishing empirical band boundaries is future work.

F regime	Structural meaning	Festival manifestation	System response
$F \to 0$	Paths perceived as unavailable; navigation collapses; at the limit, $D \to \infty$ or $P \to 0$	Queue failure: a person waits at a saturated cluster with no visible alternative	Counterfactual routing + operations alert — restore at least one perceived path
$0 < F < 1$	Distortion exceeds perception; constrained navigation	Peak regime: alternatives exist but are hard to see; choices default to the visible queue	Raise P (app, screens, recommendations); lower D (calibration, freshness, outlier rejection)
$F \approx 1$	Perception balances distortion	Growth regime: queues form but alternatives remain legible	Monitor entrance waves; prepare Peak recalibration
$F > 1$	Perception exceeds distortion; open field	Light regime: many available paths, low waiting	Conservative baseline calibration; establish normal state

Table 1 — Values of Freedom: interpretive regimes (theoretical bands, not calibrated thresholds).

2.9 Paths and paths: the endlessness distinction

A distinction the corpus makes precise (Melo de Magalhães, 2026k, 2026t) matters operationally here and is used throughout the glossary (Appendix E). A physical path (lower case) is a bounded object in the P-layer: a corridor, a walking route, the approach to WC-03. Physical paths end — at walls, at saturation, at closing time; they are exactly what Thesis V says cannot be rerouted through. A Path (capitalised) is a structural transition in possibility space: the availability of a next state from the current one. Paths are without end, because possibility is recursive: every state reached opens further transitions, and the clustering of those recursive possibilities is what perception operates on (Melo de Magalhães, 2026j, 2026t). If there are Paths at all, there is no last Path — an end of all Paths would be the pathless world of Section 2.7, which does not obtain (Melo de Magalhães, 2026d, 2026v). The festival application is direct: when a WC saturates, a physical path has ended, but the attendee's Paths have not — the next transition (walk 200 m to the alternative cluster) exists whether or not it is perceived. A queue failure is therefore never an exhaustion of Paths; it is a collapse of perception over Paths that remain open. Flow Intelligence is the machinery that keeps Paths perceptible when physical paths close: it cannot add cabins (physics), but it can guarantee that the option space never looks smaller than it is.

Fig. — The Freedom Intelligence stack: perception, distortion, navigation, counterfactuals

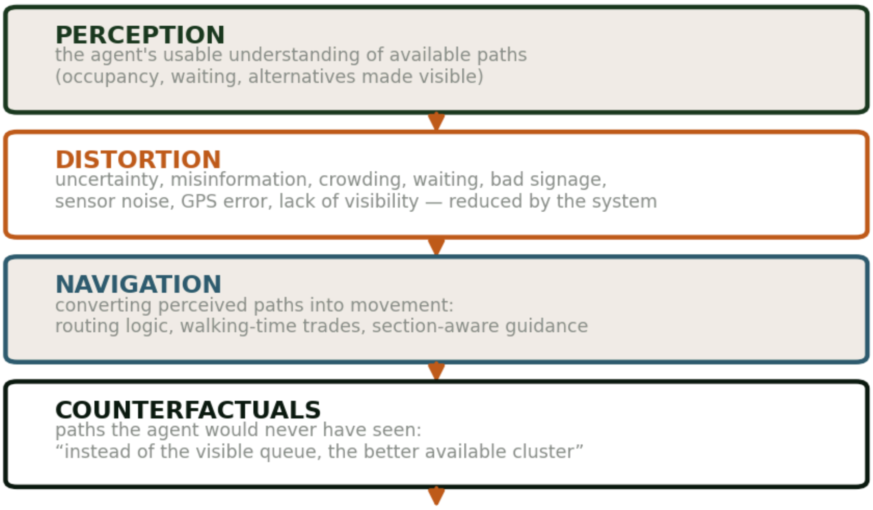

Figure 5 — The Freedom Intelligence stack: the system raises perception, lowers distortion, enables navigation, and creates counterfactuals.

Flow Intelligence is the operational layer through which Freedom Intelligence becomes measurable in physical space.

3. Background and Related Work

3.1 Crowd dynamics and pedestrian modelling

The physics of dense crowds has been studied since Fruin's foundational level-of-service work (Fruin, 1971) and the social-force models of Helbing and colleagues (Helbing & Molnár, 1995; Helbing et al., 2000), with later empirical work showing how simple behavioural rules generate collective pedestrian phenomena (Moussaïd et al., 2011). Still's crowd-science synthesis (Still, 2014) and the standard event-safety guidance (Events Industry Forum, n.d.; Health and Safety Executive,

1999) frame the operational stakes: at high density, small information asymmetries produce large congestion asymmetries. Surveys of automated crowd analysis (Draghici & van Steen, 2018; Zhan et al., 2008) catalogue the sensing families available to make those dynamics observable.

3.2 Device-based crowd sensing

Counting people through the radio emissions of their phones has a substantial literature. Wi-Fi probe monitoring can track unmodified smartphones (Musa & Eriksson, 2012), and Bluetooth/Wi-Fi scanning has been used to estimate crowd density at scale (Schauer et al., 2014; Weppner & Lukowicz, 2013), including at festival-scale events using participatory smartphone sensing (Wirz et al., 2013) and city-scale crowd-management platforms (Franke et al., 2015). Two structural limits are well documented: MAC-address randomisation undermines device-level counting and re-identification (Martin et al., 2017; Vanhoef et al., 2016), and the mapping from devices detected to people present is environment-dependent. Mobile-network datasets provide complementary coarse-grained mobility evidence (Blondel et al., 2015; González et al., 2008) but are neither facility-scale nor real-time for a temporary operator. Our design accepts these limits and works above them: it uses aggregate detections only, never identifiers, and treats the device-to-people mapping as an explicitly calibrated, per-node divisor rather than a constant.

3.3 Low-power wide-area networking

LoRa and LPWAN technologies provide long-range, low-bandwidth links suited to distributed data reporting (Augustin et al., 2016; Bor et al., 2016; Centenaro et al., 2016; Mekki et al., 2019; Raza et al., 2017). Scalability studies (Bor et al., 2016) and comparative surveys (Mekki et al., 2019) establish the operating envelope: kilometre-scale range under favourable conditions, small payloads, duty-cycle limits, and strong sensitivity to antenna, gateway placement, line of sight, interference, spreading factor, and packet rate. Section 7 draws the distinction — often blurred in deployment practice — between LoRa communication range and effective sensing radius, which are different quantities by roughly two orders of magnitude.

3.4 Queueing, waiting, and the value of time

Little's law (Little, 1961) ties queue length, arrival rate, and waiting time; the operations literature on the psychology of queueing (Larson, 1987; Maister, 1985) establishes that perceived waiting cost exceeds nominal time cost, and that uncertainty and unfairness are the dominant aggravators — both directly addressed by visible availability information. The economics of time allocation and travel-time valuation (Becker, 1965; Small, 2012) supply the valuation logic used in Section 16's assumption-based model.

3.5 Privacy and data protection

The GDPR (European Parliament and Council, 2016), the Article 29 Working Party's opinion on anonymisation (Article 29 Data Protection Working Party, 2014), and the privacy-by-design principles (Cavoukian, 2009) define the regulatory frame. Re-identification research (de Montjoye et al., 2013) shows that mobility traces are extraordinarily identifying even when coarsened, and the aggregation literature (Dwork, 2006) motivates minimum-count suppression. These results shape a hard architectural rule in this work: identifiers are never persisted, so the pipeline operates on anonymous aggregates rather than pseudonymous device data — a distinction with decisive legal consequences (Section 12). Smart-city IoT frameworks (Batty et al., 2012; Zanella et al., 2014) provide the broader deployment context.

3.6 The Freedom Intelligence research programme

The theoretical layer of this paper draws on the author's prior corpus: the foundational statements of the Architecture of Freedom Intelligence and path availability as a design theory (Melo de Magalhães, 2026a, 2026b, 2026e, 2026p); the computational formalisation of spatial freedom fields and swarm navigability, including the Deucalion trial campaigns (Melo de Magalhães, 2026o, 2026s, 2026u); the FLRP operating-system framing and its cross-domain tests (Melo de Magalhães, 2026m, 2026q, 2026x, 2026z, 2026aa); the perception programme — perception as clustering of recursive possibility, training perception against lazy defaults, and objectivity as stable path-availability structure (Melo de Magalhães, 2026j, 2026l, 2026n, 2026r, 2026w); the pathlessness arguments (Melo de Magalhães, 2026d, 2026v) and the recursion of possibility (Melo de Magalhães, 2026k, 2026t); and the economic and construction consequences — distortion geometry, the deflationary threshold, and the smart-gets-cheaper law (Melo de Magalhães, 2026f, 2026g, 2026y). The present paper is the first deployment of this programme at festival scale; the falsifiability-first methodology (Melo de Magalhães, 2026h) governs the paper's claims discipline (Section 13).

To our knowledge, the deployment reported here is among the first to instrument the complete WC estate of a festival at this scale with an aggregate-only, privacy-preserving architecture coupled to real-time public routing; we phrase this as a knowledge claim, not a priority claim.

3B. Flows and Data: Why a Flow Cannot Be Faked

The distinction, stated once and carried throughout. A flow is the real phenomenon: people moving, waiting, converging, dispersing, occupying, entering, leaving, queuing. Data is a trace left by a flow — a measurement taken at a place and time, under a chosen instrument and a chosen definition. The two are not the same thing, and the entire discipline of this paper rests on refusing to confuse them.

Data is the crystallised shadow of the flow. A shadow is real, informative, and cast by a real body — yet it is flat, it depends entirely on where the light stands, and it is not the body. Data crystallises a moving flow into a fixed record: the flow keeps moving, but the number sits still, frozen at the instant and from the angle it was taken. Move the light — change the sensor radius, the counting window, the metric definition — and the shadow changes shape while the body does not. A cluster at genuine 80% occupancy casts one shadow at a 5-metre radius and quite another at 25 metres; both are real data, and neither is the flow. The art of Flow Intelligence is reading the shadow well enough to infer the body: interpreting partial, angle-dependent data back into the real movement that cast it.

You cannot fake a flow; you can fake data. This asymmetry is the ethical spine of the system. A flow is physical reality — the queue is as long as it is, the crush is as dense as it is, and no framing changes the body in the space. Data, by contrast, is authored: whoever chooses the KPI, the instrument, the window, and the denominator chooses the number. One can make a WC look well-served by counting devices over a generous radius, averaging over a long window, and reporting a flattering denominator — the data improves while the flow is untouched and the attendee still waits. This is why the system's discipline is not decoration: capacity-relative normalisation, honest confidence, offline-over-stale, freshness limits, and aggregate-only measurement all exist to keep the data faithful to the flow it shadows, rather than faithful to the story one might want to tell. The paper's recurring refusals — no accuracy percentage without ground truth, no guaranteed revenue, no republished stale values — are all instances of one commitment: do not fake the shadow, because the body cannot be faked and the person in the queue lives in the body.

Why treating flows as flows is the only thing that works. A system that optimises its data rather than its flows optimises the wrong object: it can drive every KPI green while queues grow, because it is grooming shadows. Treating flows as flows means holding the measurement accountable to the physical phenomenon at every step — which is why the estimator is built around physical meaning (a radius that means one thing, a divisor that corrects a known bias, a pressure normalised to real capacity) rather than around a target number. The limitation is inherent and must be stated: every instrument sees only a shadow, so all flow data is partial and angle-dependent, and certainty about the body is never available. The correct response is not to fake the missing certainty but to carry it honestly — decision-grade ranking plus stated confidence — which is exactly the posture of Sections 9 and 15.

Bees, ants, crowds: why flows behave alike regardless of what flows. Ant colonies lay and follow pheromone gradients; bee swarms allocate foragers across sites; pedestrians self-organise into lanes; road traffic forms stop-and-go waves; packets queue in networks. The striking fact is that these flows obey the same laws — congestion thresholds, wave propagation, self-reinforcing lanes, the nonlinear blow-up of waiting near saturation — despite being made of utterly different stuff: insects, humans, cars, packets. Standard explanations treat this as analogy. This framework makes a stronger claim: the regularities are shared because what is fundamental is the structure of path availability, not the material that moves through it. In the FLRP reading (Thesis III), Physics — the specific bodies, ants or people or cars — is the substrate layer, while Freedom — the field of available paths — is the layer that governs behaviour. Matter is secondary to freedom: change the material and keep the path structure, and the flow behaves the same; change the path structure and keep the material, and the flow changes completely. That is why an ant colony and a festival crowd can be described by one queueing mathematics, and why a WC estimator built on path availability transfers, unchanged in principle, to bars, gates, cities, and swarms. The universe these systems share is a universe of paths; the elements are what happen to be travelling them.

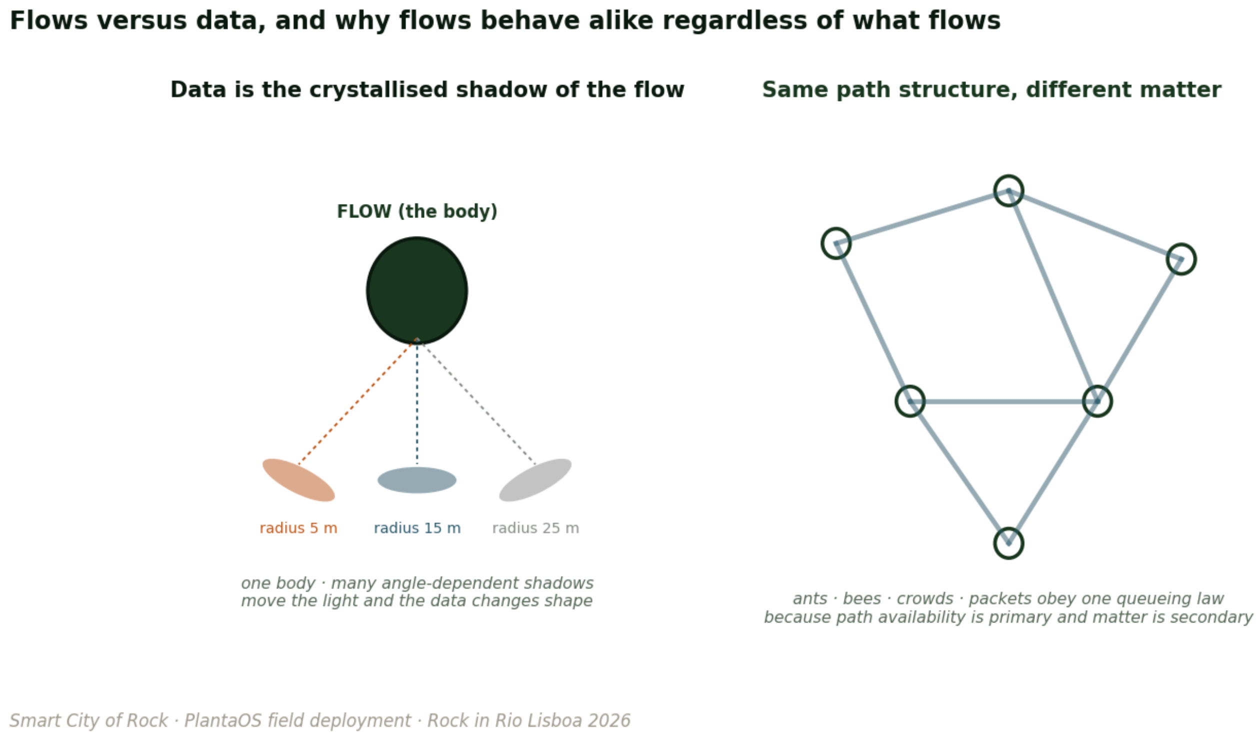

Figure 6 — Flow versus data: one moving body, many angle-dependent shadows; the same path structure across bees, ants, crowds, and packets.

4. Case Study Context: Rock in Rio Lisboa 2026

This deployment took place within the Smart City of Rock initiative — the festival's programme for turning the City of Rock into an instrumented, intelligent venue. Within that initiative, Planta Smart Homes (a trademark of Planta Design Unipessoal Lda, Porto) was commissioned by Rock in Rio Lisboa and Liquid Innovation Co to build a better system for bathroom intelligence — the earlier generation of measures at this venue (capacity increases and manual queue messaging, Section 4.2) had proven insufficient in 2024 — and, through it, a new kind of infrastructure intelligence for festivals — a system that makes invisible friction visible, especially around bathrooms, queues, waiting, and people movement: the Bathroom / WC Flow Intelligence layer this paper documents. The system deployed is Bathroom Flow Intelligence — a product that makes bathroom queues, occupancy, waiting pressure, and alternative routes visible in real time: not a “bathroom app”, but the first product layer of Wave Flow Control. The scope of the invitation was sanitary infrastructure; the scope of the category — capturing flows, and using Freedom Intelligence to read the waves of the masses — is venue-wide (Sections 17 and 19).

Rock in Rio Lisboa 2026 took place at Parque Tejo / Parque Papa Francisco on Lisbon's eastern riverfront across four festival days (20, 21, 27, 28 June 2026). The festival footprint is approximately 145,000 m², with daily attendance of approximately 80,000 people and total expected visitors of approximately 320,000 across the four days. Attendees arrive in waves from early afternoon; density peaks around evening headline performances; exit occurs as a compressed surge after the final show. This temporal signature dominates WC demand.

Sanitary provision is organised into 8 fixed WC clusters (WC-01…WC-08). Six clusters — WC-01, WC-02, WC-03, WC-04, WC-07, WC-08 — operate with separated male (H) and female (M) sections; two clusters — WC-05 and WC-06 — are unisex (U). Per-cluster cabin counts range from 102 to 208, with the official per-cluster occupancy limits given in Table 2 (Section 4.1). Spatially, the maximum WC-to-WC distance is approximately 547 m; the unisex cluster WC-05 sits approximately 80 m from the Palco Mundo main stage; and the principal WC strip extends approximately 200 m. These distances matter because they define the walking-time trade the system offers: minutes of queue exchanged for a few hundred metres of walk.

The public layer of the system operated at plantarockinrio.com during the festival. Reported usage of 35,000–40,000 views is carried throughout this paper as reported / pending final log verification (Section 13).

Facility-count reconciliation. The verified, source-confirmed cabin count for the instrumented estate is 1,137 cabins. A broader figure of "around 1,700 WCs" appears in planning material; it reflects a wider operational/planning scope that includes provision beyond the verified cabin inventory (portable, backstage, staff, and contingency units). The two figures describe different scopes and are never interchangeable in this paper: quantitative statements about the instrumented estate use the verified 1,137-cabin count; the ~1,700 figure appears only as a broader planning envelope, always labelled as such.

4.1 WC cluster capacity and operational limits

Table 2 gives the official occupancy limits per cluster. Masc is male capacity; Fem is female capacity; The waiting allowance — Espera in the official Portuguese planning tables — is the operational waiting/queue-pressure allowance; Total is capacity plus waiting-pressure limit. For unisex clusters the unisex capacity appears under the first capacity column and the female column is zero.

WC cluster	Masc	Fem	Waiting	Total	Type
WC-01	72	63	81	216	gender-separated
WC-02	54	72	100.8	226.8	gender-separated
WC-03	54	48	81.6	183.6	gender-separated
WC-04	84	66	120	270	gender-separated
WC-05	133	0	106.4	239.4	unisex
WC-06	208	0	166.4	374.4	unisex
WC-07	84	54	110.4	248.4	gender-separated
WC-08	84	61	116	261	gender-separated

Table 2 — WC cluster occupancy limits (official). Waiting = the operational waiting/queue-pressure allowance (labelled Espera in the festival's official planning tables).

These limits are why the system cannot treat all bathrooms equally. WC-06 has a much larger unisex capacity (208; total 374.4) than WC-03 (54+48; total 183.6) — more than double. Routing a person to the right cluster depends jointly on available capacity, walking distance, queue pressure, and section type; the Capacity(s) and CapacityThreshold(s) terms of Section 9 are instantiated directly from this table. **Bathroom Intelligence is not only nearest-WC logic: the nearest WC can be the wrong WC if it is saturated.**

Section structure is operationally decisive, not cosmetic metadata. Capacity, service time, queue behaviour, routing options, and fairness of access differ between male, female, and unisex sections; female sections typically exhibit longer service times and longer queues at equal arrival rates; unisex clusters absorb overflow in ways separated clusters cannot. Any estimator or routing logic that ignores the distinction produces guidance that is wrong precisely for the people waiting longest.

4.2 Queues at Rock in Rio: the public discourse, and what is different this year

The queue problem at this venue is documented, recent, and public. The 2024 edition — the festival's first at Parque Tejo — generated a 97% increase in complaints on Portugal's consumer-complaint platform relative to the 2022 edition, with queues (WCs prominently among them), access, overcrowding, and food service among the dominant categories (Portal da Queixa, 2024; Público, 2024). Attendees publicly reported waiting on the order of an hour for a WC, and press coverage identified the central WC strip — midway between the main stage and the principal food area — as the worst pressure point, while peripheral clusters moved faster (Observador, 2024; Público, 2024). The organisation responded mid-festival with structural and informational measures: WC provision at Parque Tejo already exceeded the previous venue by roughly 40% (≈490 units), and for the second weekend it added a dedicated team of ~30 people inside the WCs and ~25 outside, launched a WhatsApp channel pushing real-time information about which WCs and food areas had shorter queues, and applied a Stop&Go crowd-release protocol at the exits (Observador, 2024; Time Out Lisboa, 2024).

Two lessons follow, and they define what is different in 2026. First, capacity alone did not solve the problem: 2024 had more WCs than any previous Lisbon edition and still produced hour-long queues at specific clusters while others flowed — a textbook distribution failure, exactly the freedom-collapse pattern of Section 2 (paths existed; they were not perceived). Second, the 2024 organisation's own remedy — humans broadcasting “which WCs have shorter queues” over WhatsApp — is precisely a manual, unmeasured prototype of this system. It validated the demand for the information; it could not validate the information itself, because no measurement stood behind it. The 2026 deployment replaces broadcast guesswork with sensed, calibrated, confidence-scored cluster state, delivered continuously through the app and screens instead of episodically through a messaging channel. The difference this year is not more toilets; it is that, for the first time, the queues are measured.

4.3 Gendered waiting norms, unisex pooling, and its risks

The gender structure of the estate is a queueing decision with published mathematics behind it. Under equal-split gendered layouts, the queueing-theoretic study of toilet configurations by Van Hautegem and Rogiest (2017) computes women averaging over six minutes of waiting against seconds for men — a disparity reaching ≈40× in the waiting term — driven by two asymmetries: service time (a cabin visit takes roughly 1.5–2× a urinal visit) and capacity form (urinals pack far more service positions per metre). Mixed and unisex layouts change the mathematics through pooling: shared cabins act as one multi-server queue instead of two separate ones, and statistical multiplexing cuts women's expected waits from over six minutes to under a minute and a half in the Ghent model, with industry trials at Rock Werchter-class festivals reporting the same direction. This deployment's estate encodes both regimes — six gendered clusters where urinal capacity serves male peak flow efficiently, and two unisex clusters (WC-05, WC-06) where pooling equalises waits — and the section-aware estimator prices each correctly against its own limits (Table 6).

Unisex pooling also carries risks that a serious deployment names and designs against. Safety and harassment concerns rise in shared facilities, especially at night and under alcohol — the mitigations are fully enclosed floor-to-ceiling cabins, strong lighting, sight lines from circulation areas, and attendant presence at Peak. Comfort and cultural acceptance vary across the audience — clear signage and the continued availability of gendered alternatives preserve choice. Hygiene expectations differ around shared seats — tighter cleaning cycles on unisex blocks answer it, and the cleaning windows enter the capacity table as offline time (Section 19.3). Queue-discipline conflicts (urinal users bypassing the cabin line) need marked separate lanes where urinal pods coexist with cabins. And two crowd-safety hazards apply to any high-pressure cluster: counter-flow at entry-only geometries — WC-05's configuration, watched by its bidirectional sensing role and steward coverage — and surge crush at post-show waves, which the Post-concert regime exists to see early. The design position: pooling for efficiency, enclosure and staffing for safety, and measurement over assumption for both.

5. Problem Definition: WC Queues as Flow Friction

For the attendee, a WC queue is lost experience: minutes in line are minutes not at a show, with friends, or at concessions. For the operator, queues are a visibility problem: with 8 clusters across 145,000 m² and demand that shifts by thousands of people within minutes of a show ending, no human observer maintains an accurate picture of available capacity. For the economist, waiting is friction that suppresses spending and degrades satisfaction (Section 16 quantifies the assumption-based envelope).

Four structural asymmetries make the problem hard. Demand is spiky, arriving in programme-synchronised waves rather than smoothly. Demand is spatially uneven: clusters near stages saturate while remote clusters idle. Sections are asymmetric across male, female, and unisex provision. And information is asymmetric: an attendee sees only the queue in front of them, so they cannot trade a 20-minute wait for a 3-minute walk — unless a system tells them. The system objective follows: for every person, at any moment, estimate the pressure and availability of every WC cluster and recommend the fastest, most comfortable path to an available facility, using only anonymous, aggregate measurements of people and flows.

Without flow intelligence, the venue runs blind and friction compounds: people go to the nearest visible queue; congestion feeds on itself; underused clusters stay empty; staff react only after failure is visible; waiting becomes normalised; and the festival quietly loses value hour by hour.

With flow intelligence, the same infrastructure behaves differently: queues are detected earlier; attendees can be redirected; operators coordinate cleaning and stewarding by measured pressure; underused capacity gets used; time is returned to people; and the venue becomes less blind and more adaptive. Nothing physical changed — only visibility did.

The show is not the queue. Every minute recovered from a line is a minute returned to the reason people bought the ticket.

Figure 7 — Waiting as a tax on experience: queue time converts directly to lost experiential and economic value at venue scale.

6. Site Geometry and Geolocation Dataset

6.1 Coordinate system and venue polygon

All geospatial data uses WGS84 latitude/longitude. The venue boundary was captured as a four-vertex draft polygon — (38.794163, −9.090849), (38.792699, −9.096005), (38.789449, −9.094425), (38.787827, −9.091571) — treated throughout as an editable draft to be validated against, and where divergent replaced by, official GIS/site plans. Every figure displaying the polygon carries that caveat.

6.2 GPS limitations under festival conditions

Coordinates were collected under real festival conditions, and their error budget must be stated rather than assumed away. Sources of error include: temporary infrastructure (containers, stages, fencing) creating GNSS multipath; degradation of assisted positioning when mobile networks are saturated; Wi-Fi-positioning instability during high-density operation; manual capture error under time pressure; and occasional outlier readings with no obvious physical cause. The consequence is an architectural rule: raw sensor coordinates are research data, retained verbatim in the appendix with confidence grades; public WC points are canonical cluster markers generated by median aggregation, outlier detection, manual review, and official site-plan correction where available. Suspicious coordinates are flagged, never silently corrected.

6.3 Sensor coordinate dataset

Table 3 lists the 24 geotagged sensor seed points with cluster, section type, role, coordinates, a three-level coordinate-confidence grade (high / medium / low), and data-quality notes. Naming: H = male section, M = female section, U = unisex; suffix 1 = entrance role, suffixes 2/3 = interior role.

sensor_id	wc_cluster	section	role	latitude	longitude	confidence	notes
WC-01-H1	WC-01	male	entrance	38.793900	-9.094434	medium	Nearly identical to WC-01-H2 (≈0.5 m); pair treated as one effective point
WC-01-H2	WC-01	male	interior	38.793901	-9.094439	medium	Nearly identical to WC-01-H1

sensor_id	wc_cluster	section	role	latitude	longitude	confidence	notes
WC-02-H1	WC-02	male	entrance	38.792676	-9.094007	high	—
WC-02-H2	WC-02	male	interior	38.792676	-9.094199	high	—
WC-02-M1	WC-02	female	entrance	38.792133	-9.095900	high	—
WC-02-M2	WC-02	female	interior	38.791958	-9.095163	high	—
WC-03-H1	WC-03	male	entrance	38.791614	-9.095144	high	—
WC-03-H2	WC-03	male	interior	38.791643	-9.095504	high	—
WC-03-M1	WC-03	female	entrance	38.791665	-9.095432	medium	Very close to WC-03-M3; averaged, never summed
WC-03-M3	WC-03	female	interior	38.791711	-9.095434	medium	Very close to WC-03-M1; averaged, never summed
WC-04-H1	WC-04	male	entrance	38.790413	-9.096673	high	—
WC-04-H2	WC-04	male	interior	38.790111	-9.098251	low	OUTLIER — far west of other WC-04 sensors; flagged, excluded from canonical point
WC-04-M1	WC-04	female	entrance	38.790352	-9.096969	high	—
WC-04-M2	WC-04	female	interior	38.790346	-9.096851	high	—
WC-05-U2	WC-05	unisex	interior	38.789946	-9.094806	medium	Only interior point exists; entrance point missing — cluster confidence capped
WC-06-U1	WC-06	unisex	entrance	38.791227	-9.090976	high	—
WC-06-U2	WC-06	unisex	interior	38.791330	-9.090771	high	—
WC-07-H1	WC-07	male	entrance	38.792245	-9.091415	high	—
WC-07-H2	WC-07	male	interior	38.792280	-9.091096	high	—
WC-07-M1	WC-07	female	entrance	38.792147	-9.091450	high	—
WC-07-M2	WC-07	female	interior	38.791572	-9.095257	low	OUTLIER — recorded position falls near WC-03 (~340 m from assigned cluster); flagged, excluded
WC-08-H1	WC-08	male	entrance	38.793577	-9.091463	high	—
WC-08-M1	WC-08	female	entrance	38.793816	-9.091349	medium	Very close to WC-08-M2
WC-08-M2	WC-08	female	interior	38.793847	-9.091293	medium	Very close to WC-08-M1

Table 3 — Sensor coordinate dataset (n = 24), WGS84, with confidence grades. Raw values retained verbatim; see Appendix A.

A distinction runs through everything that follows: these 24 points are geotagged research seed points — physically verified positions constituting the paper's spatial ground truth. The operational backend maintains a separate configured node registry that may include inactive, diagnostic, test, or historical nodes; backend counts need not equal seed-point counts, and the two inventories are never conflated. Coverage is deliberately reported as unequal: WC-05 contributes a single interior point (entrance missing — cluster confidence capped), and WC-01's two points are nearly co-located, effectively one spatial observation.

Cluster	Seed points	Section logic	Coverage note
WC-01	2	male / female	pair nearly co-located — one effective point
WC-02	4	male / female	full reference topology
WC-03	4	male / female	full reference topology; M1/M3 averaged
WC-04	4	male / female	H2 flagged outlier
WC-05	1	unisex	entrance point missing — confidence capped
WC-06	2	unisex	full unisex reference topology
WC-07	4	male / female	M2 flagged outlier

Cluster	Seed points	Section logic	Coverage note
WC-08	3	male / female	M1/M2 averaged
Total	24	—	—

Table 4 — Seed-point distribution per cluster.

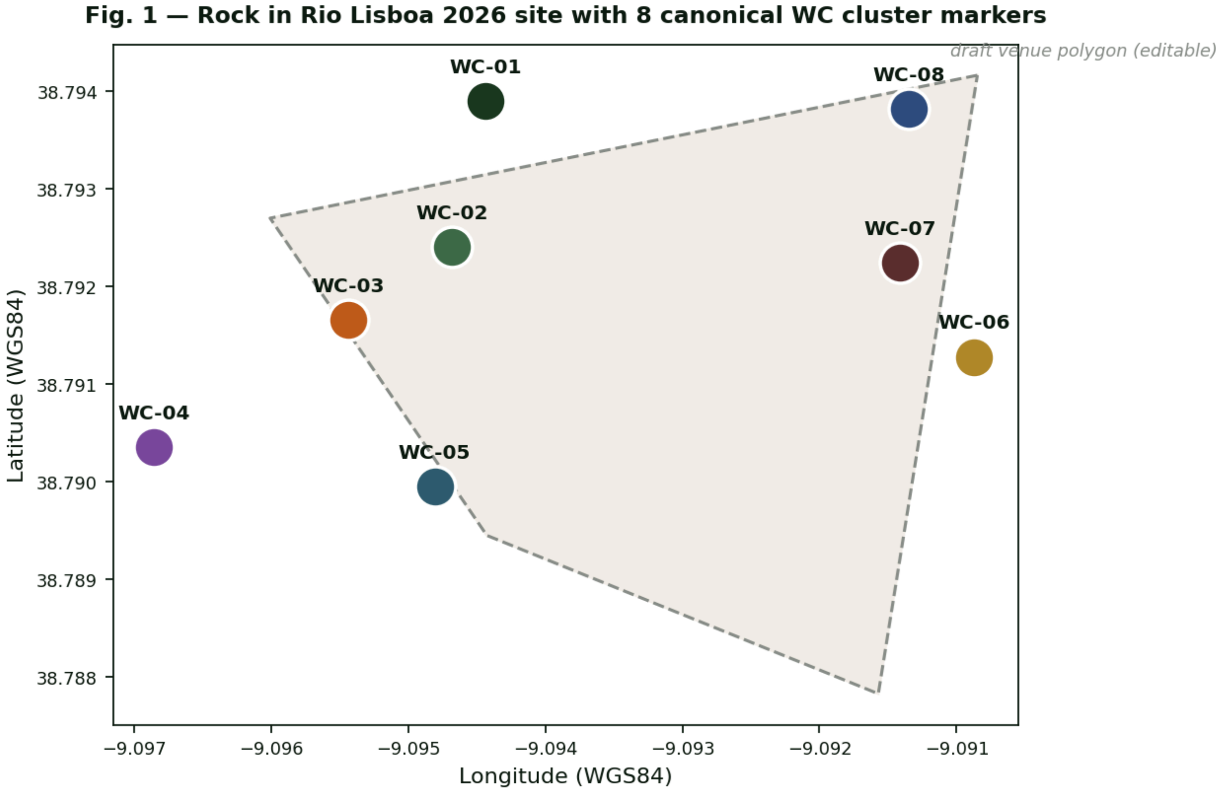

Figure 8 — Site overview: one canonical marker per WC cluster inside the draft venue polygon.

Fig. 2 — 24 geotagged sensor seed points by cluster and role

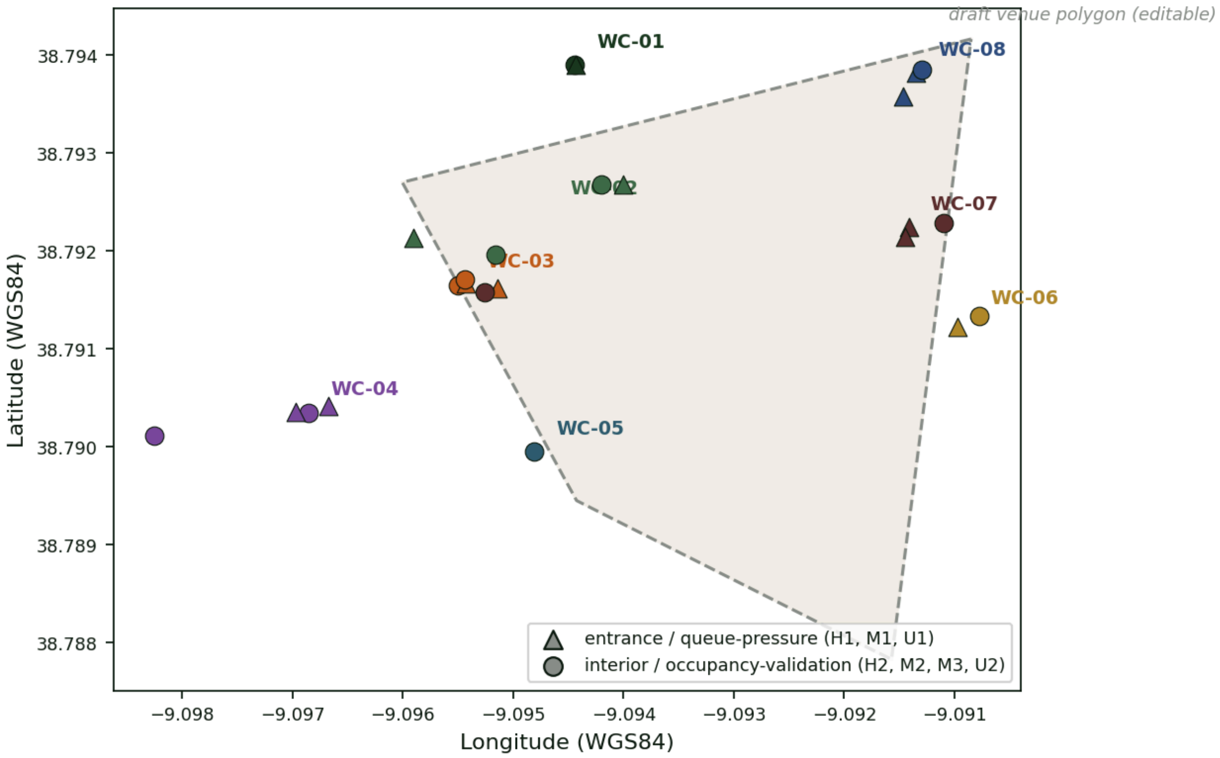

Figure 9 — All 24 sensor seed points by cluster and role (triangles = entrance, circles = interior).

6.4 Canonical WC marker generation

Public maps must show one canonical marker per cluster, not one marker per sensor: attendees navigate to facilities, not instruments, and a scatter of near-duplicate pins degrades trust. The generation method: (1) group sensors by cluster; (2) detect outliers by distance-to-cluster-median combined with expected-topology comparison (a sensor whose position contradicts its assigned cluster is suspect); (3) flag rather than silently correct — flagged points are excluded from the canonical computation but preserved in the dataset; (4) compute the median latitude/longitude of remaining points (median, not mean, so one residual outlier cannot drag the marker); (5) retain manual override so operations staff can pin the marker to the true entrance; (6) publish per-cluster confidence and warnings alongside the marker.

Fig. 4 — Canonical WC marker generation: one public marker per cluster, not per sensor

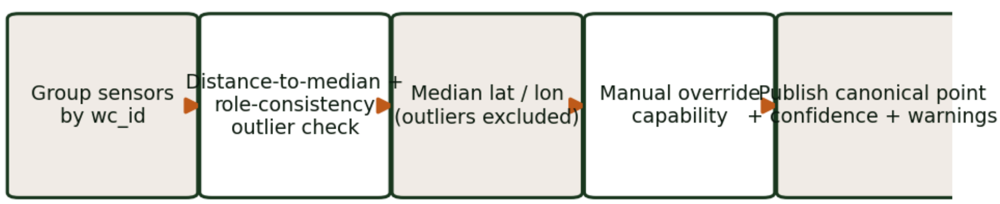

Figure 10 — Canonical marker generation pipeline.

Applying the method surfaces the known issues, all retained with flags: WC-01-H1/H2 nearly identical; WC-08-M1/M2 very close (averaged); WC-03-M1/M3 very close (averaged); WC-07-M2 recorded ~340 m from its assigned cluster, near WC-03 — flagged and excluded pending physical re-verification; WC-04-H2 displaced far west of its cluster — flagged and excluded; WC-05 lacking an entrance point; and the draft polygon pending official GIS.

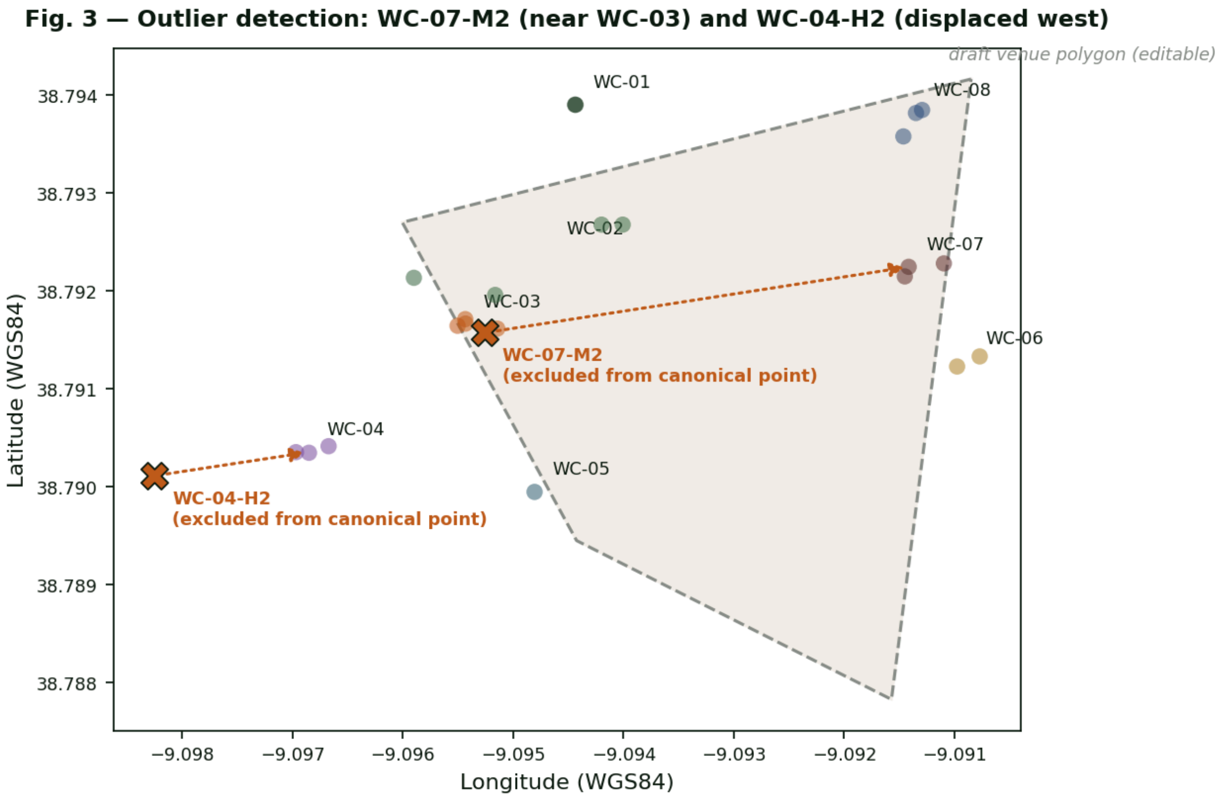

Figure 11 — GPS uncertainty and outlier handling: WC-07-M2 and WC-04-H2 flagged and excluded from canonical points (dotted arrows show assigned clusters).

7. Hardware and Sensing Architecture

7.1 Node platform

Sensor nodes are built on LilyGo LoRa32 boards: an ESP32-class microcontroller with integrated Wi-Fi/Bluetooth radio and a LoRa transceiver (868 MHz in the EU; 433/915 MHz variants exist for other regions). The ESP32 performs local aggregate radio-presence sensing — counting nearby active devices within a calibrated radius, never storing identifiers — while the LoRa radio provides low-power, long-range backhaul for small data reporting payloads. Unit cost is on the order of tens of euros per node; power is via USB, Li-Po, or commodity power banks; the platform is explicitly suited to temporary deployment: configured, mounted, and reporting within minutes, and cheap enough to replace on damage. The link is low-bandwidth but robust for small data reporting payloads; it is not suitable for video or any high-bandwidth data — a constraint that is also a privacy feature, since the pipe physically cannot carry imagery.

7.2 Communication range is not sensing radius

Two radii govern a node and they differ by roughly two orders of magnitude. LoRa communication range — how far data reporting travels to a gateway — can reach kilometres outdoors but depends on antenna quality, gateway placement, line of sight, interference, packet rate, spreading factor, duty-cycle limits, and the built environment; at a festival, temporary metal structures and 80,000 attenuating bodies compress it substantially. Effective sensing radius — the zone whose devices a node should count — must be far smaller and is calibrated per sensor role: entrance positions use a deliberately tighter radius than interior positions, because an entrance sensor that hears half the concourse converts passers-by into phantom queue. Conflating the two radii is a common deployment error; this paper keeps them strictly separate.

The sensing radius was calibratable remotely, gated by received-signal-strength (RSSI) thresholds. Table 5 gives the six operational sensitivity levels implemented in the node firmware.

Radius shown in page	Sensitivity	What it captures
5 m	−65 dBm	only phones very close — door-level sensing
10 m	−75 dBm	nearby phones
15 m	−82 dBm	medium range — door + entrance
20 m	−86 dBm	far — whole WC area
25 m	−90 dBm	very far
>25 m up to 60 m	−93 dBm	captures almost everything around

Table 5 — The six RSSI/radius sensitivity levels implemented in the LilyGo LoRa32 node firmware (remotely configurable).

The technical range spans approximately 3 m to 60 m, but the operationally useful range depends on the sensor role. Recommended logic: entrance sensors should usually use tighter ranges (5, 10, or 15 m); interior sensors can use 15 or 20 m depending on WC size; 25 m and above can overcount surrounding movement and should be used carefully; the large unisex clusters WC-05 and WC-06 may require wider sensing, but at the cost of lower confidence and stronger calibration. The failure modes are symmetric: high sensitivity creates false occupancy by capturing people passing nearby; low sensitivity misses real queue pressure.

The system became more useful not by maximizing sensor range, but by limiting range to the correct physical meaning.

7.3 Why private radio

Public mobile networks at 80,000-person density congest to effective unusability at exactly the moments data reporting matters most — show end, exit surge. Private low-power radio backhaul removes that dependency: the data reporting path is owned end-to-end, and its capacity is engineered for the payloads it carries. Power planning treats each festival day as an independent mission (banks swapped or recharged overnight, battery state tracked as node metadata); enclosures are non-metallic (ABS-class) because metallic housings attenuate the very signals the node measures; mounting is at 2–3 m on cluster structures, oriented toward the sensed zone.

7.4 Sensing topology

The reference topology is simple and strict. Each gender-separated cluster carries four sensing roles: one entrance sensor and one interior sensor per section (male: H1 entrance, H2 interior; female: M1 entrance, M2/M3 interior). Each unisex cluster carries two roles: one entrance (U1), one interior (U2). Entrance sensors estimate approach flow, queue pressure, and queue formation — the leading edge of a demand wave is visible at the entrance minutes before it converts into occupancy. Interior sensors validate actual occupancy and prevent passers-by from being booked as occupants. Physically close sensors contribute through medians and averages weighted by freshness and confidence — never blind summation, which double-counts the same devices.

Fig. — Sensor topology: gender-separated vs unisex reference design

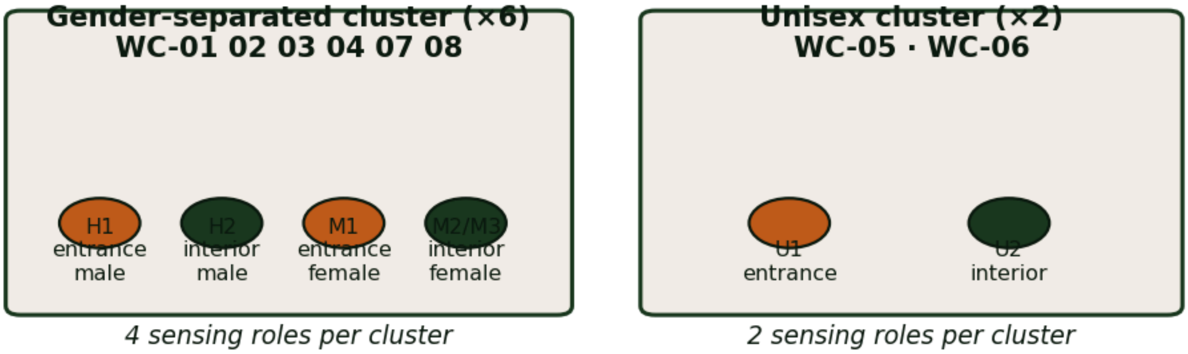

Reference design: 6 × 4 roles + 2 × 2 roles = 28 sensing roles across the estate.

Entrance roles (amber) estimate approach flow, queue pressure, and queue formation.
Interior roles (green) validate actual occupancy. Nearby sensors are averaged, never blindly summed.

Figure 12 — Sensing topology: four roles per gender-separated cluster, two per unisex cluster; 28 sensing roles across the estate.

Fig. — LilyGo LoRa32 / ESP32 node: edge sensing, LoRa backhaul, power-bank operation

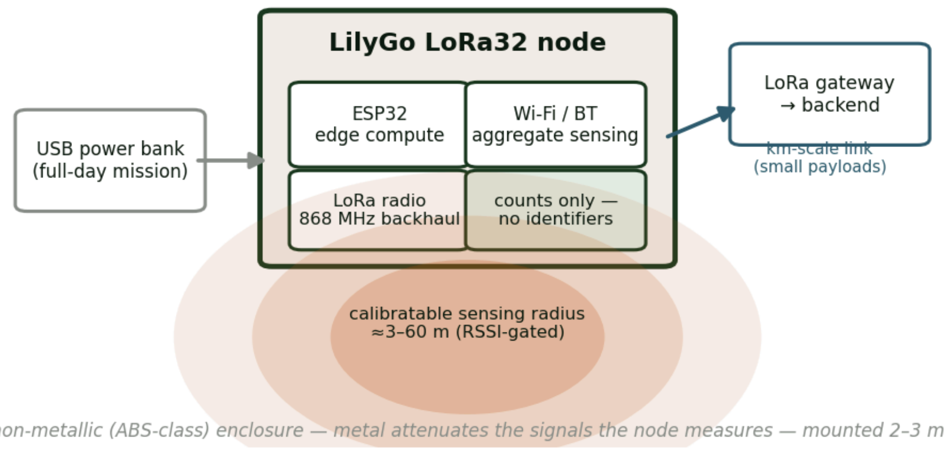

Figure 13 — LilyGo LoRa32 / ESP32 node architecture.

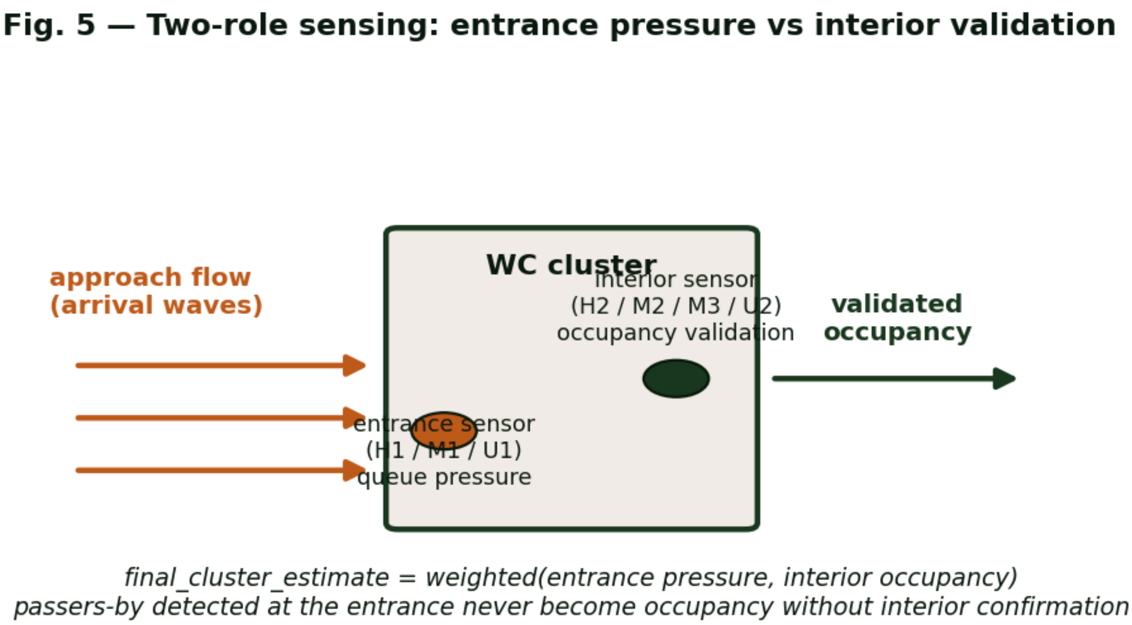

Figure 14 — Entrance vs interior sensing: pressure estimation validated by occupancy confirmation.

Section 6's 24 verified seed points document the field record against this 28-role reference design, with the gaps stated (WC-05 entrance missing; WC-01 pair co-located). Where this paper states deployed quantities, it means verified seed points; where it states per-cluster patterns, it means the reference design. The two are kept separate to avoid overclaiming field coverage.

Single-sensor mode for compact clusters. The two-role geometry is the reference design; a compact section can legitimately run on one sensor. When a cluster's physical depth is short — door to back wall within roughly 10–15 m — a single well-placed door-level node at a tight radius (5–10 m, −65 to −75 dBm) covers queue, doorway, and interior inside one physically meaningful footprint: the entrance/interior distinction collapses because both live within the sensor's resolution. WC-05 operated in exactly this mode (only its U2 role present), and short single-section blocks are the natural candidates. The trade is explicit and priced by the confidence engine: single-sensor mode preserves occupancy estimation and cross-cluster ranking, while it surrenders the entrance lead time — the wave is seen as it arrives rather than minutes early — so its confidence is capped and its weight in venue-level decisions reduced accordingly. The operational rule: two roles wherever queue prediction matters; one role suffices where the section is short, the flow is simple, and ranking is the requirement.

7.5 How the node works, and where it goes

How it works. The LilyGo LoRa32 is an ESP32-class board whose radio listens passively for the Wi-Fi and Bluetooth frames every phone emits routinely; the node transmits nothing toward people. The firmware applies an RSSI gate — six real sensitivity levels are implemented in the node code, from −65 dBm (≈5 m, door-level only) to −93 dBm (beyond 25 m, up to ≈60 m, everything around) — counts the distinct devices inside the gate, discards identifiers on the node itself, and uplinks the bare count over the private LoRa backhaul on the 868 MHz band. Remote calibration is the closing of this loop: the backend pushes a radius level and the node applies the matching threshold within seconds — which is what made negative calibration executable at distance, during show hours, without touching a single enclosure. The effective operating envelope is 3–60 m; the six levels are the operational vocabulary of Table 5.

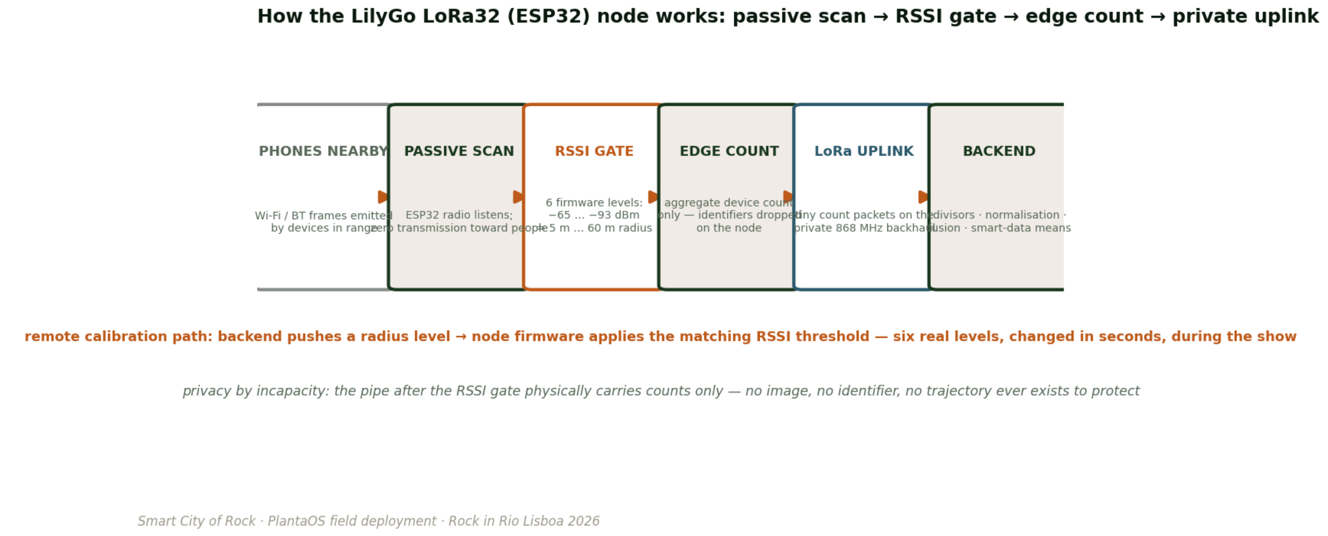

Figure 15 — How the node works: passive scan → six-level RSSI gate → edge count → private LoRa uplink, with the remote-calibration path.

Where it goes. Placement follows the two-role geometry. The entrance node mounts on or beside the door frame at 2–3 m — above heads, below structures — at a tight radius (5–10 m, −65/−75 dBm) so that its footprint means the queue and the approach. The interior node mounts on the back wall at the same height, at a wider radius (10–20 m, −75/−86 dBm) sized to the section's physical depth, so that its footprint means real occupancy. Enclosures are ventilated ABS — radio-transparent and heat-safe — mounted clear of metal sheeting and cable trays, with the power bank inside the enclosure on the 10-hour swap cycle. Physically close nodes enter the estimator through averages or medians; compact clusters run single-sensor mode (Section 7.4). The placement rule compresses to one sentence: put the sensor where its radius has exactly one physical meaning.

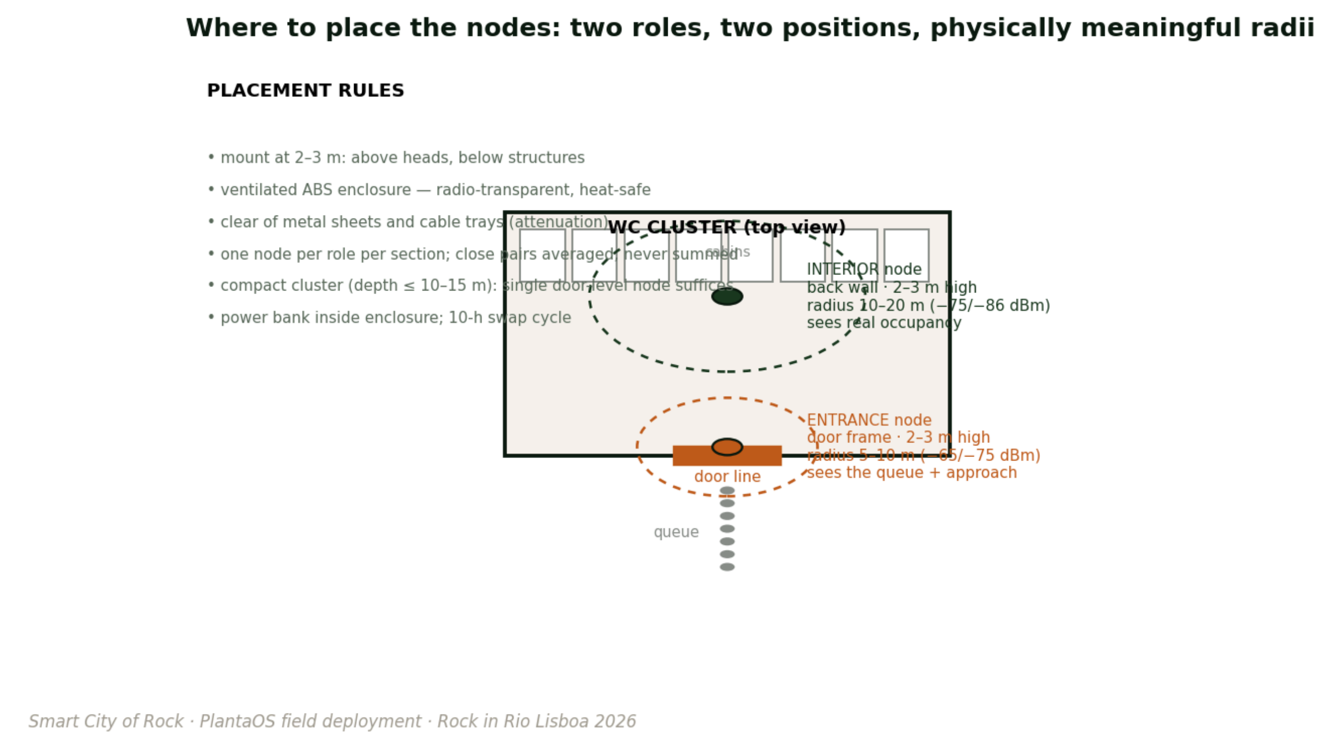

Figure 16 — Where the nodes go: entrance at the door frame, interior at the back wall, radii sized to physical meaning.

7.6 Why these sensors, and why the alternatives lose

The virtue of the LilyGo/ESP32 node is not any single specification; it is the fit between its properties and the problem's constraints. It is cheap enough to deploy densely and replace casually — a damaged node is an errand, not an incident. It is fast enough to deploy that the T−7 window suffices. It is privacy-preserving by incapacity, not merely by policy: the LoRa pipe physically cannot carry imagery, and the node computes counts at the edge. Its sensing radius is remotely calibratable in six steps (Table 5), which is what made negative calibration executable at distance during show hours. Its backhaul is private, so it kept reporting precisely when the public networks drowned. And it fails gracefully: a dead node produces an honest offline state, not a wrong number. No individually superior sensor — a better camera, a better turnstile — offers that combination under temporary-venue constraints; Section 15.1 makes the underlying argument that this combination, approximate but fresh and honest, is what queue routing actually needs.

7.7 Electrical and operational safety

The safety profile of the deployment is deliberately minimal. Nodes run on standard USB power banks at 5 V — no mains wiring at or near the WC clusters, no trailing cables through circulation areas, no work at height beyond the 2–3 m mounting. Power banks were replaced on a rolling cycle of approximately every 10 hours, sized so that each swap occurred well before depletion; battery state was tracked as node metadata, and swaps were a two-minute steward task. Enclosures are ventilated ABS-class plastic — chosen for radio transparency, with the side effect that there is no sealed metal box heating in the sun. The nodes' radio activity is standard low-power Wi-Fi/Bluetooth reception plus brief LoRa data reporting bursts, orders of magnitude below the emissions of the phones in every attendee's pocket. Across the deployment, no material safety risk was identified: low voltage, no mains, no heat concentration, no moving parts, no physical interaction with attendees. The honest formulation is not “risk-free” — nothing physical is — but that the residual risk is of the same class as any battery-powered consumer device attached to street furniture.

8. Methodology Overview

The end-to-end methodology comprises eighteen stages: (1) site mapping; (2) WC cluster identification; (3) sensor topology design; (4) sensor deployment; (5) coordinate capture; (6) canonical WC marker generation; (7) data ingestion; (8) calibration; (9) negative calibration; (10) five-minute rolling averages; (11) confidence scoring; (12) offline/failure handling; (13) SCOR data reporting export; (14) public screen/app recommendations; (15) privacy/GDPR controls; (16) economic value modelling; (17) counterfactual comparison; and (18) lessons-learned synthesis. Stages 1–6 are covered in Sections 6–7; stages 7–14 in Sections 9–11 (with the wave model of Section 10); stage 15 in Section 12; stages 16–18 in Sections 15–21. Every stage is designed to be inspectable: a value on a public screen can be traced back through cluster state and calibration to the individual samples that produced it.

9. Calibration Method

9.1 Why raw counts cannot be published

Raw device counts are not people counts. They are distorted by detection radius (a node hears devices outside the zone it should observe), received-signal-strength variability, field overlap between nearby sensors, crowd density itself (bodies attenuate signal), passers-by registering at entrances without joining queues, placement compromises imposed by mounting constraints, devices carried by the same person, nearby queues and adjacent circulation, overlapping sensor fields, high-density signal noise, GPS/geolocation uncertainty, offline or stale nodes, and time-of-day flow conditions. A system that publishes raw counts as occupancy is wrong in ways obvious to anyone standing at the facility — which destroys public trust immediately and permanently.

9.2 Estimation model: from counts to operational pressure

The system was never sensor-detects-phones → app-shows-bathroom. The actual chain is: entrance signal + interior signal + capacity limits + time wave + confidence + calibration divisor → cluster pressure → best-WC recommendation. This subsection specifies that chain precisely, because its central design choice — entrance-dominant weighting, frequently around 70% entrance / 30% interior — is easy to misread as arbitrary and is anything but.

Occupancy versus pressure. For a WC cluster there are two different questions. Question 1 — how full is the WC? — is answered primarily by the interior sensor. Question 2 — is a queue forming, or about to form? — is answered primarily by the entrance sensor. At Rock in Rio the operationally decisive question was almost always the second: where is pressure forming before it becomes a visible queue? A high-density festival does not fail only when a WC is physically full; it fails when demand forms a queue faster than the system can redistribute people. That is why the fusion is usually entrance-dominant: the entrance detects the wave before the interior becomes full. The entrance sensor reports arrivals, queue formation, wave pressure, post-concert surges, pass-through intensity, and rising demand before saturation; the interior sensor reports actual occupancy, whether the WC is really being used, whether entrance pressure is real or merely passers-by, and whether the entrance is overcounting. ***Entrance predicts the wave. Interior validates the wave.***

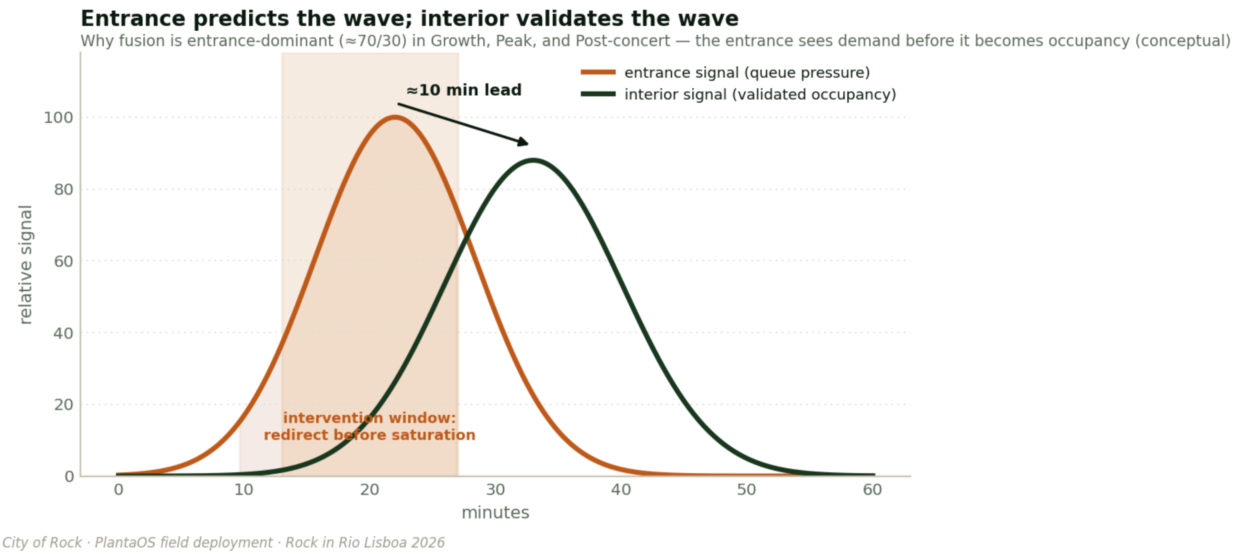

Figure 17 — Entrance predicts the wave; interior validates the wave: the lead time is the intervention window (conceptual).

Capacity-adjusted formulation. Raw counts are never fused directly, because a count of 80 people means completely different things in different clusters — WC-03 at 80 people is near pressure, while WC-06 at 80 people is still comfortable. Each calibrated signal is therefore first normalised against its section-specific limit:

$$E\_norm(s,t) = \text{calibrated entrance estimate / waiting-adjusted section limit}$$

$$I\_norm(s,t) = \text{calibrated interior estimate / physical section capacity}$$

and the section pressure is the regime-weighted fusion of the normalised signals:

$$P(s,t) = wE(t) \cdot E\_norm(s,t) + wI(t) \cdot I\_norm(s,t), \quad wE(t) + wI(t) = 1$$

$$\text{Pressure score} = \min(\,100,\, 100 \cdot P(s,t)\,)$$

The published quantity is thus a pressure score, not a raw people count, banded operationally as: 0–40 low; 40–70 medium; 70–90 high; 90+ critical. The waiting-adjusted section limits instantiating E_norm derive from the official occupancy limits of Table 2 (physical capacity plus proportional waiting allowance) and are given in Table 6.

WC	Male limit	Female limit	Unisex limit
WC-01	115.2	100.8	—
WC-02	97.2	129.6	—
WC-03	97.2	86.4	—
WC-04	151.2	118.8	—
WC-05	—	—	239.4
WC-06	—	—	374.4
WC-07	151.2	97.2	—
WC-08	151.2	109.8	—

Table 6 — Waiting-adjusted section limits used for E_norm (physical capacity plus proportional waiting allowance; derived from Table 2).

Regime-dependent weights. The weighting is not fixed across the day; it follows the wave regimes, and the frequently used 70/30 split is highly defensible in Growth, Peak, and Post-concert — while for a pure occupancy question, outside queue-risk hours, the weighting may be reversed. Table 7 gives the operating bands.

Regime	Entrance weight	Interior weight	Operational purpose
Light	30–40%	60–70%	establish baseline and avoid false positives
Growth	60–70%	30–40%	detect queue formation — the first wave
Peak	70–80%	20–30%	prevent queue collapse; queue pressure is the main risk
Post-concert	70–85%	15–30%	detect sudden demand waves arriving fast

Table 7 — Regime-dependent fusion weights (entrance-dominant, e.g. 70/30, in Growth/Peak/Post-concert; interior-dominant in Light).

Aggregation by cluster type. For gender-separated clusters, the fusion runs per section — male pressure = wE·(male entrance)_norm + wI·(male interior)_norm, and identically for the female section — and the cluster-level figure combines the sections. For unisex clusters, a single U1/U2 pair feeds the same formula. For public display, cluster pressure = 70% × worst section + 30% × average section: if the female section is saturated, the cluster is operationally problematic even if the male side is fine, and a plain average would hide exactly the failure the system exists to expose.

Every estimate carries a confidence score, computed per role:

$$C(s,t) = f(\,\text{freshness, completeness, entrance–interior agreement, stability, sensor health, geolocation confidence, time regime}\,)$$

Confidence drops when entrance and interior sensors disagree strongly (one of them is measuring something other than the queue), when samples are stale, and when a contributing sensor's geolocation is suspect (Table 3's low-confidence rows).

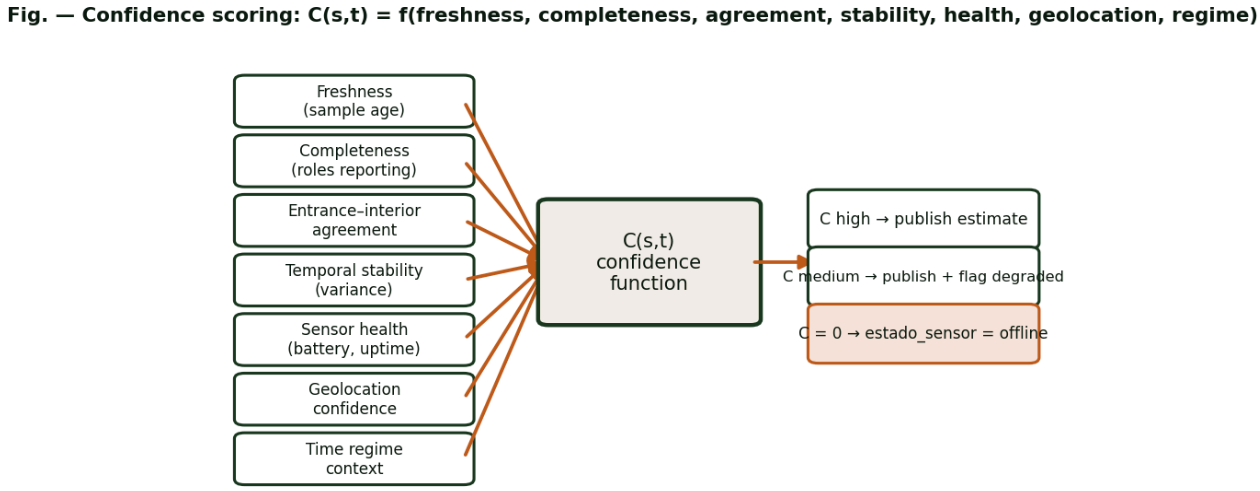

Figure 18 — Confidence scoring logic: seven inputs, one score, three publication outcomes.

Confidence-weighted fusion. The highest-quality version is not fixed 70/30; it is 70/30 adjusted by confidence. Each base weight is multiplied by its sensor's confidence and the pair renormalised:

$$w^*E = wE \cdot CE \,/\, (\, wE{\cdot}CE + wI{\cdot}CI \,) \qquad w^*I = wI \cdot CI \,/\, (\, wE{\cdot}CE + wI{\cdot}CI \,)$$

where CE and CI are the entrance and interior confidence scores from freshness, stability, sensor health, geolocation confidence, and inter-signal agreement (Section 9.1's C(s,t) applied per role). The consequence is automatic graceful degradation: if the entrance sensor is offline or noisy, its effective weight collapses and the system leans on the interior sensor — without any operator action, and without ever fabricating the missing signal. This is what made the fusion more robust than any simple average.

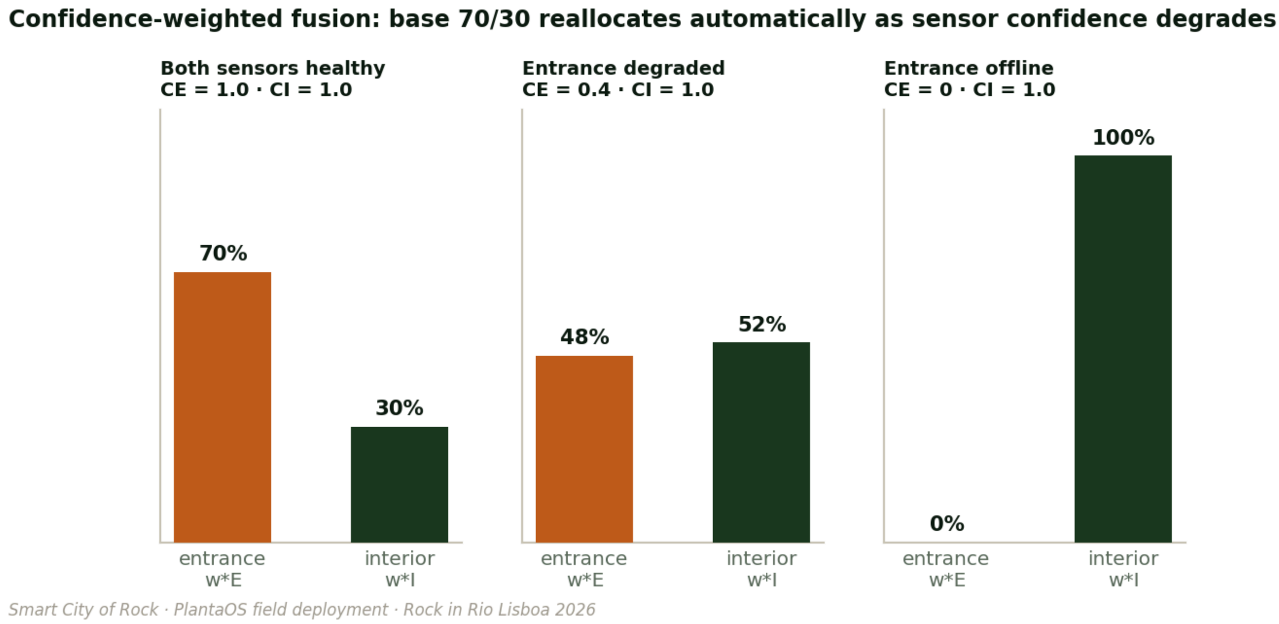

Figure 19 — Confidence-weighted fusion: the base 70/30 reallocates automatically as sensor confidence degrades.

The smart-data layer: two-stage stabilisation. Stabilisation runs in two stages, at every level of the hierarchy. Stage one — the smart-data layer: every published quantity is built from the average of the last 60 seconds of raw values, computed per sensor, per gender section, per cluster, and at the all-clusters estate level:

$$P1(x,t) = mean(\ raw(x,\ t{-}60s\ ..\ t)\ ) \quad for\ x \in \{\ sensor,\ section,\ cluster,\ estate\ \}$$

The one-minute mean absorbs detection flicker — the second-to-second churn of devices crossing the radius boundary — while remaining fast enough to track a wave front. Stage two: the five-minute rolling mean P5(s,t) of the one-minute series stabilises the externally published pressure values, and Q(s,t) = max(0, P5 − threshold) remains the operational alert trigger. Raw values move every second; the smart-data layer moves every minute; the public surfaces move on the five-minute rhythm — each level exactly as fast as its decision requires. Physically close sensors (the WC-01 pair, WC-08-M1/M2, WC-03-M1/M3) enter through averages or medians rather than sums.

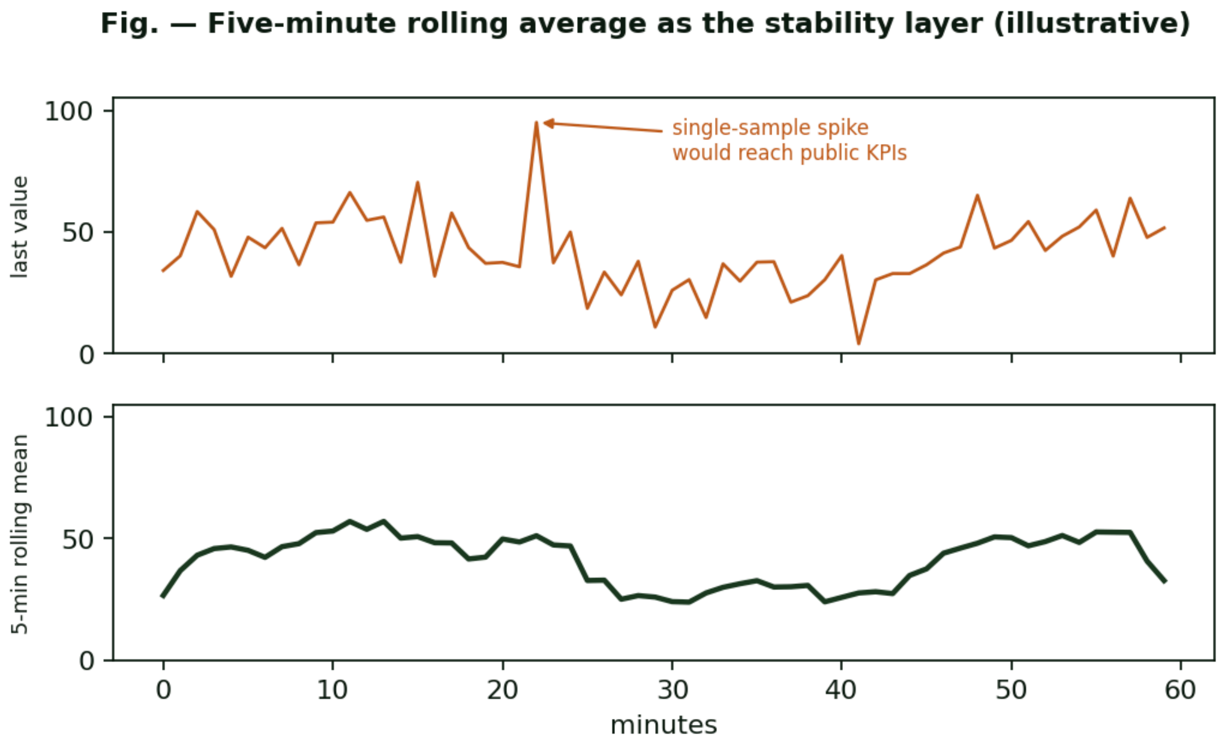

Figure 20 — The five-minute rolling average as stability layer: spikes and dropouts never reach published values (illustrative).

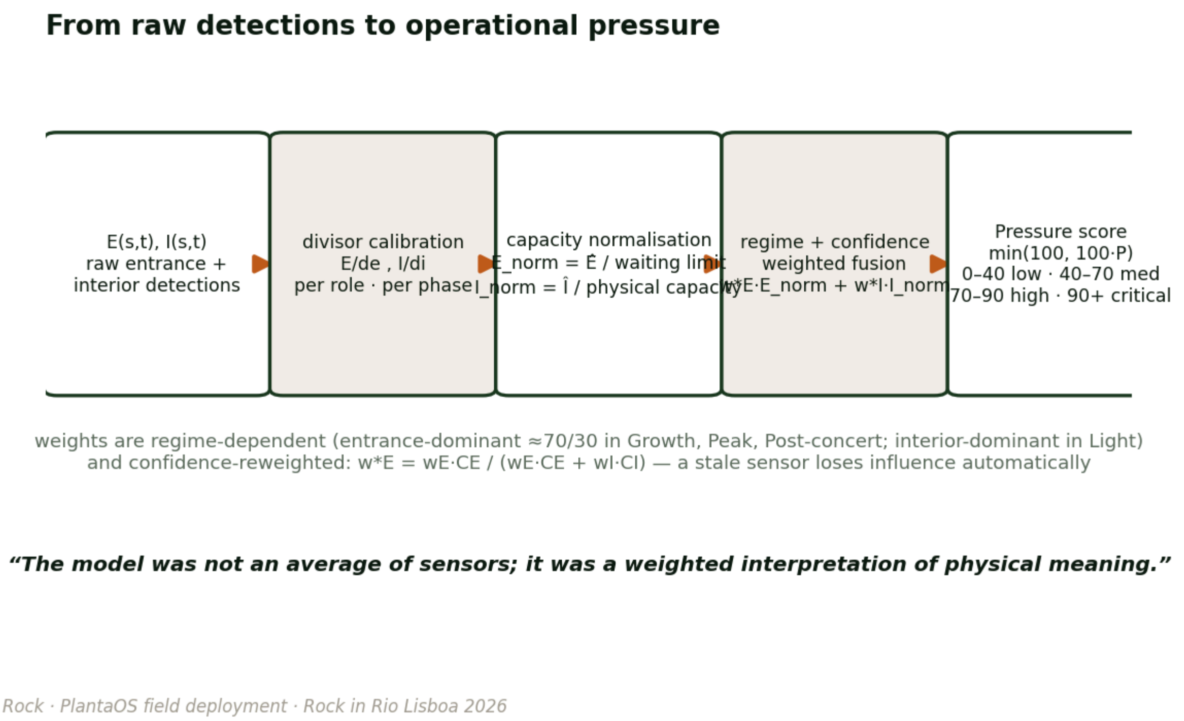

Figure 21 — The full estimation chain: detections → divisor calibration → capacity normalisation → regime + confidence weighted fusion → pressure score.

The model was not an average of sensors; it was a weighted interpretation of physical meaning. Entrance sensors predicted queue waves, interior sensors validated real occupancy, and capacity limits converted both into operational pressure.

9.3 Negative calibration

The first real tests produced a finding that shaped the entire deployment: raw counts fed through the initial formulas systematically overestimated people. Divisors set optimistically low, radii set generously wide, and stale sensors left in the fusion all pushed estimates upward, and the published occupancy was visibly inconsistent with conditions at the clusters. This was not a failure of the idea; it was a discovery of the real physics of the venue — reflective temporary structures, adjacent circulation paths, and dense crowds produce biased raw signals everywhere, and every deployment plan must budget authority and time for on-site downward correction.

Negative calibration comprised: increasing divisors; reducing the fusion weights of noisy sensors; tightening effective sensing radius; lowering published confidence when inputs degraded; excluding stale sensors from fusion; preventing old values from ever being published as current; and routing all publication through five-minute averages instead of unstable last values. The final field configuration converged on a uniformly tightened radius and increased divisors with automatic recalibration disabled, so behaviour was deterministic during show hours.

Table 8 reconstructs the calibration trajectory. Exact per-node historical logs were not retained for every change; the table is therefore labelled as reconstructed from field calibration notes and reports the configuration classes, not per-message history. The final field configuration converged on a uniform tightened radius (≈6 m) with increased divisors and automatic recalibration disabled, so behaviour was deterministic during show hours.

Sensor / cluster	Role	Initial range	Final range	Final sens.	Initial divisor	Final divisor	Reason for change	Effect on estimate	Confidence impact
Entrance nodes (H1/M1/U1), all clusters	entrance	~15 m	~6 m	≈−67 dBm	2–3	10	Passers-by and adjacent circulation captured at door	Overestimation removed; counts match queue zone	Raised — signal matches ph meaning
Interior nodes (H2/M2/M3/U2), standard clusters	interior	~20 m	~6 m	≈−67 dBm	2–3	10	Field spilled beyond WC envelope; overlap with entrance field	Occupancy no longer inflated by exterior devices	Raised — entrance/interior agreement improved
WC-05 / WC-06 (large unisex)	interior	20–25 m	~6 m	≈−67 dBm	2–3	10	Wide fields overcounted surrounding movement in open areas	Stable, conservative estimates for the largest clusters	Raised, with explicit cap (WC missing entrance)
Stale / low-freshness nodes	any	(in fusion)	excluded	—	—	—	Old samples propagated as current	Phantom occupancy eliminated	estado_sensor = offline, con = 0
Auto-recalibration	system	enabled	disabled	—	—	—	Unsupervised drift during show hours	Deterministic behaviour at Peak	Predictable, auditable confid

Table 8 — Calibration before/after (reconstructed from field calibration notes; configuration classes, not per-message logs).

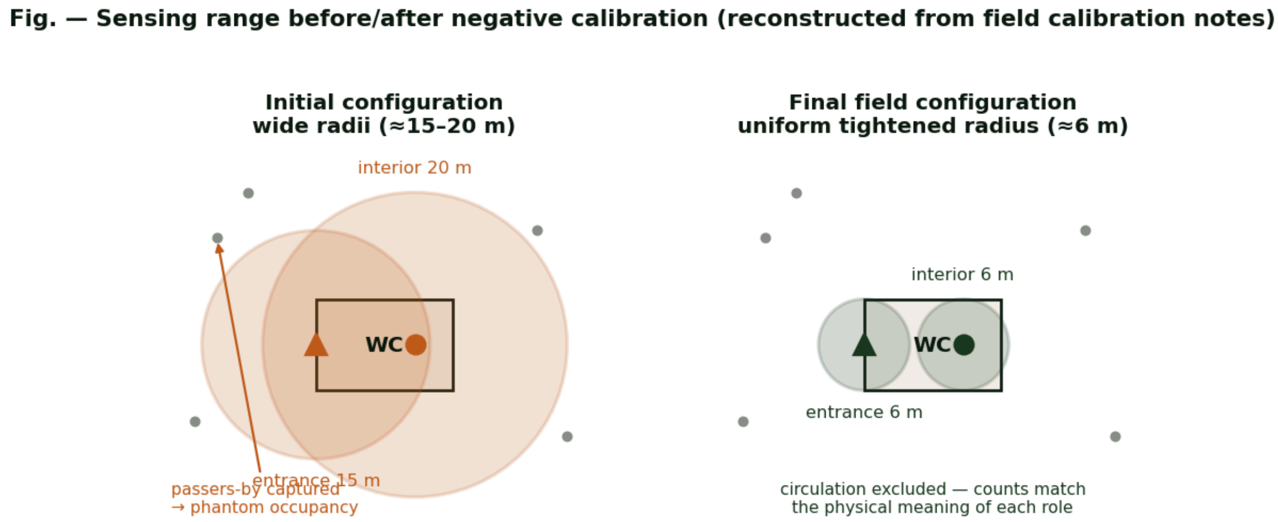

Figure 22 — Sensing range before/after negative calibration (reconstructed from field calibration notes).

The system became intelligent not by measuring more, but by learning when to believe each signal less.

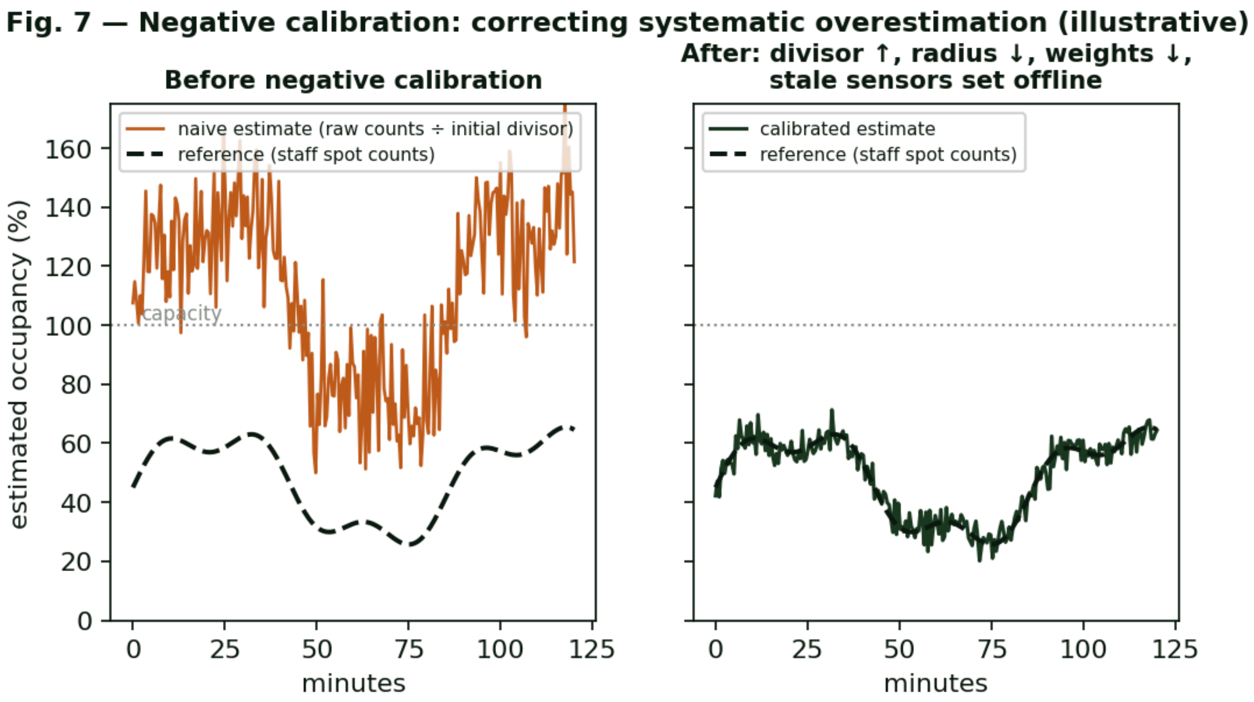

Figure 23 — Negative calibration (illustrative reconstruction): naive estimates overshoot and destabilise; after divisor increases, radius tightening, and offlining of stale sensors, estimates track staff spot counts.

9.4 Four temporal regimes

Festival demand is regime-switching; one static calibration cannot serve the whole day. Four regimes structure the time-dependent weights: Light (13:00–16:30), Growth (16:30–18:30), Peak (18:30–21:30), and Post-concert (21:30–close). Entrance weight we(t) rises during Peak and Post-concert, because demand arrives in synchronised waves whose leading edge is visible at entrances minutes before converting into occupancy — that lead time is the system's operational value. Interior weight wi(t) never falls to zero in any regime: interior validation is what prevents passers-by from becoming false occupancy, and strong entrance–interior disagreement is itself a confidence penalty. Operationally: during Light, use conservative calibration; during Growth, increase attention to entrance sensors; during Peak, entrance+interior fusion is essential; during Post-concert, freshness is critical because waves move quickly.

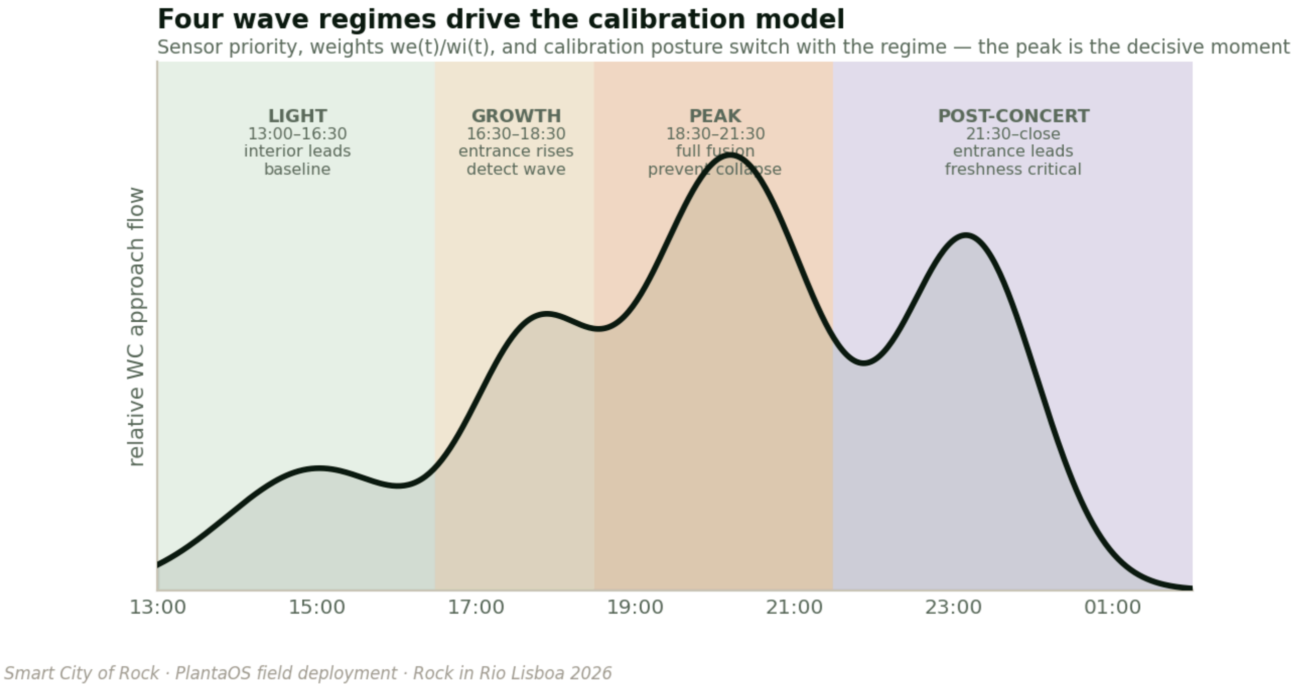

Figure 24 — Four temporal regimes over a festival day; entrance weight increases during Peak and Post-concert.

9.5 Phase-dependent divisors and fusion weights

Calibration was per time, per cluster, and per radius simultaneously — three orthogonal dials. Per time: divisors moved with the wave regime. Table 9 records the divisor bands as configuration memory / field notes — not final audited logs — and Figure 93 plots them against the regimes.

Phase	Approx. divisor	Interpretation
Opening / Light	3–5	very conservative; avoid overcount while density is low and passers-by dominate detections
Growth	2.3–2.9	begin recognising real increase; entrance signal starts to mean queue formation
Concert Peak	1.7–0.9	capture strong pressure under crowd attenuation — bodies absorb signal and phones under-detect at crush
Post-concert	1.9–2.6	fast waves, but controlled; freshness matters more than amplification

Table 9 — Phase-dependent divisor bands (configuration memory / field notes — not final audited logs).

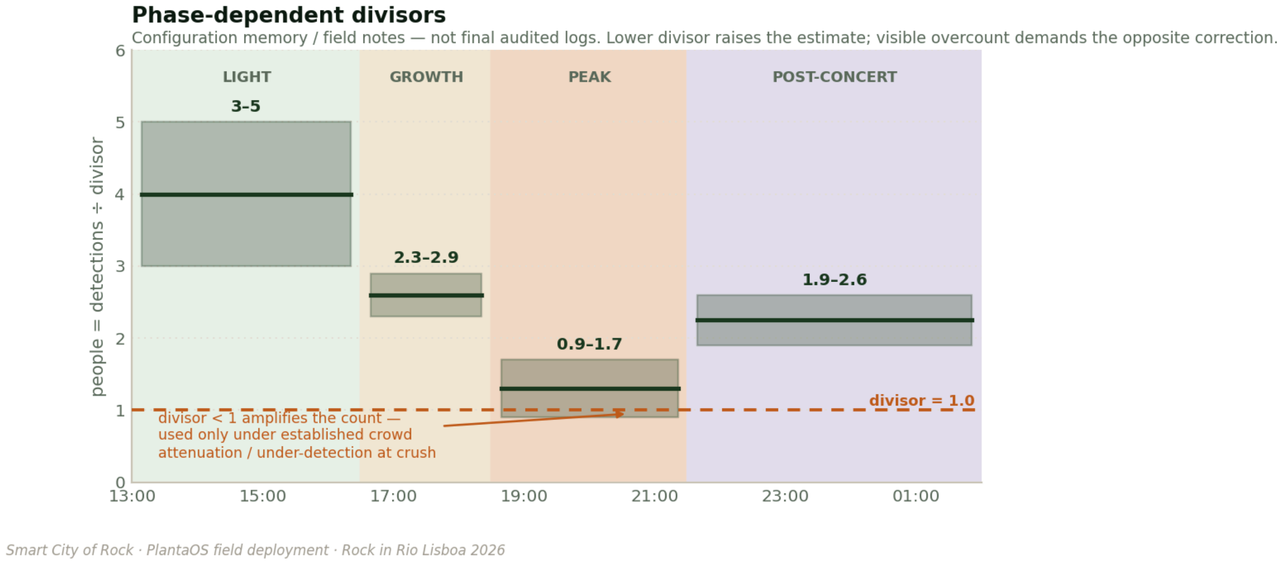

Figure 25 — Phase-dependent divisors against the wave regimes, with the sub-unity caution.

Two numerical patterns in this configuration deserve explicit explanation, because both involve values near 0.4–0.5 and they are different things. First, the fusion weights we(t) and wi(t) of Section 9.2 typically sit near 0.4/0.6 or 0.5/0.5. This is not indecision; it is the structural fact that neither signal alone is sufficient — entrance counts without interior validation book passers-by as occupants, and interior counts without entrance pressure see the wave only after it has become a queue. Near-equal weighting encodes genuine complementarity, tilted by regime: entrance weight rises toward ~0.6 at Peak and Post-concert, interior toward ~0.6 in Light.

Second, and more subtle: divisors below 1.0. A divisor divides, so a divisor of 0.9 — or conceptually 0.4–0.5 in extreme crush — amplifies the raw count rather than shrinking it. This inverts the negative-calibration direction, and it is legitimate in exactly one circumstance: established under-detection. At maximum density, human bodies attenuate radio signal (crowd attenuation), phones in pockets pressed between bodies fall below detection thresholds, and each detected device can represent more than one present person. When staff observation confirms that the sensor is seeing less than reality, dividing by less than one is the correction. The governance rule is strict and stated in the field notes: lower divisor raises the estimate, so sub-unity divisors are only valid when under-detection is established by independent observation; if overcount is visible, the correction is the opposite — raise the divisor and cut the radius. The same physical situation never justifies both.

9.6 Reconstructed calibration per cluster

Per cluster: the sensing geometry of Section 6 implies a distinct calibration posture for each cluster. Table 10 gives the per-cluster reconstruction — presented as methodological reconstruction from field notes and the verified seed-point geometry, not as audited per-node history.

WC	Structure	Expected sensors	Probable calibration	Confidence
WC-01	Male/Female	entrance + interior per section	short radius — the two seed points are nearly identical, effectively one observation; 5–10 m entrance, 10–15 m interior	medium
WC-02	Male/Female	4 complete roles	good validation cluster; 10–15 m entrance, 15–20 m interior	high
WC-03	Male/Female	4 roles, some very close	duplication risk — average/median M1/M3, never sum	high / medium
WC-04	Male/Female	4 roles, H2 suspect	exclude or down-weight H2; calibrate on M1/M2/H1	medium

WC	Structure	Expected sensors	Probable calibration	Confidence
WC-05	Unisex	ideal U1 + U2; only U2 present	low confidence; infer queue with caution; 15–20 m, never >25 m without validation	low / medium
WC-06	Unisex (largest)	U1 + U2	largest capacity; can accept 15–20 m but with strong divisor	medium / high
WC-07	Male/Female	4 roles, M2 suspect	exclude M2 as outlier; use neighbours WC-06/WC-08 as sanity check (≤15% weight)	medium
WC-08	Male/Female	3 points	good north/east pressure sentinel; note missing male interior	medium / high

Table 10 — Reconstructed calibration per cluster (methodological reconstruction from field notes and verified seed-point geometry).

Per radius: the third dial is the RSSI-gated sensing radius of Table 5, tuned per role and per cluster as above, and converging estate-wide on the tightened final configuration of Table 8. The three dials interact — a wider radius demands a larger divisor; a Peak regime justifies a smaller one — which is why calibration authority was held by a single operator with logged changes, and why automatic recalibration was disabled during show hours.

9.7 Alternative estimator formulations

The deployed estimator — linear divisor scaling with regime-dependent weights — is one member of a family. There are at least seven ways to do this calculation, each with a distinct evidence requirement: (1) Divisor scaling (deployed): people = detections/d with per-role d — transparent, field-tunable in seconds, deterministic, and auditable; its weakness is that d compresses many physical effects into one constant. (2) Regression calibration: fit detections → people against staff spot counts per cluster and regime — strictly better once enough labelled pairs exist; requires the ground-truth campaign of Section 23. (3) Flow conservation with infrared: occupancy(t) = occupancy(t−1) + entradas_ir − saidas_ir, with radio as drift correction — the strongest formulation at instrumented doorways, since IR counts are direction-resolved. (4) State-space / Kalman filtering: treat occupancy as a latent state with wave-regime dynamics and radio+IR as noisy observations — principled fusion of freshness and noise, at the cost of opacity to field operators. (5) Bayesian fusion: weight each sensor by its confidence C(s,t) as a likelihood — the natural formalisation of what the confidence engine already approximates. (6) Learned estimators: features such as RSSI histograms and dwell distributions mapped to counts — highest ceiling, highest data appetite, weakest field explainability. (7) Queueing inversion: apply Little's law — $L = \lambda W$ (Little, 1961) — to convert measured arrival rates into expected waits — the right tool for converting any of the above into the waiting-time numbers the app displays. The deployed choice was correct for a first live festival: when a number is wrong at 21:00 in front of 80,000 people, an operator must understand why in one sentence and fix it in one command. The upgrade path — regression on ground truth, IR conservation at doors, Kalman fusion above both — preserves the divisor model as its transparent fallback.

9.8 Heuristics and metaheuristics: what runs live, what runs at design time

The live layer runs deterministic heuristics only. Five are mandatory, and each is closed-form and one-sentence-explainable: the best-WC greedy ranking over the cost function distance + pressure + section type + confidence; hysteresis on public recommendations — a cluster displaces the current recommendation only past a fixed margin, which is the damping mechanism for the advertisement-inversion feedback of Section 13.4; threshold alerting on the pressure bands (40/70/90), driving steward and cleaning dispatch; the two-stage smoothing stack (60-second smart-data means feeding the five-minute published series); and confidence reweighting in closed form (Section 9.2). The boundary rule is absolute: during show hours the live tick contains zero stochastic search and zero heavy models — when a value is wrong at 21:00 in front of 80,000 people, the fix must be one command and the explanation one sentence.

The design-time layer is where metaheuristics belong. Four earn their place before doors: swarm/ACO-class search for sensor placement and cluster siting over the venue graph — the direct festival application of the FREE-guided swarm optimisation developed for buildings in the author's programme (Melo de Magalhães, 2026i, 2026o, 2026u); simulated annealing or genetic search for portable-unit allocation across clusters under the wave scenarios of the prior study (Section 19.3); scenario optimisation for Day-0 preset selection against lineup-based crowd estimates; and layout search over the gendered/unisex mix per zone, scored by the pooling economics of Section 4.3. The division of labour is the point: metaheuristics explore the combinatorial design space where time is cheap; heuristics execute in the live loop where trust is expensive.

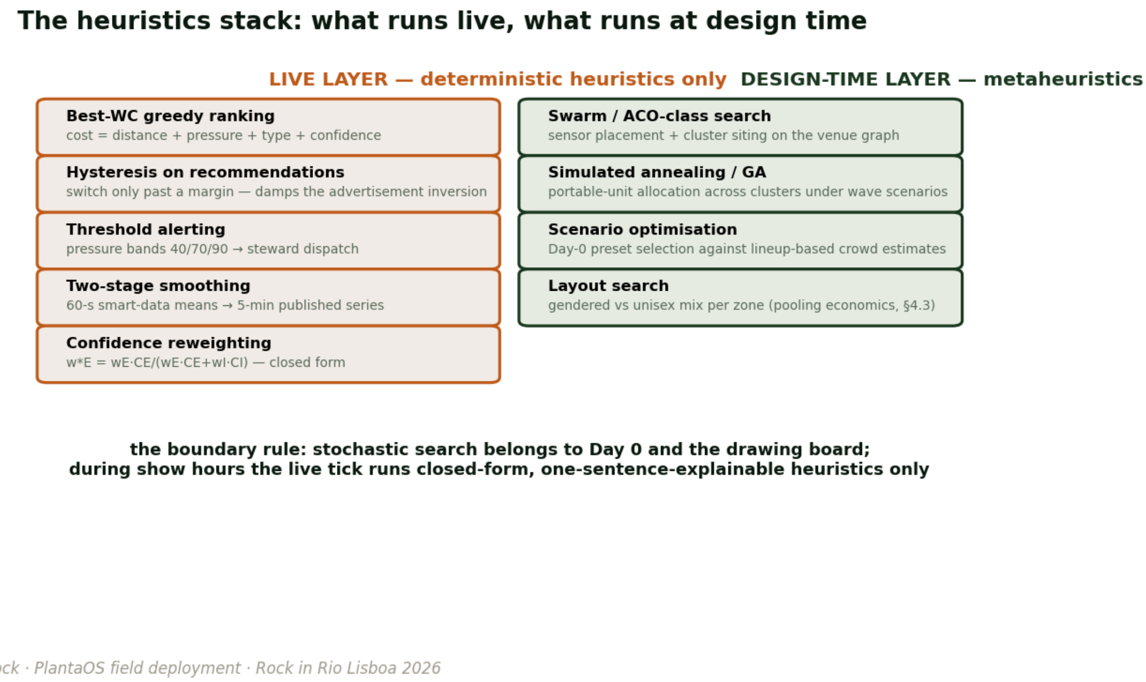

Figure 26 — The heuristics stack: deterministic heuristics live; metaheuristics at Day 0 and the drawing board.

10. Wave Flow Control: Wave-Based Calibration and Freedom Collapse

The technology category. Wave Flow Control is a technology for detecting, interpreting, and coordinating flow waves in large physical spaces — to our knowledge, a new technology category, of which the system in this paper is the first deployed instance. The word wave is load-bearing: crowds do not increase smoothly; they move in waves. At festivals, waves are triggered by gates opening, pre-concert movement, concert starts, concert endings, food and drink surges, bathroom surges, final exits, weather changes, emergencies, and transport bottlenecks. The peak is not just a bigger number — it is a phase change, the point where systems either fail or become intelligent; a system that only works during quiet periods is not a real event-intelligence system, because the value appears when the wave arrives.

If you control the waves, you do not control people; you control the conditions that make movement lighter, safer, faster, and more valuable.

Queues are not static objects; they are the residue of waves. Crowd pressure arrives in programme-synchronised surges, and the estimation problem is therefore a wave-detection problem: entrance sensors detect the beginning of a wave; interior sensors validate its result; five-minute averages smooth the wave without hiding it; negative calibration prevents false waves; confidence scoring tells operators when the system is reliable; and routing recommendations can redirect the wave before it becomes a queue failure. Each element of Sections 7–9 exists because of the wave structure — none of it is needed at a venue where demand ramps smoothly.

Operational freedom is modelled conceptually as:

$$\mathit{Operational\ Freedom}(t) = \mathit{Available\ Capacity}(t) \times \mathit{Perceived\ Alternatives}(t) \,/\, \mathit{Distortion}(t)$$

where Distortion(t) aggregates queue uncertainty, GPS error, stale sensor data, bad signage, overcrowding, lack of visible alternatives, network failure, and wrong calibration. The equation is conceptual — a design compass, not a fitted model — but it makes the intervention logic explicit. The system increases operational freedom through four applied mechanisms: (1) increasing perception — showing available WCs, waiting time, occupancy, and route options; (2) reducing distortion — calibrating noisy sensors, rejecting GPS outliers, using five-minute averages, and marking stale data offline; (3) increasing navigability — converting sensor data into actionable recommendations; and (4) creating counterfactuals — “instead of going to the nearest visible queue, go to the better available cluster.”

Wave regime	Time window	Crowd state	Main distortion	Sensor priority	Calibration behaviour	Operational action	Freedom effect
Light	13:00–16:30	Low density	Mostly low	Interior sensors validate baseline	Conservative calibration	Establish normal state	High freedom
Growth	16:30–18:30	Increasing arrivals	Rising	Entrance sensors gain importance	Adjust divisors and confidence	Detect early pressure	Freedom begins to compress
Peak	18:30–21:30	High density / concert waves	Maximum	Entrance + interior fusion required	Negative calibration essential	Redirect people; alert operations	Freedom can collapse — or be restored
Post-concert	21:30–close	Exit waves and service bursts	Sudden surges	Entrance sensors detect wave fronts; interior confirm saturation	Fast freshness checks; offline discipline	Prevent last-wave failure	Recovery and redistribution

Table 11 — Wave regimes: crowd state, distortion, sensing priority, calibration, action, and freedom effect.

The peak deserves emphasis because it is where the framework earns or loses its keep. During Light, almost any measurement approach looks adequate: distortion is low and errors are invisible. During Peak, entrance radio fields saturate, bodies attenuate signal, passers-by flood entrance counts, and naive estimates run away — exactly the conditions negative calibration (Section 9.3) was built for. A queue that forms at Peak is a freedom collapse: the wave arrived, the alternatives existed, and nobody could see them in time. The system's job is to move the intervention earlier than the collapse: the entrance sensors see the wave front minutes before it converts to occupancy, and those minutes are the window in which redistribution is still cheap.

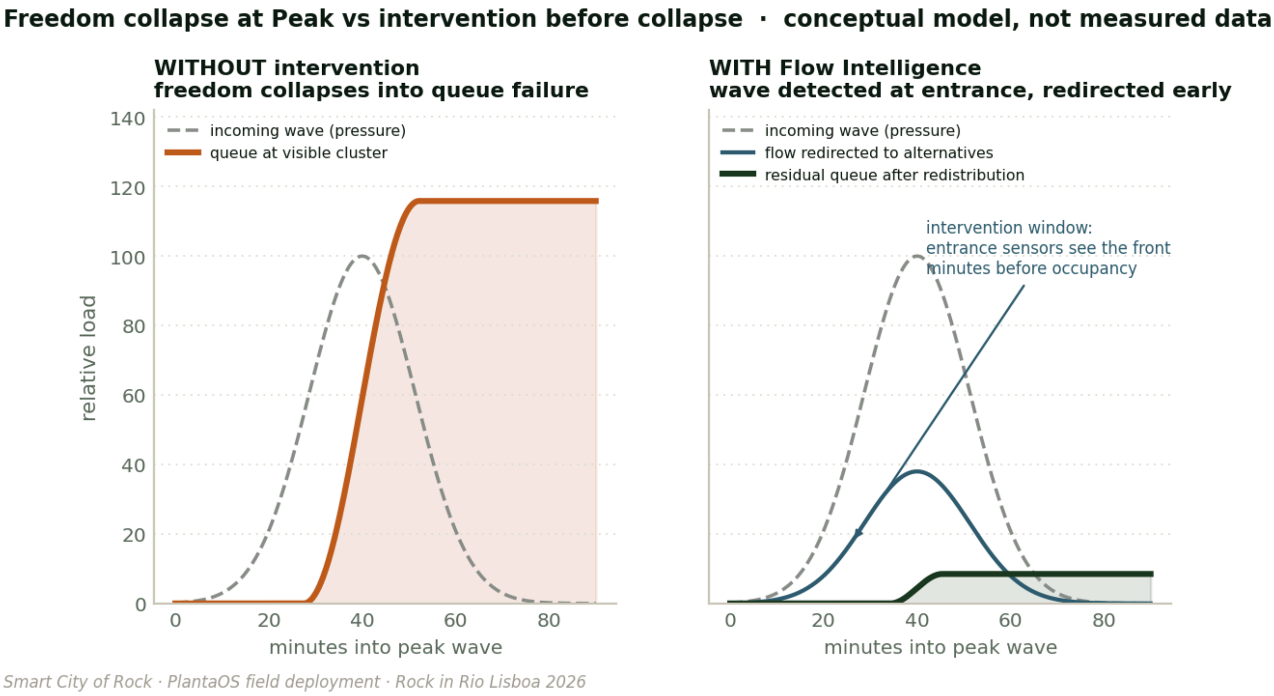

Figure 27 — Freedom collapse at peak vs Flow Intelligence intervention before collapse (conceptual model, not measured data).

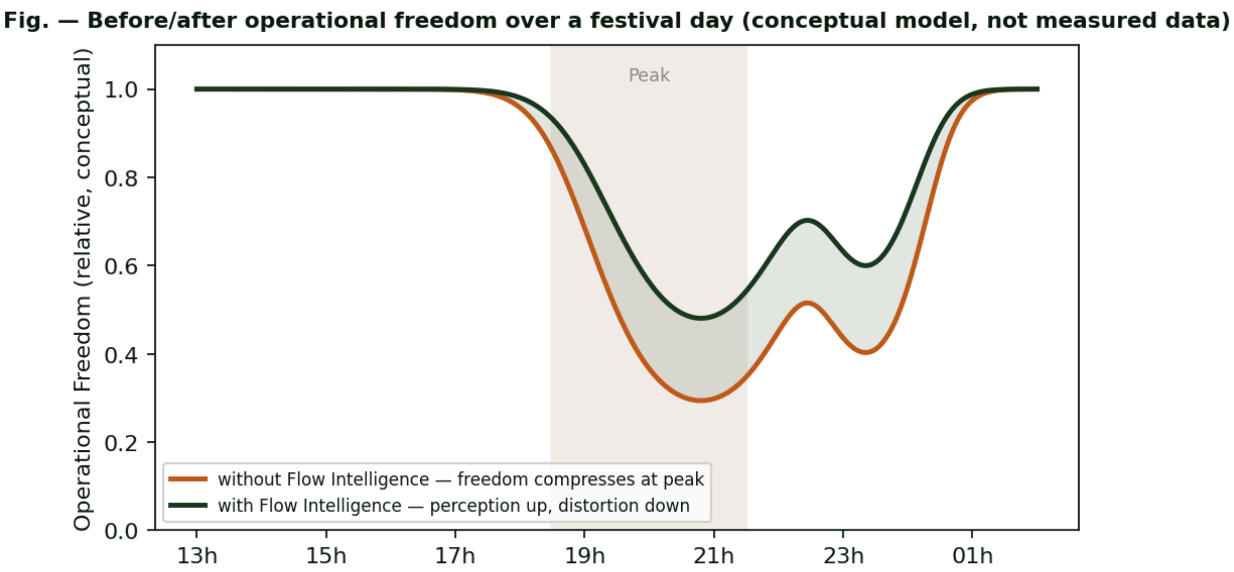

Figure 28 — Before/after operational freedom over a festival day (conceptual model): perception up, distortion down, especially during Peak and Post-concert.

11. Flow Intelligence Pipeline and External Reporting

The pipeline is: sensors → backend ingestion → calibration engine → confidence and freshness checks → cluster-level state → public screens/app → five-minute rolling average → SCOR KPIs and cluster data reporting. One estimation core serves every consumer: public recommendations, operations alerts, and external data reporting all derive from the same calibrated, confidence-scored cluster state, so no channel can disagree with another.

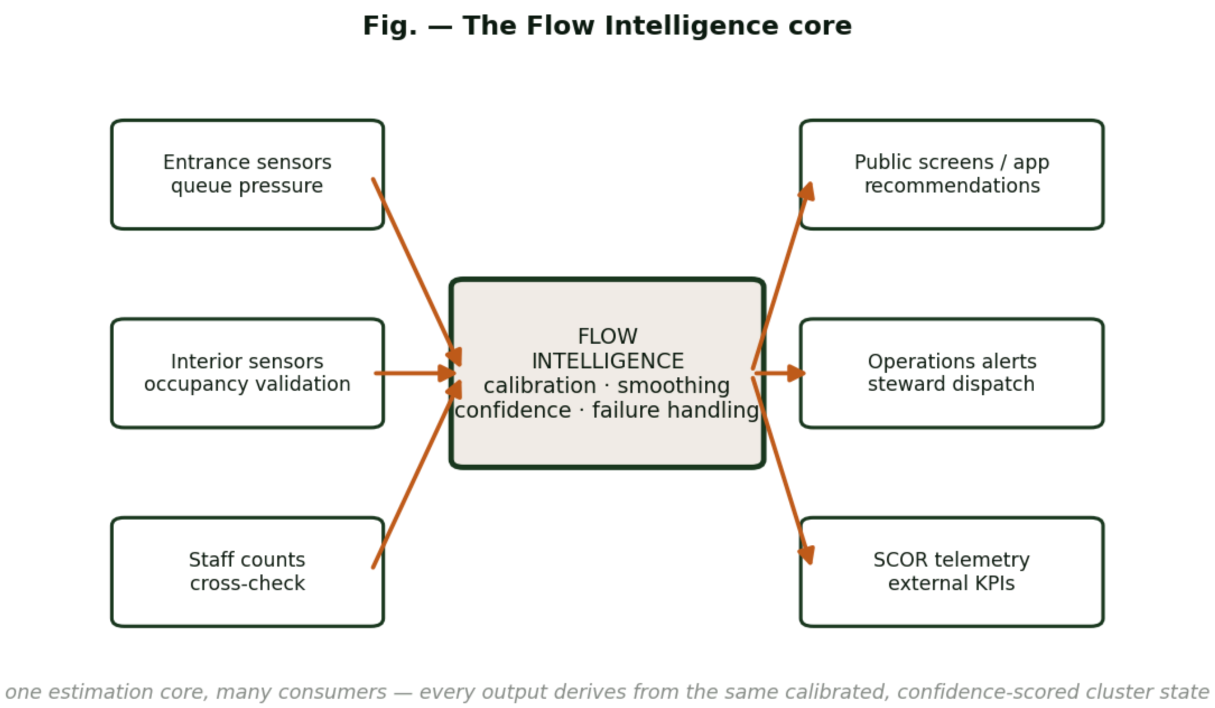

Figure 29 — The Flow Intelligence core: one estimation engine, many consumers.

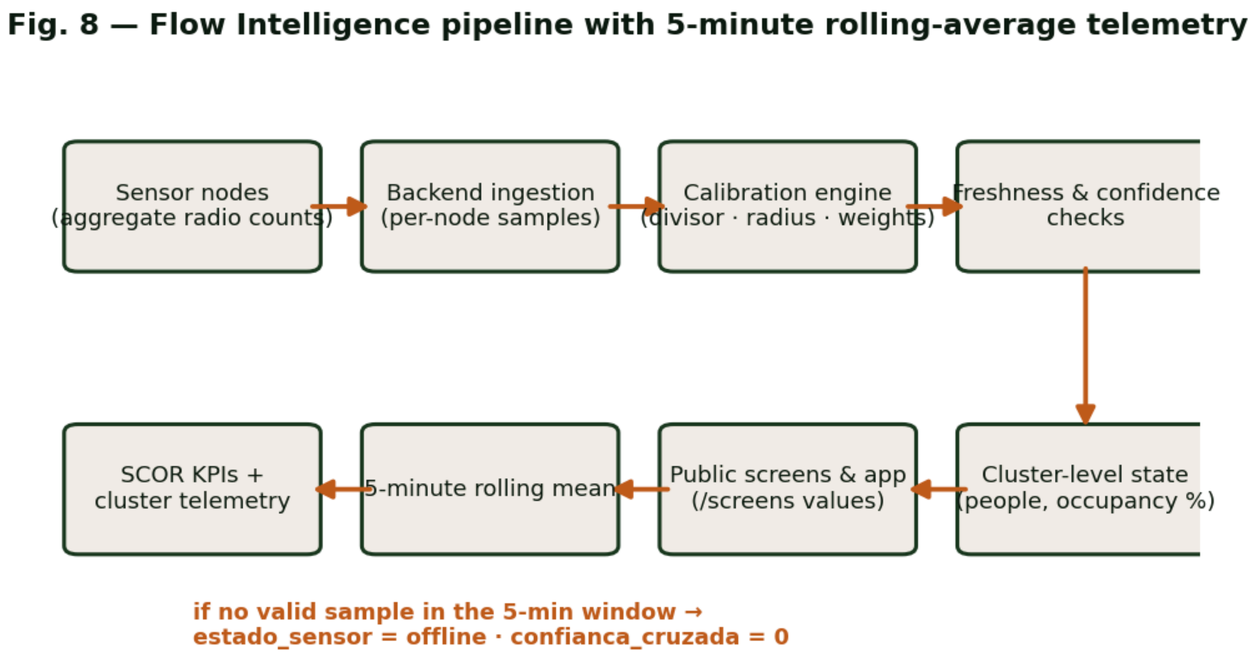

Figure 30 — System architecture from sensor nodes to public screens and SCOR data reporting.

11.1 Failure handling

Each cluster section occupies one of three data states — fresh, degraded, offline — with strict transition rules. External data reporting publishes the mean of the last five minutes, never a last-value snapshot. If fresh data exists, the calibrated average is sent. If data is stale or the window is empty, the system transmits estado_sensor = offline (or degraded) with confianca_cruzada = 0 —

it never backfills, never repeats old values as current, and never manufactures confidence. An honest "we do not know right now" is operationally superior to a stale number presented as current.

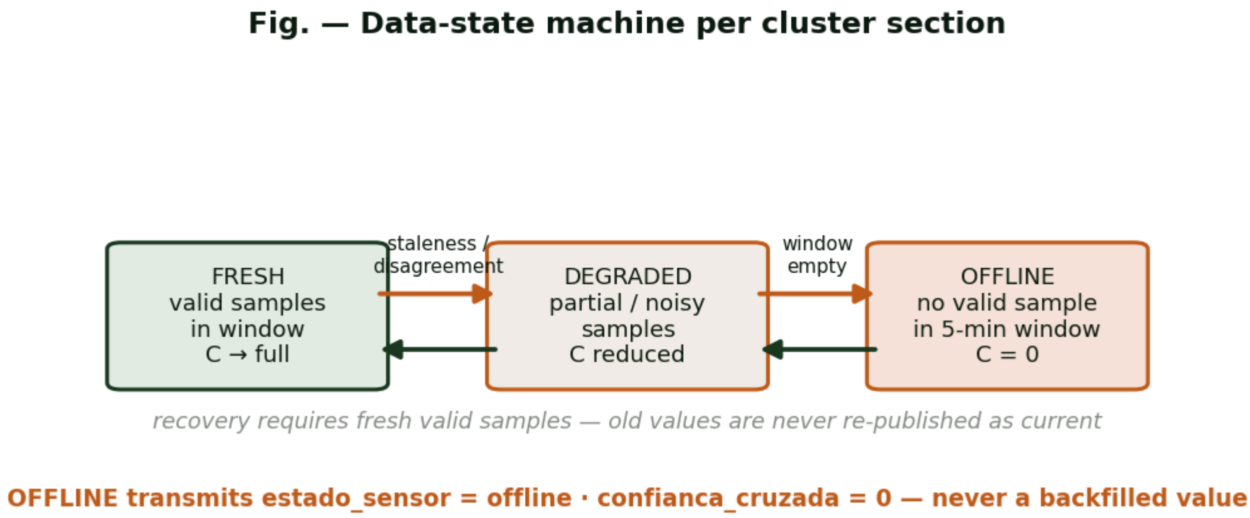

Figure 31 — Data-state machine: fresh / degraded / offline, with recovery only through fresh valid samples.

11.2 SCOR/Sensaway data schema

Integration with the SCOR/Sensaway platform transmits per-cluster-section messages on a five-minute cadence with the fields in Table 12, plus four festival-level KPIs (flow index; average WC occupancy %; active critical alerts; people redirected per day). Two prohibitions are absolute: no individual-level data is ever transmitted, and no personal trajectories are ever stored anywhere in the chain.

Field	Meaning
cluster_id	Cluster-section identifier (e.g. WC-02_M, WC-05)
timestamp	End of the 5-minute window
telemoveis_detectados	Aggregate devices detected (raw, post-dedup)
pessoas_estimadas	Calibrated people estimate P5
homens / mulheres / unissexo	Section-level estimates (male / female / unisex)
entradas_ir / saidas_ir	Infrared entry/exit counts where instrumented
ocupacao_instantanea	Occupancy % vs section capacity
contagem_prosegur	Independent security staff count when available (cross-check)
confianca_cruzada	Cross-validated confidence [0–1]
estado_sensor	online / degraded / offline

Table 12 — SCOR cluster data schema (see Appendix D for transmission rules).

11.3 The public app layer: plantarockinrio.com

The public application, served at plantarockinrio.com during the festival, is the layer through which cluster state becomes attendee behaviour. It showed, or could show: nearest WCs; occupancy; waiting level; recommended route; cluster status; gender/unisex availability; real-time pressure; and confidence or degraded/offline state where relevant. The mechanism: (1) the user opens the app; (2) the app obtains the user's location where permitted; (3) the nearest-WC calculation happens client-side where possible, so position never leaves the device; (4) the app compares distance, occupancy, and estimated wait; (5) it recommends the best WC, not necessarily the closest WC — the counterfactual routing of Section 14 made concrete; (6) public screens show the same cluster-level states for attendees without the app; and (7) the backend exports five-minute average states to SCOR/Sensaway. Because every surface reads the same confidence-scored cluster state, the app, the screens, and the operations dashboards cannot disagree. Interface evidence from the live deployment is presented in Section 13.3.

The cluster TVs. Each WC cluster carried its own screen, and the data held consistent across all of them: every cluster's TV displayed that cluster's allocation of people by gender, as occupancy from 0 to 100% per section, side by side with the venue-wide counterfactual panel pointing to the emptiest WC. The attendee standing in front of WC-02 saw WC-02's male and female bars and, beside them, where to walk instead — one confidence-scored state, rendered locally at every cluster, refreshed on the five-minute rhythm.

Figure 32 — The cluster TV, schematically: per-gender occupancy 0–100% for this cluster plus the emptiest-WC counterfactual (the deployed screens appear in Section 13.3).

The nearest bathroom is not always the best bathroom. The recommendation logic exists because of that sentence: best = distance + pressure + section type + confidence, never distance alone.

The app did not simply show where bathrooms were. It transformed bathroom choice into a real-time routing decision.

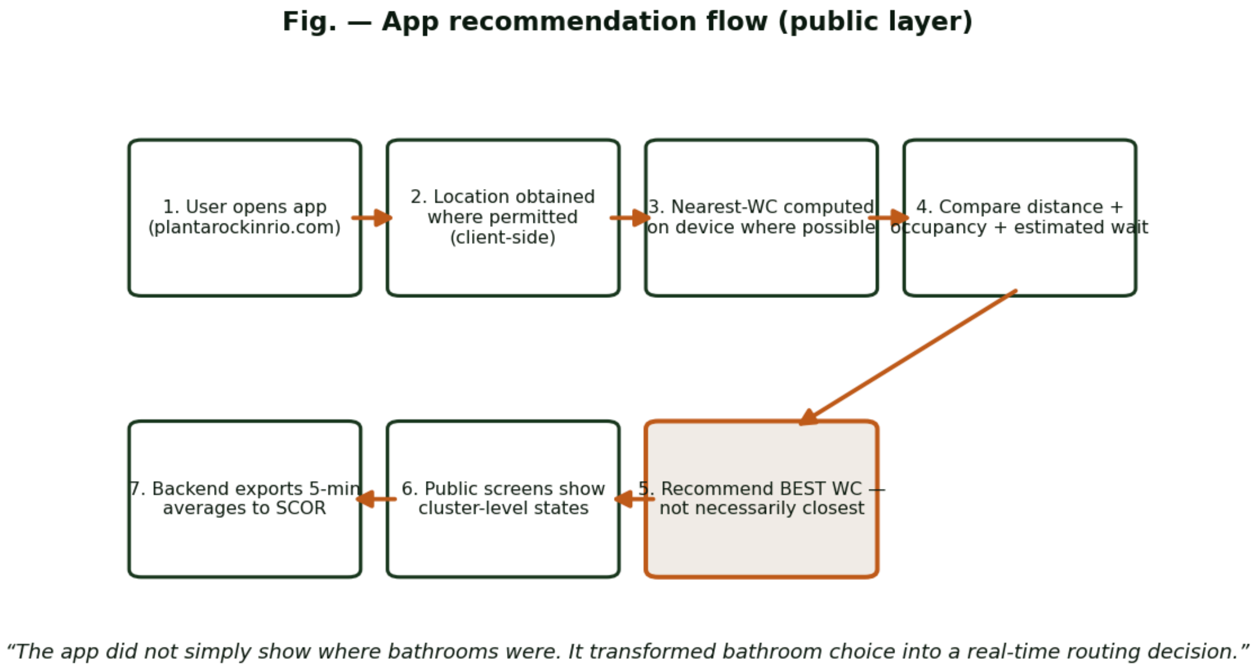

Figure 33 — App recommendation flow.

12. Privacy and GDPR

Privacy was not a legal afterthought. It was a technical constraint.

Privacy is a design constraint that shaped every layer, not a compliance appendix. The governing rule: the system counts people and flows, and nothing else. Concretely:

- No individual identification — the system's atomic datum is an aggregate count at a node over an interval; it has no concept of a person.
- No person-by-person tracking and no trajectory storage — no re-identification across nodes, no per-device histories, anywhere in the chain.
- No live user location stored where avoidable — the attendee application does not upload position to obtain a recommendation.
- Client-side nearest-WC routing where the platform permits — the computation runs on the attendee's device against the published cluster-state feed, so position never leaves the device.
- External systems receive only aggregate cluster states — the Table 12 schema, never node-level raw counts.
- Raw identifiers are never persisted — counts are computed at the edge and only counts move downstream.
- Five-minute aggregation reduces re-identification risk by temporal coarsening; small-count suppression withholds externally any window whose count falls below a threshold.

Vocabulary is legally decisive. Anonymous aggregate data — "47 devices near node X in window T" — cannot be related to an identifiable person and falls outside the GDPR's material scope. Pseudonymous device data — stored radio identifiers, even hashed — remains personal data, because re-identification is conceivable. The two are not the same thing, and architectures that persist identifiers "but hash them" are in the second category regardless of intent. This system is confined by construction to the first category, which is what permits operation at an 80,000-person venue without consent flows for sensing. Residual risk is assessed rather than dismissed: small counts at isolated nodes late at night could in principle correlate with an observed individual, mitigated by suppression thresholds, 5-minute temporal and cluster-level spatial aggregation, and retention limits on raw node samples.

Fig. 9 — Privacy-by-design architecture (people and flows only; no individual tracking)

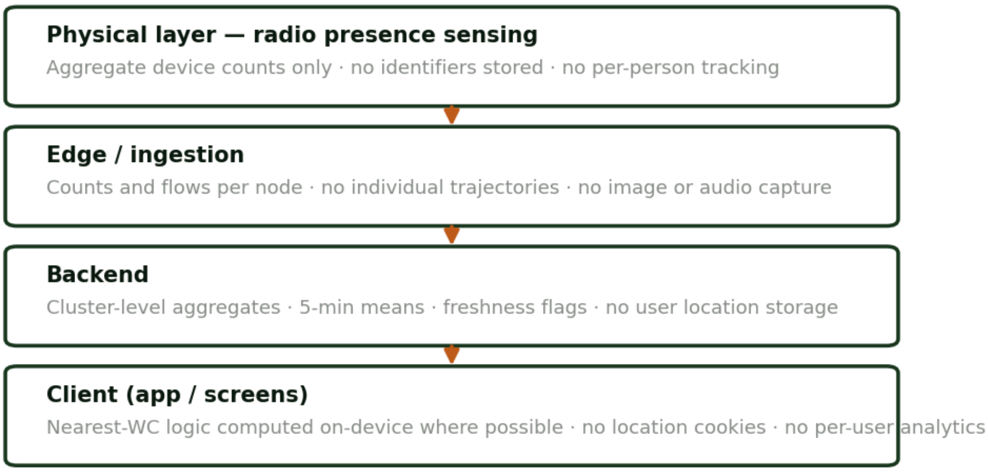

Figure 34 — Privacy-by-design architecture across physical, edge, backend, and client layers.

Fig. — Data-minimisation chain: each stage strictly reduces information

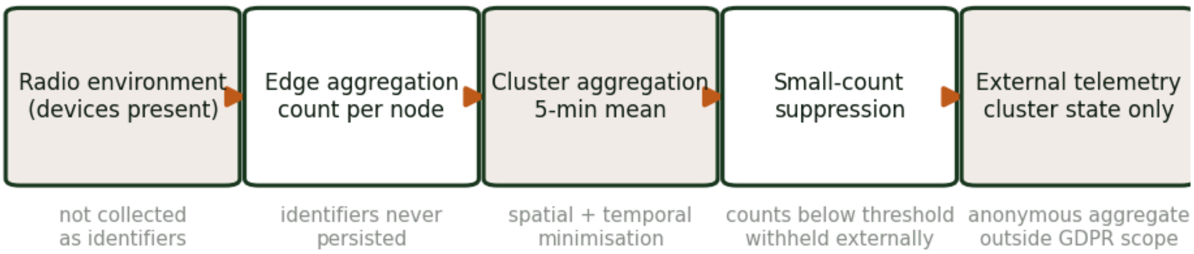

Figure 35 — Data-minimisation chain: each stage strictly reduces information; identifiers never enter the pipeline.

13. Results and Observed Outcomes

Proof of operation. The system worked in production, and that statement is field-verifiable: it ran across all four festival days at 80,000-person daily scale; the public app at plantarockinrio.com and the per-cluster TVs served live, calibrated, per-gender cluster state throughout operating hours; the SCOR feed transmitted on its five-minute cadence with auditable ingestion records; calibration converged after the Day-1 cycle, and thereafter published values tracked staff observation through

Peak and Post-concert regimes. Operational proof — the layer sensing, calibrating, and routing live at scale — is established by the deployment itself; quantified impact (exact minutes recovered per person) is the subject of the controlled ground-truth campaign of Section 23. The first claim stands on this festival; the second earns its number at the next one.

Outcomes are separated into three types — field-supported, operational, and strategic. Throughout this section, each claim states its own basis in prose: field-verified observation, backend-derived measurement, reconstruction from field notes, stated assumption, or reported figure pending verification.

- 8 WC clusters monitored across all four festival days: six gender-separated, two unisex; verified estate of 1,137 cabins; per-cluster range 102–208.
- The app and screen system was deployed (plantarockinrio.com), and SCOR/Sensaway data integration operated on the five-minute cadence; calibration was performed under real field conditions, and negative calibration was discovered and applied.
- 24 geotagged sensor seed points recorded and verified (Tables 3–4), including two flagged outliers, two near-duplicate pairs, one near-identical pair, and one cluster (WC-05) missing its entrance point — all retained with flags, none silently corrected.
- Negative calibration was required and applied estate-wide: the initial configuration systematically overestimated occupancy; the field cycle of divisor increases, radius tightening, weight reduction, and stale-sensor offlining stabilised published values, which thereafter remained plausible against staff observation through Peak and Post-concert regimes.
- The backend node registry configures nodes beyond the 24 geotagged seed points (diagnostic, inactive, historical entries); configured/active counts are reproducible from the registry at any timestamp and are reported separately from the research set.
- Valid-sample coverage varied by cluster and day; clusters with complete entrance+interior pairs sustained materially higher confidence than WC-05. Offline/stale periods were transmitted as estado_sensor = offline with confianca_cruzada = 0, never backfilled.
- SCOR transmissions ran on the five-minute cadence during operating hours; per-message success is logged and auditable in the SCOR ingestion records.
- Public screen and application usage of 35,000–40,000 views was reported for the festival period; it is included as reported usage only, and Section 13 below reconciles it against the separate Google Analytics user metric.
- Deployment scale: approximately 28–30 sensor nodes were prepared (the full 28-role reference topology plus spares); approximately 15 were active at operationally relevant moments; operational focus concentrated on WC-02, WC-03, and WC-08. These figures are field notes pending reconciliation against the backend node registry.
- Day-1 trajectory: the system functioned at opening but degraded as density rose — the overestimation problem of Section 9.3 in live form. Recovery came through the negative-calibration cycle: radius reduction, more conservative divisors, five-minute averages, offline/degraded states, and reduced trust in raw values. By the later regimes, published values tracked staff observation.
- Two usage figures circulate and they are different metrics that must never be conflated: approximately 1,900 users/visits is a Google Analytics user-level metric; 35,000–40,000 views appears to be a views/impressions-level metric. Both are carried as reported pending final log verification, and any published usage claim must state which metric it is.

13.1 Operational outcomes

Beyond the itemised evidence, the deployment produced six operational state changes: WC state became visible; queue pressure became measurable; public routing became possible; staff support became more informed; underused capacity could become visible; and sensor limitations became explicit instead of hidden — the confidence and offline discipline means the system declares what it does not know.

Positioning. By the engagement numbers available, the public application recorded at least around 35,000 views — possibly more, depending on final reporting — which indicates very strong engagement for a concrete operational use case rather than an entertainment surface. On those numbers, the deployment stands as one of the most successful startup activations of the programme in terms of public attention and practical relevance. Both statements inherit the evidence discipline above: the engagement figures are , pending final log verification, and the comparative positioning is a programme-level observation, not an audited ranking. What is not in question is category: beyond any dashboard, the deployment demonstrated a new category of festival intelligence applied to one of the hardest operational problems festivals have.

13.2 Strategic outcomes

Strategically, the deployment establishes three things: proof that festivals can have real-time WC intelligence; proof that large-event infrastructure can become adaptive rather than static; and a foundation for stadiums, airports, campuses, cities, supermarkets, and other high-density spaces (Section 17). Claims beyond this evidence base — exact queue-time reduction, guaranteed sales uplift, and estimator accuracy percentages — remain future work, awaiting the controlled ground-truth campaign of Section 23.

The status of the paper's principal claims, stated plainly: the cluster estate (eight clusters, six gender-separated and two unisex, 1,137 verified cabins), the 24 geotagged seed points with their flagged outliers, and the discovery and application of negative calibration are field-verified. The deployment of the app and screens and the five-minute SCOR cadence are field-supported and auditable in the ingestion records. The 35,000–40,000 views figure and the ~1,900 Google Analytics users figure are reported, pending final log verification, and are different metrics that must never be conflated. The deployment-scale figures (~28–30 nodes prepared, ~15 active at key moments, focus on WC-02, WC-03, and WC-08) are reconstructions from field notes pending reconciliation against the backend node registry. Queue-time reduction and the €0.2–1.0M/day economic class are stated assumptions and scenario models, not measurements or revenue. No estimator accuracy percentage is claimed, because no ground truth exists.

13.3 The product in the field: interface evidence

The following screenshots are from the live deployment: the public application at plantarockinrio.com and the venue screens co-branded Smart City of Rock · Planta Smart Homes. They are included not as decoration but as evidence: several of this paper's central claims — the RSSI bands, the divisor and weight controls, the client-side location policy, the freshness honesty, the counterfactual routing — are visible here as shipped software, not post-hoc rationalisation. Values shown are live states at the moment of capture and illustrate operating ranges, not averages.

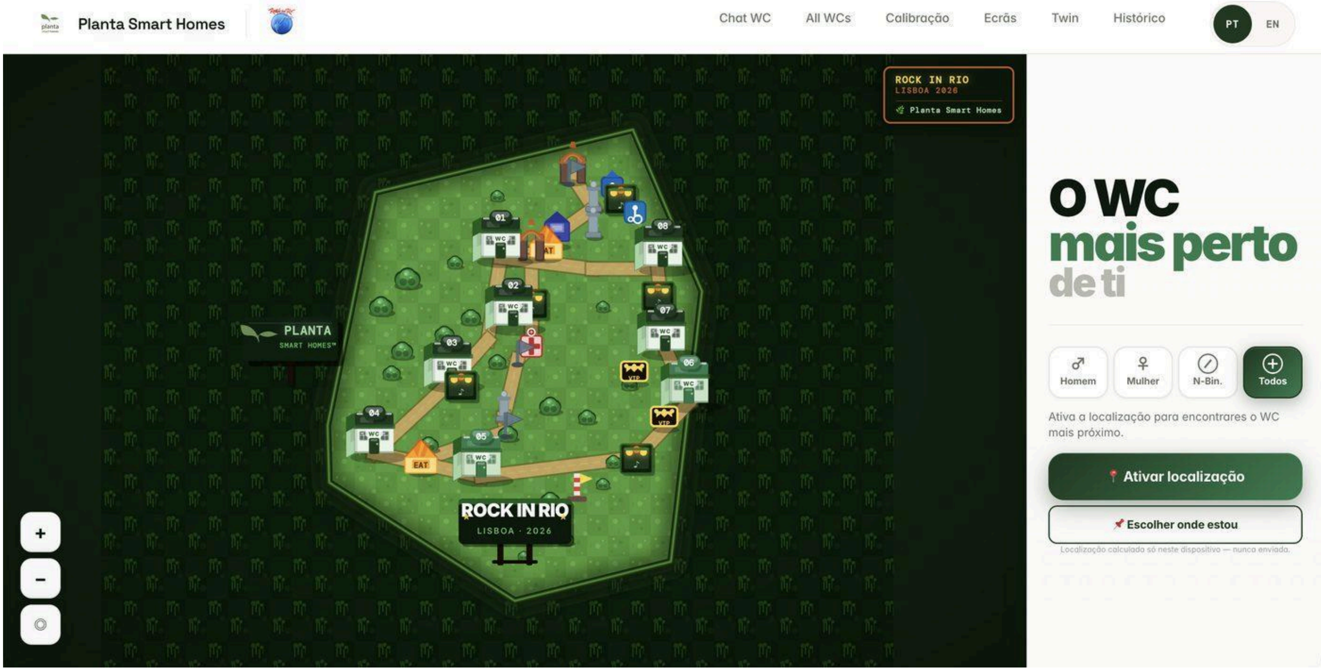

Figure 36 — The public app (plantarockinrio.com): illustrated venue map with the eight canonical WC markers, section filters (Homem / Mulher / N-Bin. / Todos), and the privacy commitment as product copy — "Localização calculada só neste dispositivo — nunca enviada" (location computed on-device only, never sent): Section 12's client-side routing, verbatim in the interface.

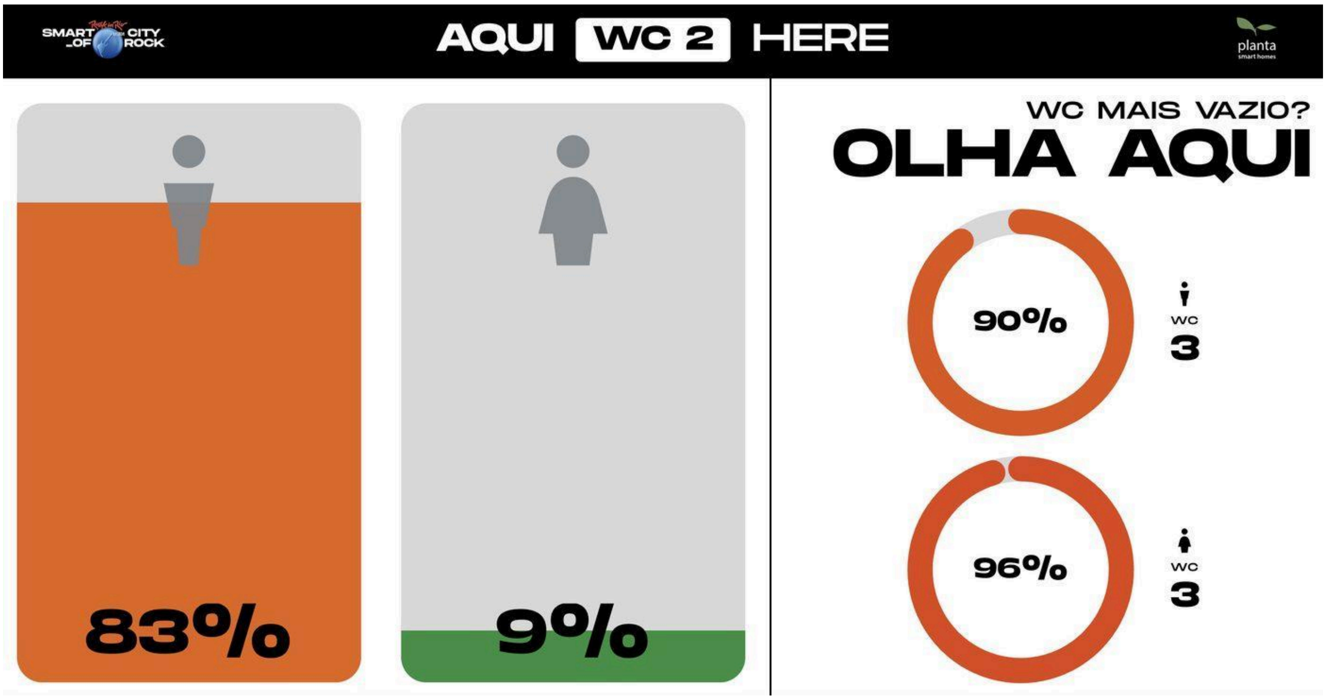

Figure 37 — Venue screen at cluster WC-02: section-level state (male 83%, female 9%) plus the counterfactual panel "WC MAIS VAZIO? OLHA AQUI" — the redirection logic of Section 14 rendered for a crowd, at cluster granularity, with zero personal data on screen.

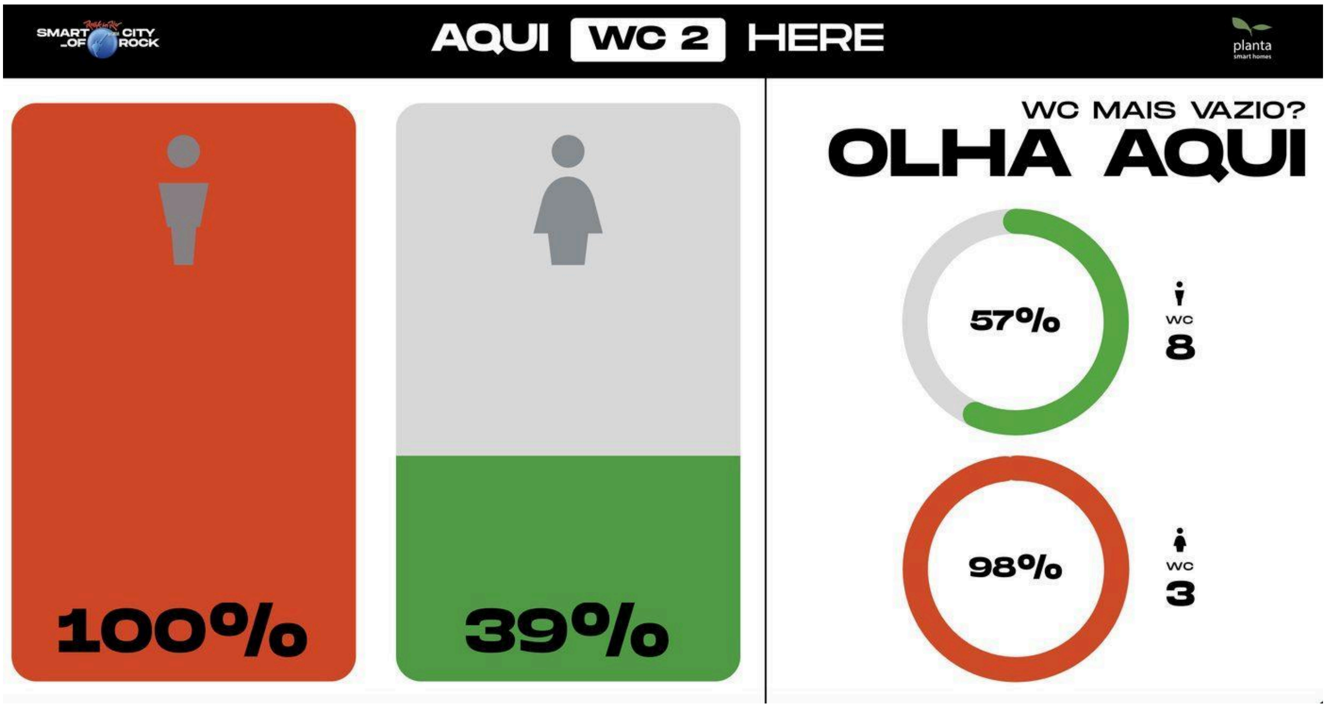

Figure 38 — Peak-state capture at WC-02: male section saturated (100%) while the female section sits at 39% — the section-asymmetry argument of Section 4 live — with the panel ranking alternatives (WC 8 at 57% vs WC 3 at 98%): the nearest WC is the wrong WC, and the screen says so.

Figure 41 — Honesty in the client: “Localização desatualizada · ±154 m · há 106 s” and “Confirma o ponto antes de seguir” — staleness and positional uncertainty surfaced to the user instead of hidden: the offline/degraded discipline of Section 11.1 extended to the attendee's own position.

Figure 42 — Live map with one canonical pin per cluster (Section 6.4) and pressure encoded in colour — WC-03 highlighted under load — alongside the on-device position dot; no sensor scatter reaches the public map.

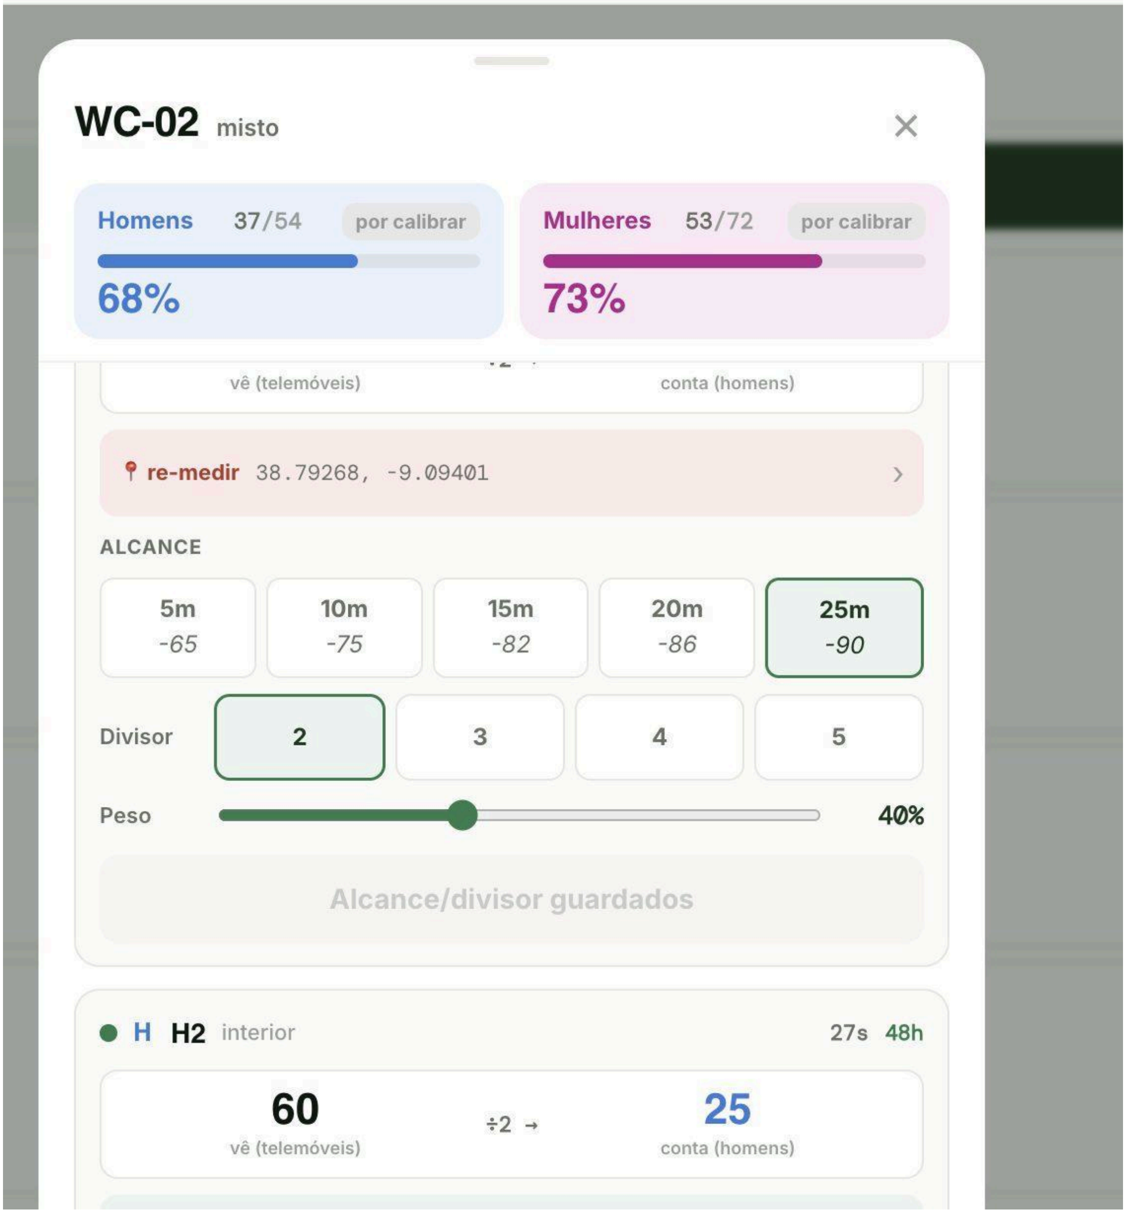

Figure 43 — The calibration surface for WC-02, and the strongest single piece of interface evidence in this paper: per-section occupancy against the official capacities of Table 2 (Homens 37/54 = 68%; Mulheres 53/72 = 73%), a "por calibrar" state flag, coordinate re-measurement (re-medir), the ALCANCE bands exactly as Table 5 (5 m/−65 … 25 m/−90 dBm), the Divisor selector (2–5) of Table 9, the Peso slider at 40% — a Section 9.2 fusion weight as a physical control — and per-node freshness (27 s). The calibration methodology of Section 9 was not a paper abstraction; it was a set of buttons operators pressed during the show.

13.4 Findings: real, and strange

The real findings. Five results earned the status of findings rather than expectations. First, overestimation is the universal starting condition: every naive configuration overcounted, at every cluster — negative calibration is a law of this sensing class, and budgeting for it is a deployment requirement. Second, the entrance lead time is real and operationally useful: entrance signals moved minutes ahead of interior saturation throughout Growth, Peak, and Post-concert, which is the entire basis of pre-emptive routing. Third, the gendered asymmetry is live and large: captures with the male section at 83–100% while the female section sat at 9–39% reproduce, in the field, exactly the disparity structure the Ghent queueing model predicts — and the unisex clusters ran visibly calmer, as pooling theory says they should. Fourth, the private backhaul earned its cost: LoRa kept reporting through the density peaks that saturated public networks. Fifth, freshness honesty preserved trust:

operators and attendees accepted 'offline' gracefully, while any stale number presented as current was corrosive.

The strange findings. Four results were counter-intuitive enough to deserve the label. The radio paradox: at maximum crowd density the sensors detect fewer devices per person present — bodies absorb signal, phones pressed between bodies drop below thresholds — so the correction inverts, and the divisor legitimately falls below 1.0 (Section 9.5): more people, weaker signal. The advertisement inversion: publicly recommending the emptiest cluster risks making it the fullest — a routing feedback loop — damped in practice by the two-stage smoothing, by partial compliance (a fraction of attendees act on any recommendation), and by ranking updates on the five-minute rhythm; the oscillation was watched for and stayed theoretical at this deployment's compliance levels. The staleness effect: users trusted an honest 'offline' more than a plausible old number — absence of data read as integrity, stale data read as deceit. And the metrics gap: 35,000–40,000 views against ≈1,900 analytics users is a 20× spread between two legitimate measurements of the same product — a reminder, inside a paper about measurement, that the first calibration problem is always the meaning of one's own metrics.

Performance statement. The system did not prove perfect queue elimination. It demonstrated that a privacy-preserving, low-cost, deployable Bathroom Flow Intelligence layer can make WC pressure visible and produce actionable recommendations under real festival constraints — and outperform static, manual, and app-only alternatives in operational actionability, privacy, and deployability. We deliberately publish no accuracy percentage: no continuously logged independent ground truth exists for the full estate, and quoting precision against spot checks would overstate the evidence. In summary: the system demonstrated a working mechanism for privacy-preserving Flow Intelligence under field constraints — 80,000 people per day, 145,000 m², 8 clusters, GPS uncertainty, mobile/Wi-Fi instability, noisy radio-presence signals, peak waves, post-concert surges, privacy constraints, external SCOR data reporting, and public-facing recommendations, all simultaneously. It did not yet prove exact causal queue reduction; future work must measure controlled queue reduction, accuracy, and economic conversion.

14. From Flow Measurement to Freedom Restoration

A queue is a freedom failure — not because waiting is always bad, but because the person lacks visible alternatives. The same physical infrastructure can produce very different freedom depending on whether people can see and navigate available paths: 1,137 cabins distributed over eight clusters are either a set of options or a set of mysteries, and the difference is information, not concrete. This is the precise difference between the deployed system and conventional solutions: static signage gives location; apps give maps; CCTV gives observation; manual counting gives delayed reports; Flow Intelligence gives navigable counterfactuals in real time.

The value of the system is not merely that it counts people; it creates better counterfactual paths.

The counterfactual structure is best seen in paired scenarios. Without Flow Intelligence, an attendee walks to the visible WC queue and waits ~20 minutes; with it, they see another WC 200 m away with lower pressure and save an estimated 10–15 minutes. Without Flow Intelligence, staff react after a queue becomes visible; with it, staff see pressure building before it becomes operational failure. Without Flow Intelligence, the venue is managed as a static map; with it, the venue becomes a live navigable system. (The per-trip time figures are scenario estimates consistent with observed cluster spacing — WC-to-WC distances up to ~547 m, alternatives typically a 2–4 minute walk — not measured averages.)

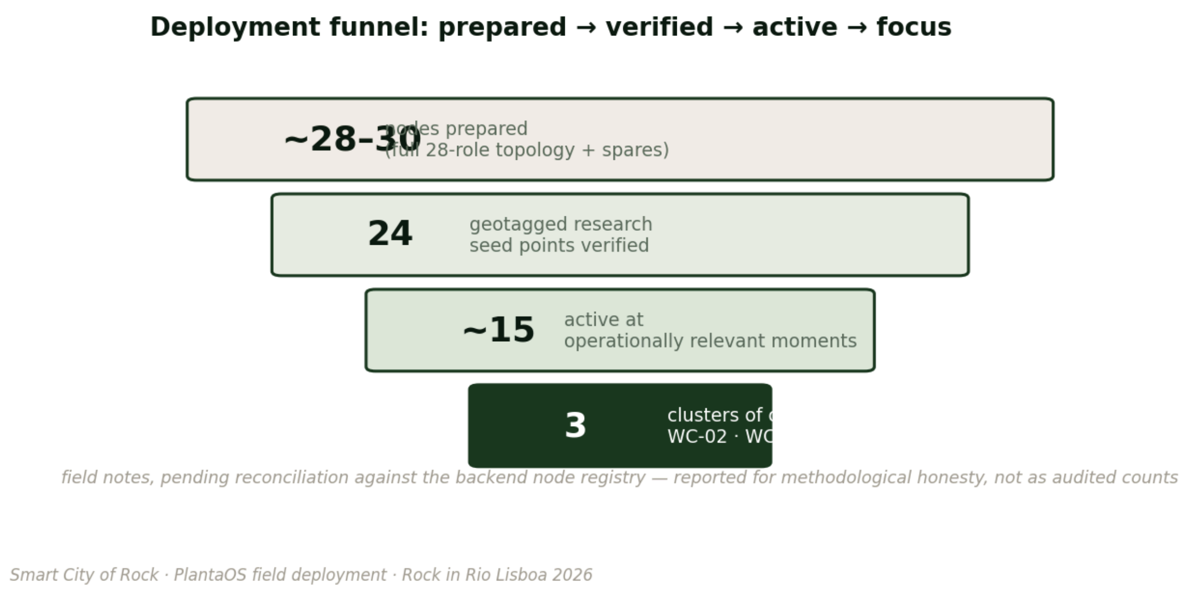

Figure 44 — Deployment funnel: prepared, verified, active, focus (field notes).

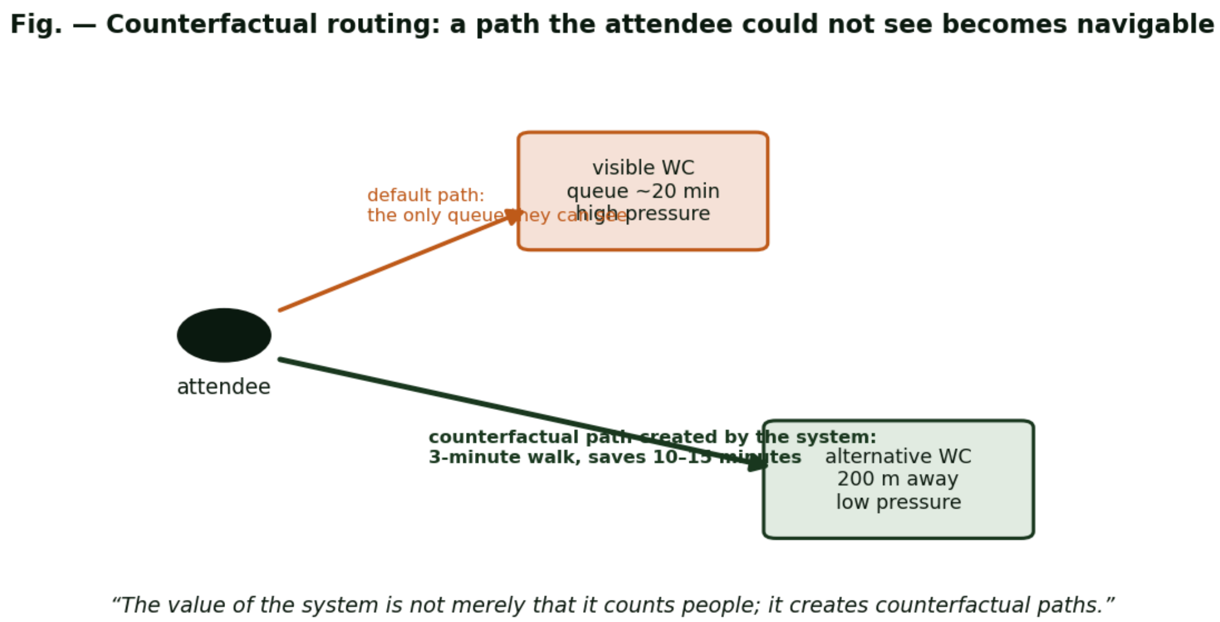

Figure 45 — Counterfactual routing: the system makes a path navigable that the attendee could not see.

The differentiation can be stated sharply: the system does not only detect crowding. It detects when freedom is about to collapse — and creates alternative paths before the collapse becomes visible as a queue.

14.1 Three counterfactual vignettes, priced

The counterfactual logic becomes concrete when run against states the system actually displayed (the Section 13.3 captures). Vignette 1 — the saturated pair. The screens record WC-03 at 98–100% while WC-08 sits at 54–57%, roughly 300 metres apart. Without the system: the WC-03 queue is the visible one, it self-reinforces, and WC-08's capacity is invisible. With the system: the "OLHA AQUI" panel and the app route the marginal attendee to WC-08 — each moved person descends the steep part of the Kingman curve (Section 16.2), and at the documented utilisation gap the convexity argument of Section 16.3 removes the majority of that person's expected wait. If redirection recovers 15 minutes each for even 5,000 people over a peak evening, that is 1,250 hours ≈ €12,500 of experience value at €10/h — from one cluster pair, one evening, zero added capacity. Vignette 2 — the gendered asymmetry. A capture shows the male section of WC-02 at 83–100% while the female section sits at 9–39%. A plain cluster average would have called WC-02 “fine”; the worst-section rule (Section 9.2) declares it problematic, and section-aware routing sends men elsewhere while women walk straight in. This is the field version of the Van Hautegem–Rogiest result: waiting disparity is a layout-and-information problem before it is a capacity problem. Vignette 3 — the post-headliner wave. The show ends; tens of thousands move at once; every entrance sensor spikes minutes before any interior saturates. Without the system, staff learn about the wave when the queues are already long — the 2024 pattern. With it, the Post-concert regime weights (70–85% entrance) fire, screens re-rank clusters, and stewards are dispatched by measured pressure while the wave is still in transit. The recovered minutes of this single daily wave are plausibly the largest line item in Table 18.

The difference is not only shorter queues. The difference is that the space stops being blind.

15. Comparison with Alternative Technologies

15.1 Why approximations are better than certainty here

A recurring objection deserves a direct answer: the system's counts are approximate, so why not deploy technology that counts exactly? Because the decision the system serves does not need exact counts — it needs ordinal correctness (which cluster is more loaded), freshness (now, not ten minutes ago), and honesty (a confidence value that admits what is unknown). A routing recommendation is correct if it sends the attendee to the genuinely better cluster; it is equally correct whether the better cluster holds 61 or 68 people. Meanwhile, every technology that approaches count certainty — turnstiles, gated lanes, calibrated video — pays for it in deployment time, physical intrusion, free-flow obstruction, and privacy cost, and each of those failure modes is disqualifying at a temporary venue (Tables 13a–13b). The trade is asymmetric: a fresh estimate within ±20% beats a perfect count that is ten minutes old, because at Peak a ten-minute-old count describes a different world. Approximation here is not a compromise reluctantly accepted; it is the epistemically honest optimum for the decision being made — provided, and this is the system's whole discipline, that the approximation carries its confidence with it and degrades to “offline” rather than to fiction.

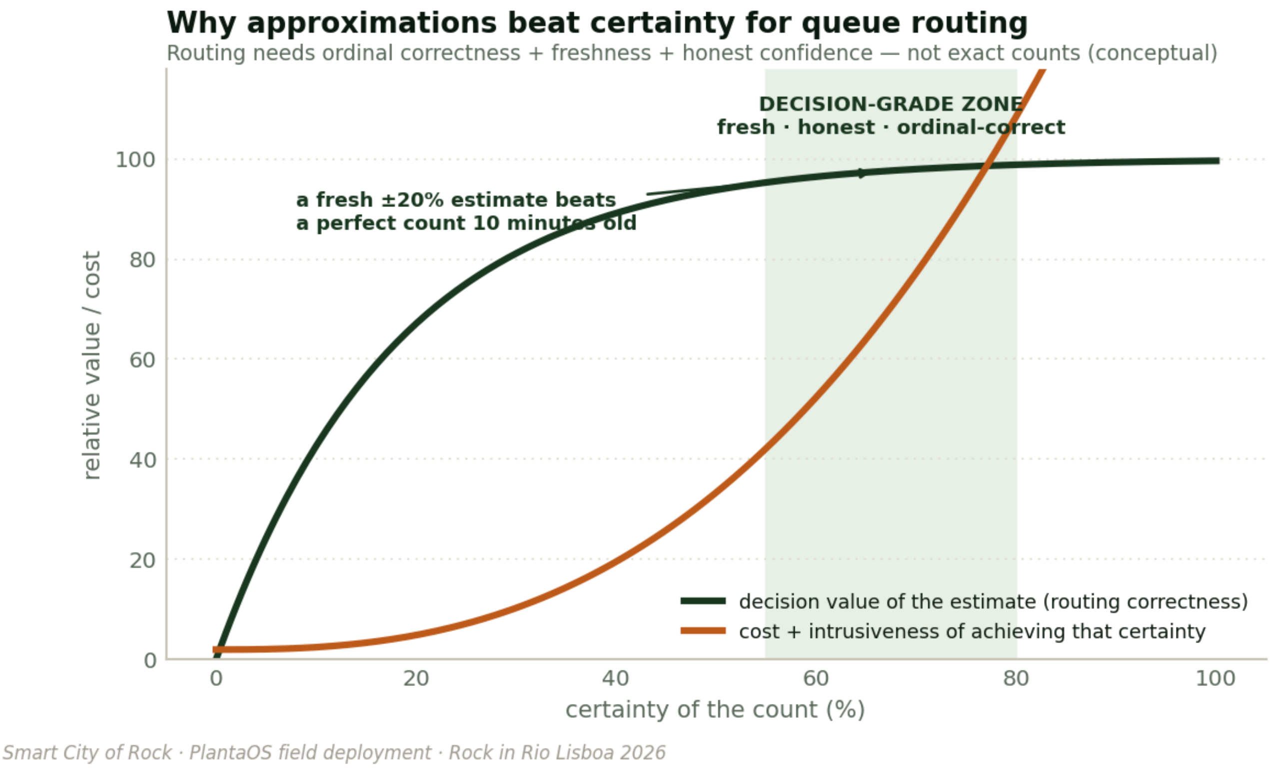

Figure 46 — Decision value vs the cost of certainty: queue routing lives in the decision-grade zone (conceptual).

Tables 13a and 13b compare ten approaches across twelve criteria relevant to a temporary high-density venue. The comparison is deliberately not a victory lap: several alternatives dominate on individual axes.

Method	Real-time	Privacy risk	GDPR burden	Cost	Deploy speed	Accuracy potential	Scalability
1. Static maps	None	None	None	Low	Fast	N/A	High
2. Static signage	None	None	None	Low	Fast	N/A	High
3. Manual staff counting	Poor (sampled)	Low	Minimal	High (labour)	Immediate	Medium at samples	Poor
4. CCTV / video analytics	Good	High (imagery)	Heavy (DPIA)	High	Slow	High when tuned	Medium
5. App-only voluntary reporting	Poor	Medium	Medium	Low	Fast	Low, biased	High
6. QR reporting	Poor	Low–medium	Light	Low	Fast	Low, sparse	High
7. Raw Wi-Fi/BT probe counting	Good	Medium (IDs if stored)	Medium–heavy if IDs kept	Low	Fast	Medium (noisy)	High
8. Turnstiles / gates	Excellent	Low	Light	High	Slow (civil works)	High	Poor at temp venues
9. LoRa / private IoT sensors	Good	Low (aggregate)	Minimal	Low–medium	Fast (days)	Medium, improvable	High
10. Planta hybrid Bathroom / WC Flow Intelligence	Good	Low (aggregate)	Minimal	Medium	Fast (days)	Medium–high after calibration	High

Table 13a — Alternative technologies: real-time capability, privacy, cost, deployment, accuracy, scale.

Method	Failure mode	Public trust	Actionability	Recommends better paths	Temp-festival fit
1. Static maps	Cannot reflect state	Medium	None	No	High
2. Static signage	Cannot reflect state	Medium	None	No	High

Method	Failure mode	Public trust	Actionability	Recommends better paths	Temp-festival fit
3. Manual counting	Fatigue; gaps; no nights	High	Low — snapshots	No	Medium
4. CCTV / video	Occlusion; lighting; privacy objection	Low–medium	High	Indirect	Poor (infrastructure)
5. App-only reporting	Participation collapses at peak	Medium	Low	Weakly	High
6. QR reporting	Nobody scans while queuing	Medium	Low	Weakly	High
7. Raw Wi-Fi/BT probes	MAC randomisation; radius spill	Medium	Medium	Not alone	High
8. Turnstiles	Blocks free flow; queues at gate	Medium	Medium	No	Poor — unusable at WCs
9. LoRa / IoT sensors	Calibration drift; power	High	High	With backend	High
10. Planta hybrid	Needs field calibration cycle	High	High — routing + alerts	Yes — native	High

Table 13b — Alternative technologies: failure modes, trust, actionability, routing capability, festival fit.

15.2 The value comparison: what each euro of each method buys

Capability comparisons (Tables 13a–13b) answer what each technology can do; procurement decisions turn on what each euro buys. Table 14 compares the methods on value terms: deployment cost class per cluster, time to operational, privacy liability carried by the operator, and — decisively — the value mechanism: whether the method merely measures the queue or acts on it. A measurement without routing recovers nothing; its value is diagnostic. Only methods that close the loop to attendees participate in the recovered-value equations of Section 16.

Method	Cost class / cluster	Time to operational	Privacy liability	Value mechanism
CCTV + computer vision	high (cameras, licensing, GDPR/DPIA)	weeks	high — image processing of crowds	measures; routing requires separate layer
Turnstiles / gated counting	high (hardware + civil works)	weeks; obstructs free flow	low	measures exactly; blocks the flow it measures
Manual counting staff	medium recurring (labour, 4 days × 8 clusters × 2 roles)	days	low	measures during shifts only; no continuity, no ro
QR self-reporting	low	days	low	sparse, biased samples; no coverage guarantee
Uncalibrated Wi-Fi analytics	low–medium	days	medium	measures wrongly under crowd conditions; false
App-only (no sensing)	low	days	low	2024's WhatsApp channel: routing without meas
Flow Intelligence for Festivals (this work)	low–medium	T−7 install + Day 0 calibration	minimal — aggregate-only, no identifiers	measures, calibrates, and routes: participates di $N \cdot \Delta t \cdot v$

Table 14 — Value comparison of methods: only closed-loop methods participate in the recovered-value model of Section 16.

15.3 Why this, not that: the design-decision record

Every architecture is a sequence of rejections. Table 15 records the principal decisions of this system, the rejected alternative, and the reason — so that a reader (or a buyer) can audit the logic rather than trust it.

Decision taken	Alternative rejected	Why
Aggregate radio presence counts	Cameras / computer vision	privacy by incapacity, not policy; deployable in days; no DPIA burden; no image ever exists
Cluster-level state	Individual tracking	routing needs cluster pressure, not persons; GDPR-minimal; ethically cleaner; equally actionable

Decision taken	Alternative rejected	Why
Divisor calibration model	Learned (ML) estimator	one-sentence explainability at 21:00 in front of 80,000; field-tunable in one command; no training da existed
Entrance-dominant fusion (≈70/30)	Occupancy-only estimation	occupancy sees the wave only after it becomes a queue; the entrance sees it minutes earlier (Secti 9.2)
Capacity-normalised pressure	Raw count comparison	80 people means different things in WC-03 and WC-06; normalisation is what makes ranking mean
Five-minute rolling publication	Real-time raw feed	stability beats immediacy for public surfaces; spikes and dropouts never reach attendees
Offline/degraded states	Repeating last known values	freshness honesty preserves trust; a stale number is a lie with confidence
LoRa private backhaul	Public 4G/Wi-Fi backhaul	public networks saturate exactly at Peak — the moment the system must not fail
USB power banks, 10-h swap	Mains wiring / generators	no cabling through crowds, no permits, no diesel; sustainability and safety in one decision (Section
Best-WC routing	Nearest-WC routing	the nearest WC is frequently the saturated one; best = distance + pressure + type + confidence

Table 15 — Design decisions: what was chosen, what was rejected, and why.

Honest reading: CCTV can be strong on raw accuracy but is weak on privacy and deployment burden at a temporary venue. Turnstiles count precisely but are intrusive and unsuitable for free-flow WC access. Manual counting is simple but far too slow. Static signage never adapts. App-only reporting is participation-biased and fails at exactly the peak moments it is needed. Raw Wi-Fi/Bluetooth probe counting is inexpensive but noisy, and if identifiers are stored it inherits GDPR burden. The deployed system is not perfect; its advantage is orchestration: anonymous aggregate sensing + entrance/inside topology + calibration + wave regimes + confidence + five-minute smoothing + app routing + operational data reporting, functioning together. Static maps and signage — the incumbent baseline at most festivals — score high on cost and deployment but provide zero state and zero recommendations, which is precisely the gap this system fills.

16. Economic Value Model through Freedom Intelligence

At Rock in Rio-class scale — on the order of 80,000 to 100,000 people per day — even fractions of an hour recovered per person aggregate into economically significant totals. This section replaces intuition with a formal model built from six equation families, each drawn from an established literature: value-of-time economics, queueing theory, convexity analysis, aggregate valuation, the psychology of waiting, and spend displacement. Every quantity below is a scenario-based model of potential recovered value under explicitly stated assumptions, with each assumption declared at the point of use.

16.1 The value of an hour: three converging anchors

The price of festival time can be anchored three independent ways, and the anchors converge. First, the ticket-implied rate — the revealed price the attendee has already paid per hour of experience:

$$v\_ticket = P\_ticket \,/\, H\_event \qquad (€90–€100)\,/\,(10–12\ h)\ \approx\ €7.50–€10.00\ per\ hour$$

Second, the wage-based value of time from the economics of time allocation (Becker, 1965; DeSerpa, 1971): appraisal practice across Europe values non-work time at roughly 30% of the gross wage rate, per the largest meta-analysis of time valuations conducted — 3,109 valuations from 389 European studies (Wardman, Chintakayala, & de Jong, 2016):

$$v\_wage = \beta \cdot w, \quad \beta \approx 0.30$$

Third, the waiting premium: the same appraisal literature consistently values waiting and queuing time at a multiple of ordinary in-vehicle or activity time — conventionally on the order of 2× — because waiting is unproductive, uncomfortable, and involuntary (Abrantes & Wardman, 2011; Wardman et al., 2016; Small, 2012):

$$v\_wait = m \cdot v, \quad m \approx 2.0\text{–}2.5$$

This paper's headline calculations use v = €7.50–€10.00/hour — the ticket-implied anchor — and deliberately ignore the waiting premium. The scenario values below are therefore conservative lower bounds: pricing queue time at the ordinary experience rate when the literature says queue time is worth a multiple of it.

16.2 The queueing mechanics: why Peak is a phase change, with an equation

Waiting is governed by queueing theory, and two canonical results carry the entire operational argument. Little's law, $L = \lambda W$ (Little, 1961), links the three observables — queue length L, arrival rate λ, and expected wait W — so that any two determine the third: it is what converts the system's entrance-pressure estimates into expected waiting times. And the heavy-traffic approximation (Kingman, 1961; Kleinrock, 1975) gives the expected queue wait for a section with c service positions, service time $\tau = 1/\mu$, and utilisation $\rho = \lambda/(c\mu)$:

$$Wq \approx (\rho / (1 - \rho)) \cdot ((Ca^2 + Cs^2) / 2) \cdot \tau$$

This equation is the economics of the whole system in one line, because it exposes the only two levers that exist. The first factor, $\rho/(1-\rho)$, explodes nonlinearly as utilisation approaches 1 — from 3× at $\rho = 0.75$ to 19× at $\rho = 0.95$ — which is the mathematical content of the claim that Peak is a phase change rather than a bigger number: near saturation, small increments of demand produce enormous increments of waiting. The second factor carries the arrival variability Ca^2 — and festival demand arrives in waves, which is precisely high Ca^2. Wave Flow Control operates on both levers simultaneously: redistribution lowers ρ at the overloaded section, and early wave detection with pre-emptive routing smooths arrivals, lowering Ca^2. Nothing in the venue's concrete changes; both factors of the waiting equation drop.

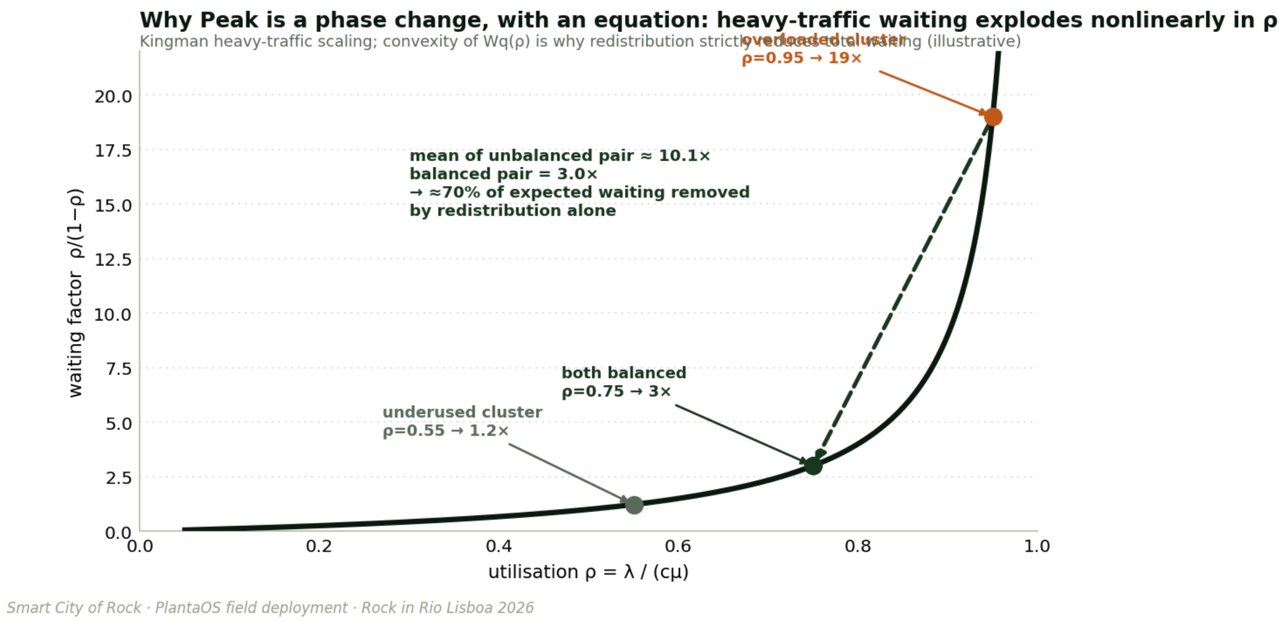

Figure 47 — The heavy-traffic waiting factor ρ/(1−ρ): nonlinear blow-up near saturation, and the effect of balancing an overloaded/underused pair (illustrative).

16.3 Why redistribution creates value: the convexity argument

Because $Wq(\rho)$ is convex, Jensen's inequality applies directly: for K sections with utilisations $\rho_1 \ldots \rho_K$ and mean $\bar{\rho}$,

$$(1/K)\ \Sigma\ Wq(\rho_k)\ \ \geq\ \ Wq(\bar{\rho})$$

with strict inequality whenever loads are unequal. Total demand and total capacity fixed, imbalance alone creates waiting — and rebalancing alone removes it. An illustrative pair from the deployment's operating range: sections at $\rho = 0.95$ and $\rho = 0.55$ have waiting factors 19.0 and 1.2, mean ≈ 10.1; balanced at $\rho = 0.75$ each has 3.0 — roughly 70% of the expected waiting removed by redistribution alone, with zero added capacity. This is the theorem behind the screens' “WC MAIS VAZIO? OLHA AQUI” panel: every attendee moved from the saturated section to the underused one descends the steep part of the curve. It is also the formal justification for the counterfactual sections of this paper: the value of the system is not that it builds toilets; it is that it operates the convexity.

16.4 Aggregate recovered value and its sensitivity

The aggregate model prices the recovered time across the crowd:

$$V = N \cdot \Delta t \cdot v \qquad \text{and per festival:} \quad V\_total = \Sigma\_d\ \ N\_d \cdot \Delta t\_d \cdot v$$

where N is daily attendance, Δt the average waiting time avoided per person per day, and v the hourly value from 16.1. The model is linear, so its sensitivity is a constant:

$$\partial V/\partial \Delta t = N \cdot v \quad \rightarrow \quad \text{at } N = 80{,}000 \text{ and } v = €10/h\text{:}\ \approx €13{,}333 \text{ per minute saved per day}$$

What is a normal wait — and why 60 minutes is not the ceiling. The scenario grid needs empirical anchors, and they exist at three levels. At the level of a single well-run queue, behavioural evidence puts the tolerance threshold at roughly six minutes before people balk, and industry planning treats anything beyond ~20 minutes as a visible failure; surge-window analyses of dance-festival sanitation describe 40-minute lines at headliner breaks as the canonical nightmare scenario. At the level of gendered asymmetry, the queueing-theoretic study of toilet layouts by Van Hautegem and Rogiest (2017) shows women averaging over six minutes against seconds for men under equal-split gendered layouts — a roughly 40-fold disparity in the waiting term — and shows mixed/unisex layouts cutting women's waits to under a minute and a half; industry trials at Rock Werchter-class festivals report the same direction (six minutes to under two). This is the formal validation of this deployment's two unisex clusters (WC-05, WC-06) and of the worst-section display rule. And at the level of this venue, the 2024 edition's press record reports waits on the order of an hour at the worst WC clusters (Observador, 2024; Público, 2024), with additional queues at food, entry, and multi-hour post-show exit waits (Leak.pt, 2026). Summed across queue types — WC, bar, food, gates, exit — a festival-goer's cumulative daily waiting on a failure day plausibly reaches 1.5 to 2.5 hours. The 60-minute column of the grids is therefore not an optimistic ceiling; it is the middle of the documented range, and the extended rows of Table 18 price the upper range explicitly. Realistic targets differ by regime: 15–30 minutes/person of recovery is the credible objective for a normally loaded day; 60–120+ minutes is the recovery envelope on the failure days the 2024 record documents.

Every minute of average waiting removed at Rock in Rio scale is worth roughly thirteen thousand euros of experience value per day — the unit in which calibration effort, steward placement, and screen positioning should be priced. Tables 16–18 span the scenario grid: attendance from 50,000 to 150,000, recoveries from 5 to 60 minutes, and both value anchors (€10/h and €9/h); the shaded class marks this deployment's scale.

Attendance	5 min	15 min	30 min	60 min
50,000	€41.7k	€125.0k	€250.0k	€500.0k
80,000	€66.7k	€200.0k	€400.0k	€800.0k
100,000	€83.3k	€250.0k	€500.0k	€1,000.0k
150,000	€125.0k	€375.0k	€750.0k	€1,500.0k

Table 16 — Potential recovered value per day at €10/hour (assumption-based).

Attendance	5 min	15 min	30 min	60 min
50,000	€37.5k	€112.5k	€225.0k	€450.0k
80,000	€60.0k	€180.0k	€360.0k	€720.0k
100,000	€75.0k	€225.0k	€450.0k	€900.0k
150,000	€112.5k	€337.5k	€675.0k	€1,350.0k

Table 17 — Potential recovered value per day at €9/hour ticket-experience rate (assumption-based).

Attendance	Time saved/person	Value/hour	Potential value/day
80,000	15 min	€10/h	€200,000
80,000	30 min	€10/h	€400,000
80,000	60 min	€10/h	€800,000
100,000	15 min	€10/h	€250,000
100,000	30 min	€10/h	€500,000
100,000	60 min	€10/h	€1,000,000
80,000	90 min (all-queue failure day)	€10/h	€1,200,000
80,000	120 min (all-queue failure day)	€10/h	€1,600,000
100,000	120 min (all-queue failure day)	€10/h	€2,000,000
320,000 over 4 days	30 min	€10/h	€1,600,000 total
320,000 over 4 days	60 min	€10/h	€3,200,000 total
320,000 over 4 days	90 min (all-queue)	€10/h	€4,800,000 total

Table 18 — Scenario examples including the four-day festival totals (assumption-based scenarios of potential recovered value).

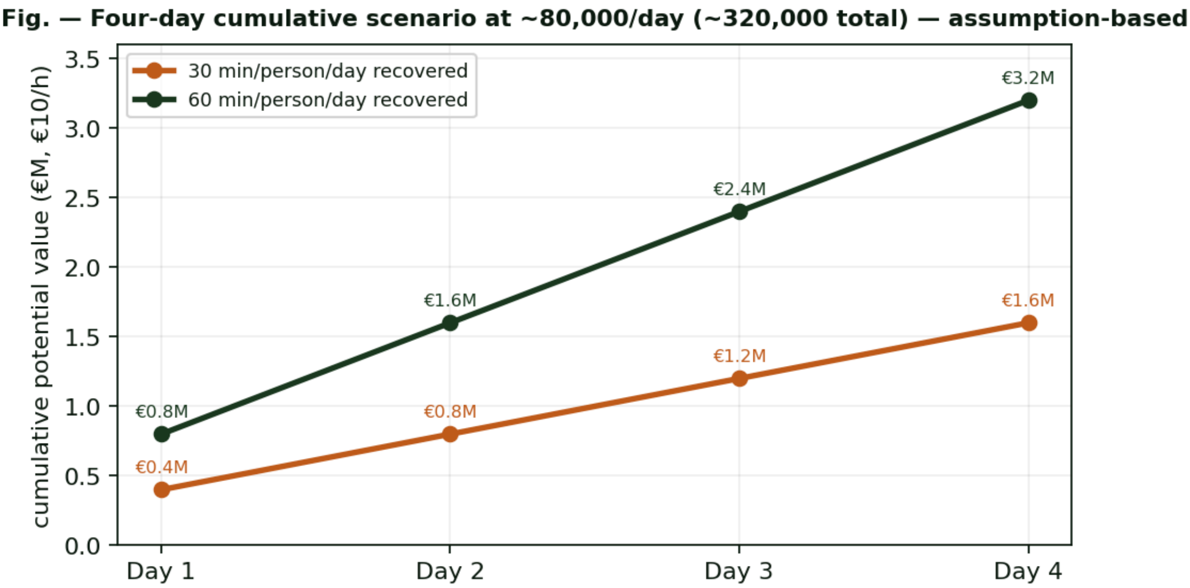

Figure 48 — Four-day cumulative value scenario at Rock in Rio scale (assumption-based).

What changes if Bathroom Intelligence is implemented by default: people spend less time in WC queues and more time at concerts and near food, drink, and merchandise; staff respond before queues become failures; cleaning is allocated by real pressure; underused WCs become visible and overused WCs are relieved; the festival uses existing infrastructure better; perceived quality improves; complaints and friction decrease; and the operator can reduce the need to overbuild capacity blindly.

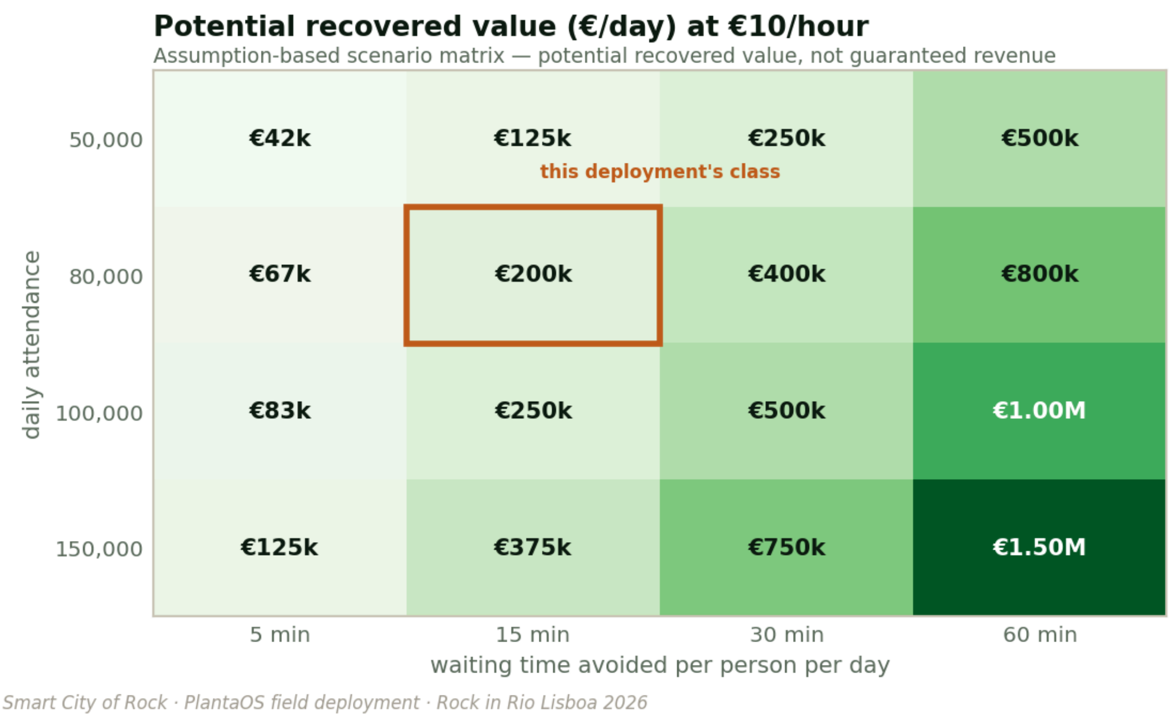

Figure 49 — Economic value heatmap at €10/hour across attendance and time-recovered scenarios.

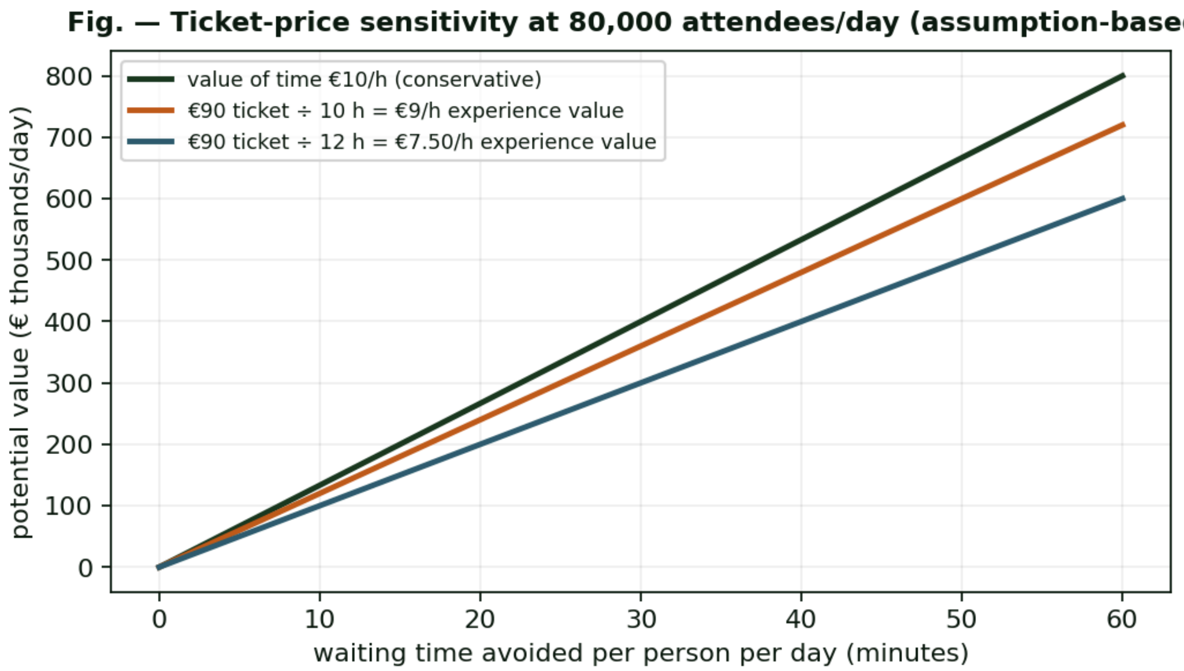

Figure 50 — Ticket-price sensitivity at 80,000 attendees/day: €10/h vs €9/h vs €7.50/h.

Secondary effects sit on top of the direct time value and are listed qualitatively, unquantified: more time near concerts; more time near food and beverage; more time near merchandise; less frustration; less staff escalation; better use of existing infrastructure; less overbuilding of capacity; better cleaning and steward allocation; and higher perceived festival quality. Honest caveats: the model assumes the system causally reduces waiting (plausible via redistribution, not yet measured against a controlled baseline), assumes uniform time valuation, and ignores displacement effects (a person routed to a farther cluster spends walking time). These caveats define the measurement agenda; they do not invalidate the direction of the estimate.

16.5 Perceived waiting: the information term

The psychology of waiting adds a term the engineering misses: waits of uncertain length feel longer than known, finite waits, and unexplained waits feel longer than explained ones (Maister, 1985; Larson, 1987). Model the experienced cost of a wait as

$$C\_perceived = \theta(I) \cdot \Delta t \cdot v, \qquad \theta\_uninformed > \theta\_informed \geq 1$$

where θ is a perception multiplier that decreases with information I. This term is why the system creates value even in the moments when it cannot yet shorten a queue: publishing live state converts an uncertain wait into a known one, and a known wait with a visible alternative into a choice. The app and screens attack θ directly; the routing attacks Δt. Both terms are real economics, and only one of them requires moving anybody.

16.6 Spend displacement and secondary value

Time in a queue is also time not spent at bars, food, and merchandise. With s the average on-site spend rate per attendee-hour, the displaced-spend model is

$$\Delta S = N \cdot \Delta t \cdot s$$

s is venue-confidential, so this paper leaves ΔS symbolic — but its sign is not in doubt, and it accrues to the operator rather than the attendee, which matters for who pays for the system. The remaining secondary effects — better staff allocation, cleaning dispatched by measured pressure rather than fixed rounds, higher perceived quality, safety margins at exits, better utilisation of existing infrastructure and hence reduced overbuilding, and sponsor-grade analytics — are directionally positive and deliberately left unpriced. The equations of this section price only time; everything else is margin.

The financial value of Wave Flow Control comes from recovering time before it is lost.

17. Transferability to Global Events

Nothing in the architecture is festival-specific: the primitive is anonymous counts of people and flows → calibrated pressure estimates → routing to available capacity. We therefore model transfer by scenario class rather than by naming other events, and we invent no numbers for venues we have not measured. Four scenario classes, read directly from Tables 16–18: a 50,000-person concert, an 80,000-person festival day (this deployment's class), a 100,000-person mega-festival day, and a 150,000-person extreme festival/stadium/city event. At global festival scale, even 5–15 minutes of recovered waiting time per person can represent six-figure daily value; at the largest scales and a full recovered hour, the envelope reaches seven figures per day. The same logic applies beyond events — stadiums, city centres and transit hubs, supermarkets, campuses, airports — wherever people queue for distributed capacity they cannot see.

17.1 From WC queues to all queues

Nothing in the estimator knows it is measuring toilets. Entrance pressure, interior validation, capacity thresholds, wave regimes, confidence, and counterfactual routing apply unchanged to every queue in the venue: bars, food outlets (the single largest complaint category alongside access in 2024 (Portal da Queixa, 2024; Público, 2024)), merchandise, lockers, entrances, and — critically — exits, where the 2024 edition's Stop&Go releases produced multi-hour waits for transport (Leak.pt, 2026; Observador, 2024). What would happen if the concept covered all queues? The venue would become a single load-balanced service system: every service point a cluster with a capacity and a waiting allowance, every attendee holding the counterfactual for every wait, operations seeing every wave front minutes early, and the venue's total waiting tax — currently paid invisibly across a dozen queue types — measured, priced, and systematically reduced. The WC deployment is the proof of mechanism; the mechanism is queue-generic.

The mapping, flow by flow. Porting the estimator to another festival flow is a change of capacity table and wave calendar, per flow type: Bars — service positions are taps/tills, service time $\tau \approx$ 30–60 s per order, waves synchronised to set breaks; entrance sensing at the bar frontage predicts the rush, interior sensing along the counter validates it. Food courts — stalls as servers, $\tau \approx$ 2–4 min, the pre-headliner dinner wave dominating; the 2024 record shows food among the largest complaint categories, making it the natural second module. Gates — lanes as servers with $\tau \approx$ 3–6 s per ticket read, one great morning ingress wave; entrance sensing far up the approach gives the longest lead time in the venue. Exits — egress lanes as servers, the single largest synchronised wave of the day, where the 2024 Stop&Go releases produced multi-hour transport waits; pressure sensing at exit approaches converts release timing from guesswork into measurement. Merchandise and lockers — tills/counters as servers with pronounced open/close waves; the WC-07 locker convergence already sits inside this deployment's geometry. Water refills and charging points — slow-service positions whose queues grow invisibly during heat peaks, directly coupled to the safety layer. Every one of these inherits the full stack unchanged: two-role (or single-sensor) geometry, one-minute smart-data means, five-minute publication, capacity normalisation, regime weights, confidence, and routing.

Fig. — Scenario-based scalability: even 5–15 minutes recovered represents six-figure daily value at scale (assumption-based)

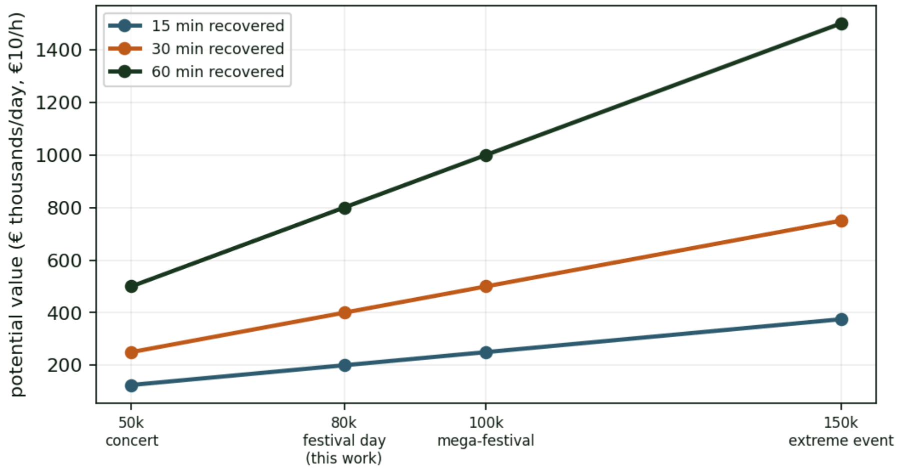

Figure 51 — Scenario-based scalability across event classes (assumption-based).

Fig. 11 — Transferable deployments of the Flow Intelligence layer

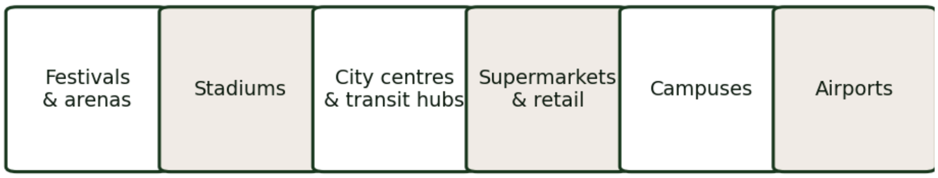

same primitive everywhere: anonymous counts of people and flows → pressure estimates → routing to available capacity

Figure 52 — Transferable deployments: festivals, stadiums, cities, supermarkets, airports, campuses.

17.2 The intelligence-first festival: structures follow flows

A sharper counterfactual than retrofitting an existing venue: suppose the festival is a Freedom Intelligence / Flow Intelligence festival first — the layer arrives before the structures, and the site plan is drawn under F = P/D from the first sketch. Three classes of savings open up, all scenario-based. First, capacity right-sizing: static provision ratios (of the 1:75 female / 1:400 male class) size for a worst case that measurement replaces — with wave-calibrated demand curves per zone, the portable-unit budget follows measured need; at major-festival scale (hundreds of rented units, at unit rental costs in the low hundreds of euros per event), even a 10–20% right-sizing is a five-figure saving per edition before a single queue forms, and it is exactly the smart-substitutes-for-mass mechanism of the deflationary-threshold and intelligent-Midas results (Melo de Magalhães, 2026f, 2026y). Second, placement by flows: clusters sited on the measured gate-stage-food-exit corridors rather than on leftover space — the Jensen gains of Section 16.3 designed in rather than recovered later, with the utilisation spread across clusters narrowed from the first hour of the first day. Third, operations designed to regimes: cleaning windows, steward rosters, and power logistics scheduled against the wave calendar instead of fixed rounds — capacity offline when demand is light, fully available when the wave arrives. The retrofit deployment documented in this paper recovers value from a fixed layout; the intelligence-first festival captures the same value plus the capital it would otherwise have buried in overbuilt, mis-placed provision. Structures follow flows.

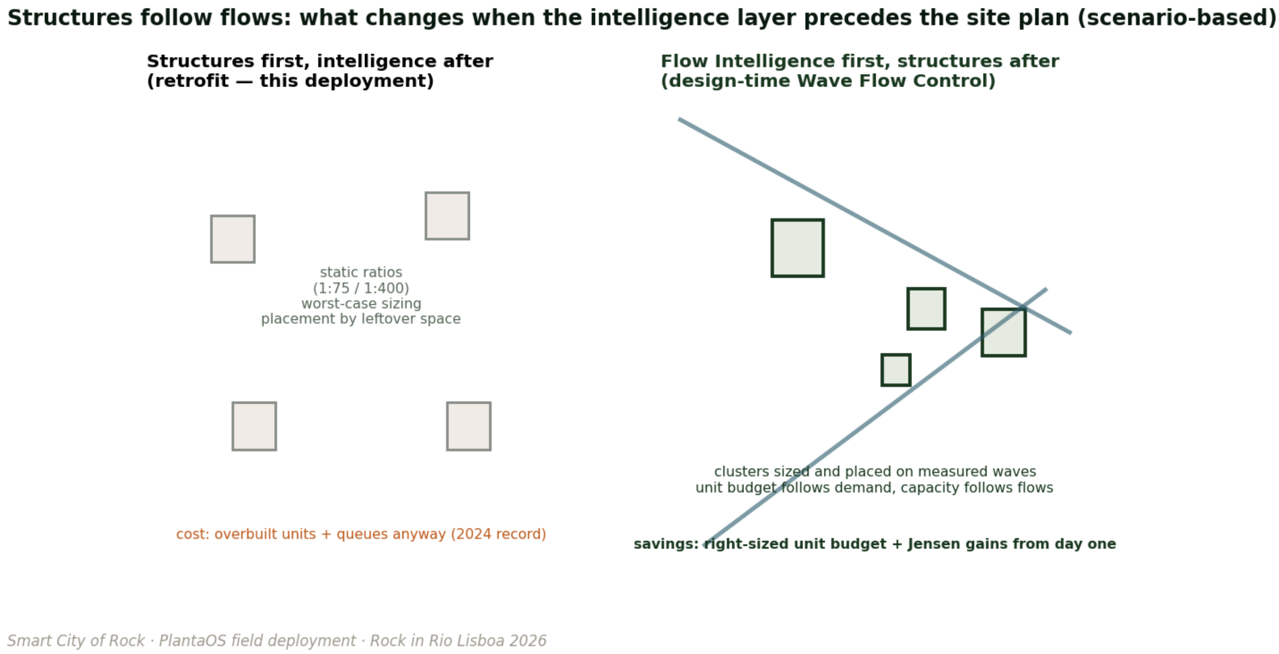

Figure 53 — Structures follow flows: the intelligence-first festival sizes and places capacity on measured waves (scenario-based).

17.3 Target clients for Wave Flow Control

The following are target segments — explicitly not existing clients; named organisations are archetypes of scale, not references. Festivals and music events (Rock in Rio-type, Primavera Sound-type, Tomorrowland-type, Glastonbury-type festivals, large concert promoters, temporary city-scale event operators): they have queues, toilets, bars, food courts, entrances, exits, and temporary infrastructure — the full wave set. Stadiums and sports clubs (Real Madrid-class stadiums; Benfica, Sporting, and FC Porto-class venues; UEFA/FIFA event operators; arena operators): stadiums have half-time waves, entrance waves, exit waves, bathroom waves, bar waves, and security bottlenecks — the same waves on a weekly cadence. Smart cities and municipalities (Lisboa, Porto, Cascais, Braga, and smart-city programmes generally): cities need to understand flows in parks, waterfronts, public events, transport hubs, markets, beaches, and tourism zones. Airports and transport hubs: bathroom queues, security queues, boarding waves, baggage waves, and retail congestion. Universities and campuses: class-change waves, cafeteria queues, library occupancy, event peaks. Retail and supermarkets: checkout waves, parking waves, footfall and occupancy. Hospitals and public services: waiting rooms, triage, cafeterias, elevators, service queues — all flow systems. Real estate and smart buildings: occupancy, shared areas, bathrooms, elevators, and cleaning optimisation. In every segment the estimator is unchanged; only the capacity table and the wave calendar differ.

Any place with people, waiting, capacity, and invisible alternatives is a client for Wave Flow Control.

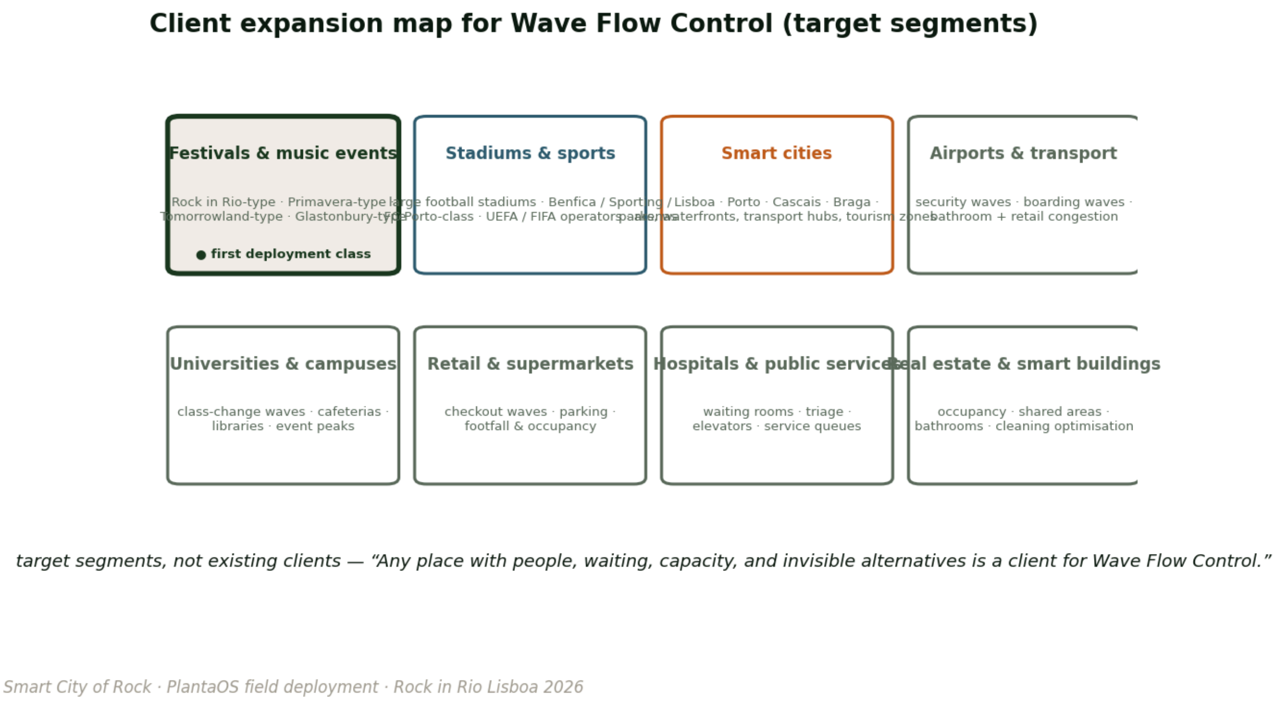

Figure 54 — Client expansion map for Wave Flow Control: target segments, not existing clients.

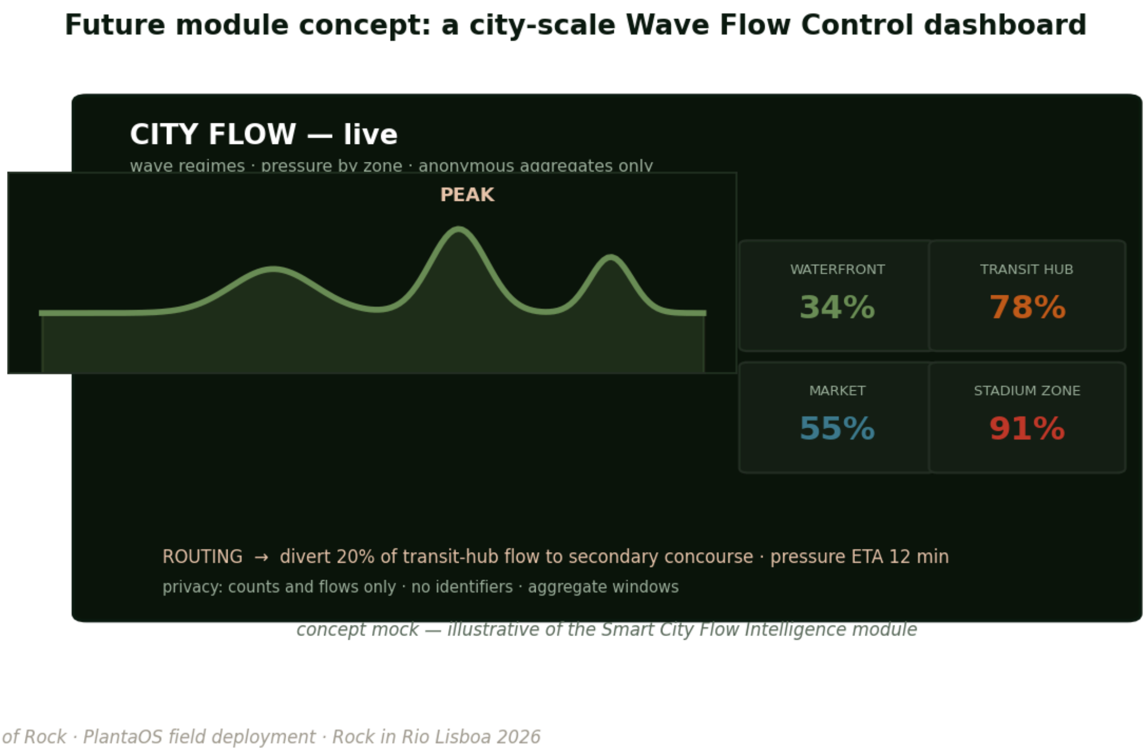

Figure 55 — Future module concept: a city-scale Wave Flow Control dashboard (concept mock, not a shipped product).

18. Sustainability

Energy and materials. The deployment's energy architecture is itself a sustainability and risk decision. Each node draws on the order of a watt from a standard USB power bank at 5 V, swapped on a rolling ~10-hour cycle from a small charged spare pool — a two-minute steward task per node. The consequences compound: no mains cabling through crowd areas (no trenching, no permits, no trip hazards), no diesel generation for the sensing layer, no work at height beyond 2–3 m mounting, and hardware that is recovered at end of event and redeployed at the next one. The power banks, nodes, gateways, and ABS enclosures are a reusable kit, not per-event waste; the marginal material footprint of instrumenting an additional festival approaches zero. The risk profile follows the same logic (Section 7.6): low voltage, no heat concentration, no moving parts — the residual risk class of a battery-powered consumer device, which is also why the insurance and safety-compliance burden stays minimal. Sustainability and operational risk are usually a trade-off; the power-bank architecture buys both at once.

Flow intelligence is a sustainability instrument because it improves the use of infrastructure that already exists. The conventional response to queues is supply-side: more units, more signage, more staff, sized for a worst case that lasts ninety minutes per day. Real-time redistribution attacks the same problem from the demand side, filling under-used clusters before anyone builds new ones. Concretely, it can reduce overbuilding of sanitary capacity, unnecessary static signage, staff guesswork, wasted attendee walking, crowd friction, and inefficient spatial distribution of demand. Sustainable value here is precisely better use of what already exists.

19. Coordination Playbook for Future Deployments

The following playbook condenses the deployment into a replicable schedule.

Before the phased schedule, the operational core fits in ten steps: (1) map the venue; (2) define all WC clusters and service points; (3) classify sections by type — male, female, unisex; (4) place entrance and interior sensors per section; (5) geotag every node; (6) set initial radius/sensitivity bands per role (Table 5); (7) run calibration through Light, Growth, Peak, and Post-concert; (8) publish real-time cluster state to the staff dashboard and to public screens/app; (9) compute 5-minute rolling averages for stable external reporting; and (10) use routing logic to redistribute demand. Everything else in this section is the professionalisation of those ten steps.

Phase	Actions
T−90 days	Obtain official site plan; identify all WC clusters; define canonical markers; estimate section capacities; define integration requirements (data reporting consumers, screen/app surfaces).
T−30 days	Prepare LilyGo LoRa32/ESP32 nodes; configure ranges and RSSI sensitivity levels per role (Table 5); test power banks for full-day missions; test LoRa communication and gateway coverage; prepare dashboards; prepare app routing logic; define SCOR schema and cadence.
T−7 days	Physically validate all WC coordinates; place entrance and inside sensors per section; verify GPS and detect coordinate outliers; test the public app end-to-end; test SCOR data reporting; define calibration defaults per role.
Event day	Install early; run baseline calibration; use the Light period for baseline; use Growth for first wave detection; recalibrate before Peak; run Peak mode during concerts (negative-calibration authority active); switch to Post-concert mode after headline shows; monitor offline/degraded sensors; export five-minute averages; public screens/app live.
Post-event	Verify view logs; export sensor logs; compare with staff observations; compute estimated queue pressure per cluster; document failures honestly; improve calibration; publish the anonymised dataset if possible.

Table 19 — Deployment coordination playbook.

19.1 Flow Intelligence for Festivals by Planta Smart Homes: how this is sold as a whole

The festival-facing product is named Flow Intelligence for Festivals by Planta Smart Homes — the festival module of the Planta Wave Flow Control platform: a deployable intelligence layer for large events and physical spaces. Its first module is WC / Bathroom Flow Intelligence, the system this paper documents; the roadmap modules are Bar, Food, Gate, Parking, Transport, Security, Emergency, Retail, and Smart City Flow Intelligence — each the same estimator with a different capacity table and wave calendar (Section 17.1). What is sold: pre-event venue mapping; the sensor kit; the app/screen layer; the operations dashboard; API/SCOR integration; real-time recommendations; the post-event report; data-driven operational improvements; and annual or per-event licensing. The festival sales proposition fits in one sentence: "Give us the site map and WC locations. We install a sensing and routing layer that helps reduce queues, redistribute demand, and recover experience time."

19.2 What a festival hires — and what it does not

The engagement is deliberately simple on both sides of the line. What the festival hires: a per-event licence for an outcome layer — N instrumented WC clusters with section-aware sensing; the calibration service (Day 0 plus live calibration authority through every event day); the public routing layer (app of the plantarockinrio.com class plus venue screens); the operations dashboard with pressure alerts; SCOR/KPI export; and the post-event report with the full aggregate dataset. What the festival does not hire — and this is half the sales argument: no permanent infrastructure and no civil works; no cameras and no personal-data processing liability (aggregate-only by physical incapacity, Section 12); no replacement of its stewards, cleaners, or security — the system informs them; no accuracy warranty on counts, because none would be honest (Section 22) — the deliverable is decision-grade ranking, freshness, and stated confidence, not audited headcounts; and no lock-in: the kit leaves when the event ends. Procurement runs in five steps: (1) the festival supplies the site map, WC locations, and section capacities; (2) Planta returns a cluster-count proposal priced per instrumented cluster per festival day against the value scenarios of Tables 16–18; (3) install at T−7; (4) Day 0 calibration (19.3); (5) event operations under the daily protocol (19.4), closing with the report. "Give us the site map and WC locations" is not a slogan; it is literally step one of the contract.

19.3 Day 0: the calibration day, and the prior study it consumes

One full day before doors is reserved for calibration, and it is the highest-leverage day of the deployment. Day 0 consumes a prior study assembled with the festival in the preceding weeks — all of it approximations, none of it accuracy, exactly as the estimator itself: lineup-based crowd estimates per hour (which headliners, on which stages, pulling how many people toward which zones); expected movement patterns (gate-to-stage corridors, stage-to-food loops, exit routes); cleaning schedules per cluster (a cabin block under cleaning is capacity offline — it must enter the capacity table as a time-varying term); section capacities and waiting limits (Tables 2 and 6); and the wave calendar mapping the published programme onto the four regimes. From these, per-cluster presets are loaded: initial radii, divisors, weights, and cleaning windows.

The calibration itself can be executed at four quality levels, in descending order of preference and cost: Option A — press-day live calibration: a real crowd of thousands (Rock in Rio's press event is the natural vehicle) exercises the estate; staff spot counts at every cluster anchor the divisors; the best option because the radio environment is real. Option B — staff-wave rehearsal: 50–150 staff walk scripted waves through each cluster — entrance-only passes, interior dwell, mixed flows — giving controlled ground truth at low density; calibrates the geometry and roles but not crowd attenuation. Option C — synthetic device testing: crates of phones/beacons at measured distances validate the RSSI bands per node; cheapest, calibrates radii only. Option D — historical-analog presets: load the per-cluster presets of Table 10 from a comparable prior deployment and treat the first Light regime of Day 1 as the calibration window. Options compose: a well-run Day 0 is C in the morning, B at midday, A if a press event exists — with D always loaded as the fallback floor. What

Day 0 explicitly does not attempt is accuracy: its exit criterion is that every section's estimate is ordinal-correct and directionally responsive under test flows, with confidence honest about the rest.

19.4 The four-phase daily calibration protocol

During the event, calibration is not an installation-time constant; it is a daily operational cycle synchronised to the wave regimes. Table 20 states the protocol as run: one owner, logged changes, automation off during shows.

Phase	Window	Operator actions	Checks
Light	13:00–16:30	load Day presets; verify all nodes online and fresh; interior-dominant weights (≈30–40/60–70); high divisors (3–5); baseline drift review at 15:00	zero-crowd sanity: no section above ‘low’ without vis cause; flag any node whose baseline moved since yesterday
Growth	16:30–18:30	switch to entrance-dominant weights (≈60–70/30–40); divisors to growth band (2.3–2.9); first staff spot-count round; adjust per-cluster divisors where observation disagrees	first-wave detection confirmed at gate-adjacent clus entrance leads interior in every gendered pair
Peak	18:30–21:30	weights 70–80/20–30; peak divisor band (1.7–0.9 — sub-unity only under confirmed crowd attenuation, Section 9.5); calibration freeze except logged single-command corrections; stewards dispatched by pressure alerts (90+)	published values plausible against staff observation the three focus clusters; offline discipline verified — stale values republished
Post-concert	21:30–close	weights 70–85/15–30; divisors 1.9–2.6; screens re-rank on the show-end wave; exit-adjacent clusters watched; begin day-close export	wave detected at entrances before interior saturatio day log complete: every divisor/radius change with t owner, reason

Table 20 — The four-phase daily calibration protocol (one owner, logged changes, automation off during shows).

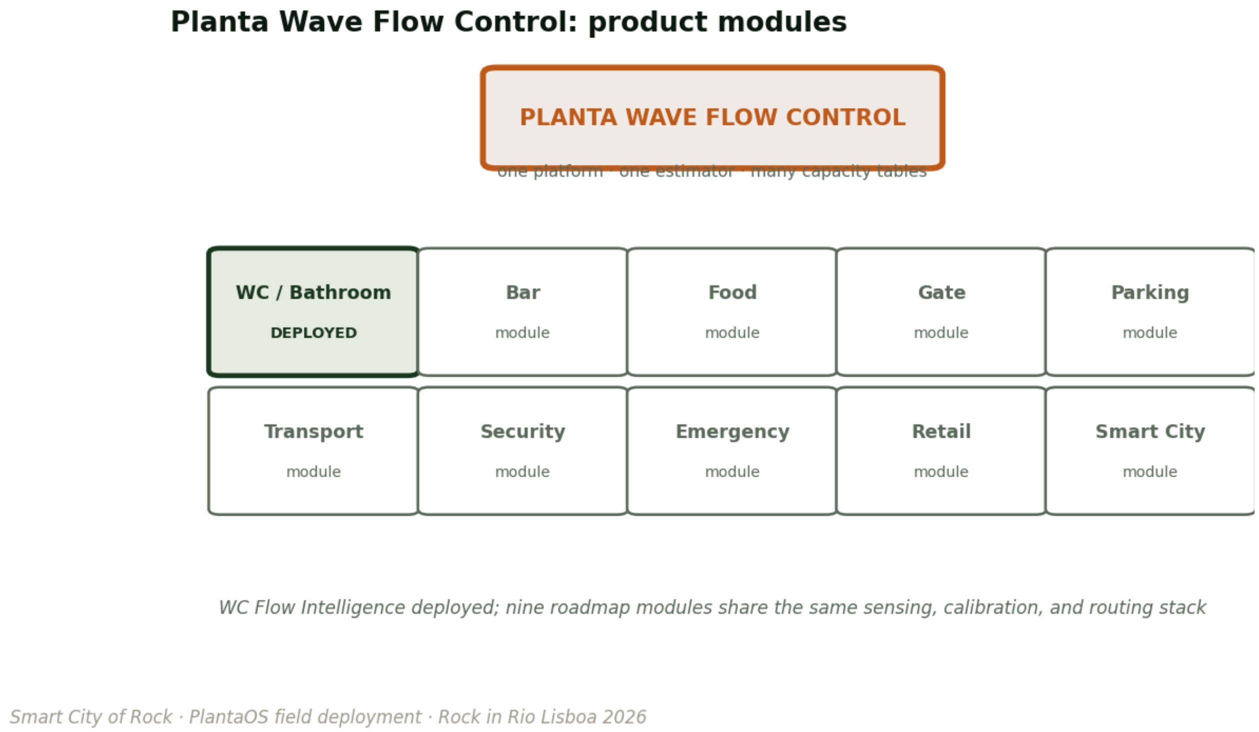

Figure 56 — Planta Wave Flow Control product modules: WC Flow Intelligence deployed; nine roadmap modules on the same estimator.

The offer is not a sensor purchase; it is an outcome contract with six inseparable layers — hardware kit per cluster (nodes per role, gateways, power banks on the 10-hour swap cycle, enclosures, mounting); the backend and calibration engine (ingestion, divisors, radii, wave regimes, confidence, offline discipline, and negative-calibration authority); the public layer (attendee app of the plantarockinrio.com class plus venue screens, routing to the best WC rather than the nearest); the operations layer (dashboards, queue-pressure alerts, steward and cleaning dispatch by measured pressure); the data layer (SCOR/KPI integration and auditable five-minute aggregate export); and the service layer (the T−90 playbook, on-site calibration through the event, and the post-event dataset and report). Selling any layer alone reproduces a failure mode this paper has already documented: sensors without calibration authority overestimate; an app without measurement is 2024's WhatsApp channel; dashboards without the public layer inform operators while attendees still queue blind. The commercial unit is the instrumented cluster per festival day; the price anchors against the recovered-value scenarios of Tables 16–18 and the smart-gets-cheaper logic of the wider programme (Melo de Magalhães, 2026y) — the buyer is not purchasing hardware, but the measured reduction of a waiting tax it is currently paying invisibly.

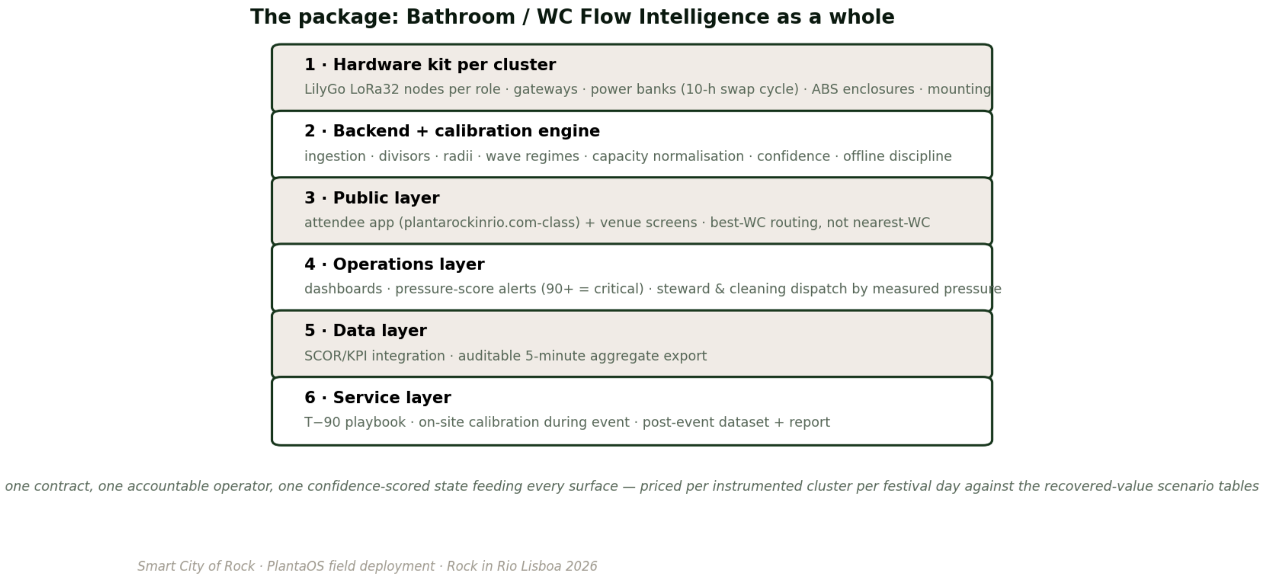

Figure 57 — The package: six inseparable layers of the Bathroom / WC Flow Intelligence offer.

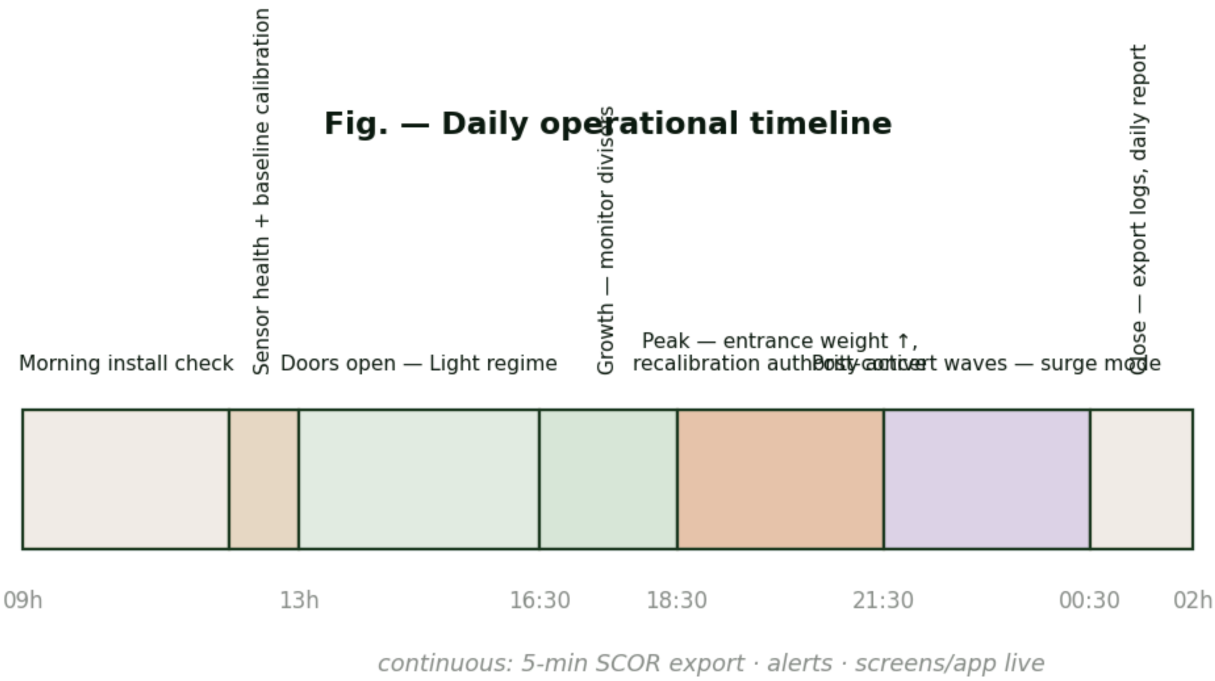

Figure 58 — Daily operational timeline from morning install to close-of-day export.

#	Sensor deployment checklist
1	Node flashed with correct cluster/section/role identity
2	Radius and divisor set per role (entrance tighter than interior)
3	Non-metallic enclosure; mounted 2–3 m, oriented to sensed zone
4	Power bank charged and sized for full day + margin; battery state registered
5	GPS captured, graded, checked against cluster median
6	First samples visible in ingestion within 5 minutes of mounting
7	Entrance–interior pair agreement sanity-checked before doors

Table 21 — Sensor deployment checklist.

#	Failure-handling checklist
1	Stale sensor → excluded from fusion automatically; alert to operations
2	Empty 5-min window → estado_sensor = offline, confianca_cruzada = 0
3	Strong entrance–interior disagreement → confidence penalty; investigate placement
4	Suspect geolocation → excluded from canonical point; flagged, never silently corrected
5	Overestimation observed → negative-calibration cycle (divisor ↑, radius ↓, weights ↓)
6	Never backfill; never republish old values as current; never manufacture confidence

Table 22 — Failure-handling checklist.

20. Scientific Novelty

In what sense is this a new technology? The question deserves a precise answer, because 'new' is the most abused word in technology writing. Not new: radio presence sensing (Musa & Eriksson, 2012; Schauer et al., 2014), LoRa backhaul (Augustin et al., 2016), queueing theory (Little, 1961; Kingman, 1961), or festival apps. New, to our knowledge, as a deployed combination under one operator at a major festival: section-aware sensing with per-role radio geometry; field-executed negative calibration with regime-dependent divisors and weights; capacity-normalised, confidence-reweighted pressure fusion; privacy by physical incapacity rather than policy; and a closed public routing loop (app + screens) driven by that state in real time. Each element is defensible alone; the integrated, calibrated, closed-loop layer is the category claim of Section 10 — and it is a Planta Smart Homes technology, productised as Flow Intelligence for Festivals by Planta Smart Homes on the Planta Wave Flow Control platform, built on the Freedom Intelligence theoretical programme (Melo de Magalhães, 2026a–2026ab). The distinction from every alternative in Tables 13a–15 is structural, not incremental: the alternatives measure or inform; this layer measures, calibrates, and routes.

20.1 Why Freedom Intelligence, specifically

A fair question is whether the theory earns its place — whether any competent engineering team without it would have built the same system. The design record says no, because the theory dictated the non-obvious choices. $F = P/D$ says value comes from raising perception or lowering distortion, never from adding mass: hence routing over construction, and the convexity economics of Section 16.3. The perception term dictated that invisible capacity is the real inventory: hence the counterfactual panels, and best-WC over nearest-WC. The distortion term dictated that uncertainty is a first-class cost: hence confidence scores, offline honesty, and the $\theta(I)$ information term of Section 16.5 — an engineering team optimising counts would have optimised accuracy instead, and Section 15.1 argues that is the wrong objective. The wave theses dictated regime-switching calibration rather than a single tuned constant. And Thesis V — intelligence as making distortion lighter — is, sentence for sentence, the product specification of the screens. The theory is not decoration on the system; it is the reason the system has the shape it has.

The novelty of this work is not one sensor, one dashboard, or one app — each component exists somewhere in prior art. The novelty is the integrated mechanism, operating simultaneously in a live high-density environment: section-aware WC topology; gender/unisex logic; entrance/interior sensing roles; wave-based time regimes; negative calibration; GPS outlier rejection; five-minute rolling averages; confidence scoring; aggregate privacy-preserving data reporting; counterfactual routing; public recommendations; operational dashboards; and SCOR/KPI integration. Removing any element degrades the whole: without section awareness, routing is unfair; without entrance/interior pairing, waves are seen too late or passers-by become occupants; without negative calibration, estimates run away at exactly the peak moments that matter; without confidence scoring and offline discipline, trust collapses; without aggregation, the privacy envelope breaks; without counterfactual routing, measurement never becomes freedom. Freedom Intelligence explains why this particular combination is the right one: every element either raises perception or lowers distortion, and the wave model dictates when each matters most.

21. Lessons Learned

1. Calibrate down, not up. Every naive configuration overestimated. Budget authority and time for on-site negative calibration, and treat it as method, not failure.
2. Freshness beats continuity. Publishing offline with zero confidence when data is stale preserves trust; repeating the last value destroys it.
3. Two roles are the minimum viable geometry. Entrance-only sensing books passers-by as occupants; interior-only sensing sees the wave only after it is a queue. The entrance/interior pair is the unit of deployment.
4. Canonical points are a product decision, not a GIS detail. Users trust one pin per facility; the pipeline must actively suppress sensor scatter and flagged outliers.
5. Geotags lie under festival conditions. Two of 24 positions were materially wrong; distance-to-median plus topology comparison caught both before they reached the public map.
6. Section structure is operationally decisive. Male, female, and unisex sections queue differently; fair routing requires modelling them separately.
7. Private radio was the correct bet. Public mobile networks were effectively unavailable at exactly the peak moments data reporting mattered most.
8. What was good, explicitly: the two-role topology held; the five-minute stability layer held; the offline discipline preserved trust through degradation; the private backhaul kept reporting when public networks failed; the per-phase divisor logic gave operators a one-command response to every observed bias; and the day-1 degradation was recovered in the field rather than hidden.
9. For the next festivals: deploy the full 28-role topology with spares from day 0 and keep ~15-active as the floor, not the plan; add the missing WC-05 entrance role first; load the per-cluster presets of Table 10 as defaults; run a density rehearsal (press day) before the first public wave; define the analytics metrics (users vs views) in writing before doors; and keep divisor governance exactly as it ended — one owner, logged changes, automation off during shows.

10. LoRa range is not sensing radius. Keeping the two radii separate — kilometres vs metres — prevented a class of silent estimation errors.

22. Limitations

- No continuous ground truth: accuracy is argued from consistency and staff spot checks; no accuracy percentage is claimed.
- Uneven spatial coverage: WC-05 has one interior point and no entrance sensing; WC-01 is effectively one spatial observation; confidence is capped and flagged accordingly.
- Geolocation defects: two flagged outliers remain unresolved pending physical re-verification; the venue polygon is a draft pending official GIS; GPS coordinates carry the festival-condition error budget of Section 6.2 and are not presented as perfect.
- Device-to-people uncertainty: divisors absorb but do not eliminate devices-per-person ambiguity; its demographic and temporal variation at this venue is uncharacterised.
- Causal effect unmeasured: the economic model assumes the system reduces waiting; no controlled on/off comparison has been run.
- Single-venue evidence: transfer claims are architectural arguments, not demonstrated results.
- Usage figures unverified: the reported 35,000–40,000 views remain pending final log verification and are not treated as a result.
- Sensing-range discipline: meaningful range beats maximum range — ranges above ~20–25 m overcount surrounding movement; the RSSI levels of Table 5 are role-bound, and zero sensor error is not claimed.
- Theory status: the Law of Freedom and the FLRP layering are theoretical/design-science models that organised the system's design; they are not proven physical laws, and the conceptual figures derived from them (freedom collapse, before/after operational freedom) are models, not measured data.
- Scope discipline: 1,137 (verified cabins) and ~1,700 (broader planning scope) are different quantities; anonymous aggregate counts are not the same thing as raw device tracking; LoRa communication range is not sensing radius. Wherever these pairs could be confused, this paper separates them explicitly.

23. Future Work

1. Ground-truth subset: instrument one cluster with continuous independent counting for a full day to quantify estimator accuracy per regime.
2. Controlled value measurement: alternate routing recommendations on/off across comparable windows to estimate the causal effect on queue length and dwell time.
3. Complete the sensing geometry: add the missing WC-05 entrance point; re-verify the two flagged outliers; resolve near-duplicate pairs.
4. Learn the time-dependent weights we(t), wi(t) and confidence function C(s,t) from data rather than fixing them by regime.
5. Verify reported usage against server logs and publish the reconciled figure.
6. Multi-venue replication — stadium and transit-hub pilots — to convert Section 17's architectural transfer argument into evidence.

7. Test the framework's central prediction directly: interventions on visibility and distortion should measurably change flow distribution across clusters — a falsifiable claim the current deployment motivates but does not settle.
8. Publish the anonymised cluster-level dataset and calibration configuration for reproducibility.

23B. Why This Matters: Liquid Innovation, Rock World, and Everyone

For Liquid Innovation Co. Liquid Innovation curates the Smart City of Rock — the mandate to make the City of Rock genuinely intelligent rather than merely branded so. A calibrated, privacy-preserving flow layer is a concrete instance of that mandate delivered at scale: it converts “smart city” from a theme into a measured system, gives Liquid a repeatable module to bring to future editions and partners, and demonstrates that the innovation programme can commission and ship real infrastructure intelligence, not only activations. The bathroom-intelligence commission was the entry point precisely because it is the hardest, most complained-about, least glamorous flow — solving it is the strongest possible proof that the Smart City of Rock is real.

For Rock in Rio and Rock World. For the operator, the queue is a direct threat to the product: the 2024 record shows waiting as a top complaint category, and every minute lost to a line is a minute subtracted from the show the ticket was bought for. A flow layer defends the core asset — attendee experience — and does so on the operator’s own terms: no cameras, no personal-data liability, no permanent civil works, no replacement of existing staff, and a kit that leaves when the event ends. It turns spatial blindness into an operations dashboard, lets stewarding and cleaning be dispatched by measured pressure instead of guesswork, and yields a post-event dataset that informs next year’s site plan. In economic terms it defends the recovered-value envelope of Section 16; in brand terms it makes the festival visibly more caring about the parts of the experience that usually go unmanaged.

For everyone else. The reason this matters beyond one festival is that waiting is a near-universal, largely unmeasured tax on human time. Stadiums, airports, hospitals, campuses, supermarkets, transport hubs, and cities all impose it, all mostly blind to it, and all constrained — rightly — by privacy limits that rule out surveillance-grade counting. A method that recovers time through aggregate, privacy-preserving measurement and routing, rather than through more construction or more cameras, is a general instrument for giving people back their time. Framed through Freedom Intelligence, the stake is larger than convenience: waiting is a partial loss of navigable freedom, and a system that reduces it returns freedom — measured, honestly, at scale. That is why the case study is a bathroom at a festival and the claim is about physical space in general.

24. Conclusion

What was shown, and where. The claims of this paper trace to their evidence as follows. That queues are freedom collapses and crowds move in waves: Sections 2 and 10, with the four-regime model field-operated per Table 20. That the estimator works as designed — entrance predicts, interior validates, capacity normalises, confidence reweights: Sections 9.2–9.7, with the controls visible as shipped software in Section 13.3. That calibration must be negative, phased, and per-cluster: Sections 9.3–9.6 and the Day 0 / daily protocol of Sections 19.3–19.4. That privacy by incapacity is deployable at 80,000-person scale: Sections 7 and 12, with the client-side location policy verbatim in the product. That the economics are large and lower-bounded: Section 16's six equation families, anchored to the ticket price, the European value-of-time meta-analysis, and the

2024 waiting record of this very venue. That redistribution alone creates most of the value: the convexity argument of 16.3 and the priced vignettes of 14.1. That this is, to our knowledge, a new deployable category: Section 20's precise statement of what is and is not new. And that a festival can buy it as a bounded, per-event engagement with an explicit list of what it is not buying: Section 19.2. Every claim excluded by the paper's discipline — exact queue reduction, accuracy percentages, guaranteed revenue — is catalogued in Sections 13 and 22.

This paper connected three levels. Freedom Intelligence explains the problem: action depends on perceived path availability, and freedom collapses when paths become invisible, distorted, congested, or unnavigable. Flow Intelligence implements the solution: a privacy-preserving cyber-physical layer that senses, calibrates, averages, interprets, and guides flows without identifying or tracking anyone. The wave model explains why timing matters: pressure arrives in programme-synchronised waves, the peak is the decisive moment, and a system that works only in the light period is not a Flow Intelligence system. Rock in Rio Lisboa demonstrates the system under real high-density pressure: approximately 80,000 people per day, four days, 145,000 m², eight WC clusters, a verified 1,137 cabins, GPS uncertainty, saturated networks, and strict privacy constraints — simultaneously.

A high-density event is a live freedom system. The more distortion it accumulates, the more people wait, crowd, guess, and lose options. Flow Intelligence reduces distortion by making hidden paths visible. In this sense, the Rock in Rio deployment is not only a WC-routing case study; it is an early demonstration of how physical spaces can become more free when they become more intelligent. The queues were always there. What changed is that they could be seen — early, honestly, and anonymously — and that, for the first time at this venue, every person standing in one had a counterfactual. In the vocabulary of the wider programme (Melo de Magalhães, 2026a, 2026p, 2026q): the physical paths of Parque Tejo were fixed, but the Paths never were — and this deployment is what it looks like when a venue is engineered so that no one's option space looks smaller than it is. The innovation is not only sensing. The innovation is orchestration under uncertainty: using calibrated sensor range, gender/unisex topology, entrance/inside sensing, wave regimes, negative calibration, five-minute averages, app routing, privacy-preserving data reporting, and economic reasoning to turn invisible WC queues into real-time operational intelligence. The festival of the future is not only programmed — it is orchestrated. The final thesis of this paper is that Planta Smart Homes created what is, to our knowledge, a new deployable technology category: Wave Flow Control for high-density physical spaces. The Rock in Rio Lisboa deployment shows how Freedom Intelligence becomes operational through Flow Intelligence — sensing invisible waves, reducing distortion, creating better paths, and turning temporary event infrastructure into an adaptive system.

Acknowledgments

The author thanks the Rock World operations and technical teams, Liquid Innovation, Sensaway/SCOR, and the festival security and stewarding staff whose independent counts supported cross-validation.

References

References follow APA (7th ed.) style.

Article 29 Data Protection Working Party. (2014). Opinion 05/2014 on anonymisation techniques (WP216). European Commission.

Abrantes, P. A. L., & Wardman, M. (2011). Meta-analysis of UK values of travel time: An update. Transportation Research Part A: Policy and Practice, 45(1), 1–17.

Augustin, A., Yi, J., Clausen, T., & Townsley, W. M. (2016). A study of LoRa: Long range & low power networks for the Internet of Things. Sensors, 16(9), 1466.

Batty, M., Axhausen, K. W., Giannotti, F., Pozdnoukhov, A., Bazzani, A., Wachowicz, M., Ouzounis, G., & Portugali, Y. (2012). Smart cities of the future. The European Physical Journal Special Topics, 214, 481–518.

Becker, G. S. (1965). A theory of the allocation of time. The Economic Journal, 75(299), 493–517.

Blondel, V. D., Decuyper, A., & Krings, G. (2015). A survey of results on mobile phone datasets analysis. EPJ Data Science, 4, Article 10.

Bor, M. C., Roedig, U., Voigt, T., & Alonso, J. M. (2016). Do LoRa low-power wide-area networks scale? In Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (pp. 59–67). ACM.

Cavoukian, A. (2009). Privacy by design: The 7 foundational principles. Information and Privacy Commissioner of Ontario.

Centenaro, M., Vangelista, L., Zanella, A., & Zorzi, M. (2016). Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios. IEEE Wireless Communications, 23(5), 60–67.

de Montjoye, Y.-A., Hidalgo, C. A., Verleysen, M., & Blondel, V. D. (2013). Unique in the crowd: The privacy bounds of human mobility. Scientific Reports, 3, Article 1376.

DeSerpa, A. C. (1971). A theory of the economics of time. The Economic Journal, 81(324), 828–846.

Draghici, A., & van Steen, M. (2018). A survey of techniques for automatically sensing the behavior of a crowd. ACM Computing Surveys, 51(1), Article 21.

Dwork, C. (2006). Differential privacy. In Automata, languages and programming: ICALP 2006 (pp. 1–12). Springer.

European Parliament and Council. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union.

Events Industry Forum. (n.d.). The purple guide to health, safety and welfare at music and other events.

Franke, T., Lukowicz, P., & Blanke, U. (2015). Smart crowds in smart cities: Real life, city scale deployments of a smartphone based participatory crowd management platform. Journal of Internet Services and Applications, 6, Article 27.

Fruin, J. J. (1971). Pedestrian planning and design. Metropolitan Association of Urban Designers and Environmental Planners.

González, M. C., Hidalgo, C. A., & Barabási, A.-L. (2008). Understanding individual human mobility patterns. Nature, 453, 779–782.

Health and Safety Executive. (1999). The event safety guide (HSG195). HSE Books.

Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. Nature, 407, 487–490.

Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. Physical Review E, 51(5), 4282–4286.

Kingman, J. F. C. (1961). The single server queue in heavy traffic. Mathematical Proceedings of the Cambridge Philosophical Society, 57(4), 902–904.

Kleinrock, L. (1975). Queueing systems, Volume 1: Theory. Wiley.

Larson, R. C. (1987). Perspectives on queues: Social justice and the psychology of queueing. Operations Research, 35(6), 895–905.

Leak.pt. (2026). Fui ao Rock in Rio e tenho uma mão cheia de queixas. https://www.leak.pt

Little, J. D. C. (1961). A proof for the queuing formula: $L = \lambda W$. Operations Research, 9(3), 383–387.

Maister, D. H. (1985). The psychology of waiting lines. In J. A. Czepiel, M. R. Solomon, & C. F. Surprenant (Eds.), The service encounter (pp. 113–123). Lexington Books.

Martin, J., Mayberry, T., Donahue, C., Foppe, L., Brown, L., Riggins, C., Rye, E. C., & Brown, D. (2017). A study of MAC address randomization in mobile devices and when it fails. Proceedings on Privacy Enhancing Technologies, 2017(4), 365–383.

Mekki, K., Bajic, E., Chaxel, F., & Meyer, F. (2019). A comparative study of LPWAN technologies for large-scale IoT deployment. ICT Express, 5(1), 1–7.

Melo de Magalhães, G. (2026a). The architecture of Freedom Intelligence. Zenodo. https://doi.org/10.5281/zenodo.18636095

Melo de Magalhães, G. (2026b). The architecture of Freedom Intelligence: A design theory of path availability. SSRN. https://doi.org/10.2139/ssrn.6298638

Melo de Magalhães, G. (2026c). Beyond the data paradigm: Freedom Intelligence and the structural laws of navigability. Preprints. https://doi.org/10.20944/preprints202603.0429.v1

Melo de Magalhães, G. (2026d). Can there be a pathlessness world? Zenodo. https://doi.org/10.5281/zenodo.20030652

Melo de Magalhães, G. (2026e). Design as path engineering: Freedom, perception, and the architecture of possibility. Zenodo. https://doi.org/10.5281/zenodo.18773081

Melo de Magalhães, G. (2026f). Distortion geometry and the deflationary threshold: Why construction resists intelligence — and when it will not. Preprints. https://doi.org/10.20944/preprints202604.0183.v1

Melo de Magalhães, G. (2026g). The Einstein gap — A navigation law for spacetime, the five theses of the Architecture of Freedom Intelligence, and the engineering of smart space. Zenodo. https://doi.org/10.5281/zenodo.19842456

Melo de Magalhães, G. (2026h). A falsifiability-first investigation across classical physics, quantum mechanics, biology, computation, cognition, and the full mechanics of perception. SSRN. https://doi.org/10.2139/ssrn.6718778

Melo de Magalhães, G. (2026i). Freedom Intelligence for smart buildings: AFI-guided swarm optimisation of multi-domain comfort in industrial training facilities. SSRN. https://doi.org/10.2139/ssrn.6382981

Melo de Magalhães, G. (2026j). The general law of navigability: Perception as clustering of recursive possibility in the Architecture of Freedom Intelligence. SSRN. https://doi.org/10.2139/ssrn.6721198

Melo de Magalhães, G. (2026k). The geometry of bounded choice. Zenodo. https://doi.org/10.5281/zenodo.20073594

Melo de Magalhães, G. (2026l). Introducing perception in navigation systems. SSRN. https://doi.org/10.2139/ssrn.6712378

Melo de Magalhães, G. (2026m). Is the universe an operating system governed by freedom? Zenodo. https://doi.org/10.5281/zenodo.20054905

Melo de Magalhães, G. (2026n). Objectivity as stable path availability structure: Subjectivity, intersubjectivity, perceived freedom, and operative freedom in the Architecture of Freedom Intelligence. SSRN. https://doi.org/10.2139/ssrn.6712479

Melo de Magalhães, G. (2026o). Observer-dependent navigability in swarm intelligence: A path-theoretic decomposition of performance into perception and distortion. Preprints. https://doi.org/10.20944/preprints202603.0945.v1

Melo de Magalhães, G. (2026p). Path availability as the first condition of reality. Zenodo. https://doi.org/10.5281/zenodo.19661379

Melo de Magalhães, G. (2026q). The path availability TOE — Theory of everything of freedom. SSRN. https://doi.org/10.2139/ssrn.6612438

Melo de Magalhães, G. (2026r). Seeing clearly: Differentiating real possibilities to improve perception and choosing the path with the least structural distortion. Zenodo. https://doi.org/10.5281/zenodo.20031067

Melo de Magalhães, G. (2026s). Spatial freedom fields: A rigorous computational framework for the Architecture of Freedom Intelligence (AFI). Zenodo. https://doi.org/10.5281/zenodo.18845574

Melo de Magalhães, G. (2026t). The strange case of possibilities is recursion. SSRN. https://doi.org/10.2139/ssrn.6718600

Melo de Magalhães, G. (2026u). A structural law of swarm intelligence: Path-theoretic unification of navigability, motion, and direction. SSRN. https://doi.org/10.2139/ssrn.6304936

Melo de Magalhães, G. (2026v). There is no pathlessness. SSRN. https://doi.org/10.2139/ssrn.6712279

Melo de Magalhães, G. (2026w). Training perception for the lightest paths — Why intelligence requires going against lazy perception. SSRN. https://doi.org/10.2139/ssrn.6712481

Melo de Magalhães, G. (2026x). Unified theory of navigability — Do all systems navigate in the same way? SSRN. https://doi.org/10.2139/ssrn.6612058

Melo de Magalhães, G. (2026y). What gets smarter gets cheaper — The law of intelligent Midas. SSRN. https://doi.org/10.2139/ssrn.6507680

Melo de Magalhães, G. (2026z). What if reality is an operating system? The Freedom FLRP unified framework. Zenodo. https://doi.org/10.5281/zenodo.19684561

Melo de Magalhães, G. (2026aa). Why is there change? Freedom as the operating system of reality. Zenodo. https://doi.org/10.5281/zenodo.19410804

Melo de Magalhães, G. (2026ab). Why is there something instead of nothing? Because there is freedom. Zenodo. https://doi.org/10.5281/zenodo.19973624

Moussaïd, M., Helbing, D., & Theraulaz, G. (2011). How simple rules determine pedestrian behavior and crowd disasters. Proceedings of the National Academy of Sciences, 108(17), 6884–6888.

Musa, A. B. M., & Eriksson, J. (2012). Tracking unmodified smartphones using Wi-Fi monitors. In Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (pp. 281–294). ACM.

Observador. (2024, June 20). O que muda no segundo fim de semana de Rock in Rio Lisboa: Equipa para orientar o público para filas menores e casas de banho reforçadas. https://observador.pt

Portal da Queixa. (2024). Complaint records for Rock in Rio Lisboa, 2024 edition. https://portaldaqueixa.com

Público. (2024, June 19). Queixas do Rock in Rio Lisboa aumentam 97% em relação à edição anterior. https://www.publico.pt

Raza, U., Kulkarni, P., & Sooriyabandara, M. (2017). Low power wide area networks: An overview. IEEE Communications Surveys & Tutorials, 19(2), 855–873.

Schauer, L., Werner, M., & Marcus, P. (2014). Estimating crowd densities and pedestrian flows using Wi-Fi and Bluetooth. In Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems (pp. 171–177).

Small, K. A. (2012). Valuation of travel time. Economics of Transportation, 1(1–2), 2–14.

Still, G. K. (2014). Introduction to crowd science. CRC Press.

Time Out Lisboa. (2024, June). Rock in Rio Lisboa: O que muda para o fim-de-semana de 22 e 23 de Julho. https://www.timeout.pt

Van Hautegem, K., & Rogiest, W. (2017). No more queueing at the ladies' room: A queueing-theoretic study of toilet layouts and waiting-time disparity. Ghent University.

Vanhoef, M., Matte, C., Cunche, M., Cardoso, L. S., & Piessens, F. (2016). Why MAC address randomization is not enough: An analysis of Wi-Fi network discovery mechanisms. In Proceedings of the 11th ACM Asia Conference on Computer and Communications Security (pp. 413–424). ACM.

Wardman, M., & Ibáñez, J. N. (2012). The congestion multiplier: Variations in motorists' valuations of travel time with traffic conditions. Transportation Research Part A: Policy and Practice, 46(1), 213–225.

Wardman, M., Chintakayala, V. P. K., & de Jong, G. (2016). Values of travel time in Europe: Review and meta-analysis. Transportation Research Part A: Policy and Practice, 94, 93–111.

Weppner, J., & Lukowicz, P. (2013). Bluetooth based collaborative crowd density estimation with mobile phones. In Proceedings of the IEEE International Conference on Pervasive Computing and Communications (pp. 193–200). IEEE.

Wirz, M., Franke, T., Roggen, D., Mitleton-Kelly, E., Lukowicz, P., & Tröster, G. (2013). Probing crowd density through smartphones in city-scale mass gatherings. EPJ Data Science, 2, Article 5.

Zanella, A., Bui, N., Castellani, A., Vangelista, L., & Zorzi, M. (2014). Internet of Things for smart cities. IEEE Internet of Things Journal, 1(1), 22–32.

Zhan, B., Monekosso, D. N., Remagnino, P., Velastin, S. A., & Xu, L.-Q. (2008). Crowd analysis: A survey. Machine Vision and Applications, 19, 345–357.

Appendix A — Raw Sensor Coordinates

Table 3 (Section 6.3) is the authoritative raw coordinate record: all 24 seed points are reproduced verbatim as captured in the field, including the two flagged outliers (WC-04-H2, WC-07-M2) and all near-duplicate pairs, with three-level confidence grades. No coordinate has been silently corrected; canonical markers (Section 6.4) are derived values, and the derivation is reproducible from Table 1 alone.

Appendix B — Formula Summary

$$P(s,t) = we(t) \cdot E(s,t)/de + wi(t) \cdot I(s,t)/di$$

$$Occ(s,t) = 100 \cdot P(s,t) \,/\, Capacity(s)$$

$$E\_norm = \text{calibrated entrance / waiting-adjusted limit}\ ;\ \ I\_norm = \text{calibrated interior / physical capacity}$$

$$P(s,t) = wE(t) \cdot E\_norm + wI(t) \cdot I\_norm\ ;\ \ \text{Pressure score} = \min(100,\ 100 \cdot P)$$

$$w^*E = wE \cdot CE \,/\, (wE \cdot CE + wI \cdot CI)\ ;\ \ w^*I = wI \cdot CI \,/\, (wE \cdot CE + wI \cdot CI)$$

$$\text{Cluster display pressure} = 0.70 \times \text{worst section} + 0.30 \times \text{average section}$$

$$P1(x,t) = mean(raw(x,\ t-60s..t)),\ \ x \in \{sensor, section, cluster, estate\}\ ;\ \ P5 = \text{5-min rolling mean of } P1$$

$$v\_ticket = P\_ticket/H\ ;\ \ v\_wage = \beta \cdot w\ (\beta \approx 0.30)\ ;\ \ v\_wait = m \cdot v\ (m \approx 2\text{–}2.5)$$

$$L = \lambda W\ ;\ \ Wq \approx (\rho/(1-\rho)) \cdot ((Ca^2+Cs^2)/2) \cdot \tau\ ;\ \ (1/K)\Sigma Wq(\rho\_k) \geq Wq(\bar{\rho})$$

$$V = N \cdot \Delta t \cdot v\ ;\ \ \partial V/\partial \Delta t = N \cdot v\ ;\ \ C\_perceived = \theta(I) \cdot \Delta t \cdot v\ ;\ \ \Delta S = N \cdot \Delta t \cdot s$$

$$P5(s,t) = mean(\ P(s,k)\ ),\ k \in \text{last 5 minutes}$$

$$Q(s,t) = \max(\ 0,\ P5(s,t) - CapacityThreshold(s)\ )$$

$$C(s,t) = f(\ freshness, completeness, agreement, stability, health, geolocation, regime\ )$$

$$Value = attendance \times hours\_recovered \times value\_per\_hour$$

$$Freedom = (\ Perception \,/\, Distortion\ ) \text{ ^ } \alpha\infty \quad \text{(theoretical model)}$$

$$\text{Operational Freedom}(t) = \text{Available Capacity}(t) \times \text{Perceived Alternatives}(t) \,/\, \text{Distortion}(t)\ \text{(conceptual)}$$

Regimes: Light 13:00–16:30 · Growth 16:30–18:30 · Peak 18:30–21:30 · Post-concert 21:30–close. Entrance weight we(t) increases in Peak and Post-concert; interior weight wi(t) never reaches zero.

Appendix C — Sensor and Failure Checklists

Tables 21 and 22 (Section 19) constitute the operational checklists: seven deployment checks per node before doors, and six failure-handling rules enforced automatically by the pipeline. The two non-negotiable rules are repeated here for emphasis: never republish old values as current, and never manufacture confidence.

Appendix D — SCOR Data Schema and Transmission Rules

Schema: cluster_id, timestamp, telemoveis_detectados, pessoas_estimadas, homens, mulheres, unissexo, entradas_ir, saidas_ir, ocupacao_instantanea, contagem_prosegur, confianca_cruzada, estado_sensor (Table 12). Transmission rules: values are 5-minute rolling averages, never last-value snapshots; if fresh data exists, send the calibrated average; if data is stale, send the offline/degraded state with confianca_cruzada = 0; never backfill fake confidence; never send individual-level data; never store personal trajectories.

Appendix E — Glossary of Freedom Intelligence and Flow Intelligence

Terms are defined as used in this paper; bracketed citations point to the corpus where each is developed in full.

Freedom (F). The structural availability of paths: the condition in which transitions between states actually exist, filtered through perceived distortion. Measured as $F = P/D$. A measurable structural fact (Melo de Magalhães, 2026a, 2026p).

Perception (P). An agent's usable understanding of available paths — not raw sensing, but the clustered, differentiated reading of real possibilities (Melo de Magalhães, 2026j, 2026l, 2026r). In this deployment: what the app, screens, and dashboards make visible.

Distortion (D). Everything that degrades the usability of paths: uncertainty, misinformation, crowding, waiting, bad signage, sensor noise, GPS error, network failure, lack of visible alternatives. The quantity every subsystem of this paper works to reduce.

F = P/D (Law of Freedom). Canonical form derived from monotonicity, scale-covariance, and separability; recovers Ohm/Fourier/Fick/Darcy in the passive regime ($R^2 > 0.99$); agent-level $R^2 = 0.885$ over 57,518 trials. Dimensionless by construction: P and D are ratios measured in the same units, so F is a pure number — which is what lets one law compare a corridor, a queue, and a circuit. Validated computationally on Deucalion across 57,518 agent-level trials, reproduced three independent times (Melo de Magalhães, 2026c, 2026o, 2026u).

α∞ (recursive limiting exponent). Exponent in the recursive extension Freedom = $(P/D)^{\alpha\infty}$, inspired by Lambert-W / self-referential limit structures; expresses perception acting on its own output. Fitted agent-level values ≈ 1.19–1.29 (Melo de Magalhães, 2026o, 2026t).

Values of Freedom. Interpretive regimes of the ratio: $F \to 0$ (collapse), $0 < F < 1$ (constrained navigation), $F \approx 1$ (balance), $F > 1$ (open field). Theoretical bands mapped to festival regimes in Table 1; not calibrated thresholds.

Path (capitalised). A structural transition in possibility space: the availability of a next state from the current one. Paths are without end — possibility is recursive; every reached state opens further transitions (Melo de Magalhães, 2026k, 2026t).

path (physical, lower case). A bounded route in the P-layer: a corridor, an approach to a cluster. Physical paths end — at walls, saturation, closing time. Thesis V: they are the distortion that cannot be rerouted through.

Path availability. The primitive of the framework: whether transitions exist from here. The first condition of any navigable system (Melo de Magalhães, 2026p).

Pathlessness. The counterfactual of a world without Paths: no transitions, hence no dynamics, information, causation, or observation. Argued incoherent/unobservable — not a possible world (Melo de Magalhães, 2026d, 2026v).

Navigability. The degree to which a system's states can be traversed by its agents; the cross-domain property the corpus unifies (Melo de Magalhães, 2026x).

Lightest path. The path of least perceived resistance — not shortest, not fastest. What ants, swarms, crowds, and currents converge on without coordination (Melo de Magalhães, 2026u, 2026w).

Freedom field. The spatial distribution of available paths that each agent reads locally; the structural condition of swarm intelligence (Melo de Magalhães, 2026s).

Freedom Intelligence. The research programme and design-science framework: the theory of path availability, perception, distortion, navigation, and counterfactual choice (Melo de Magalhães, 2026a–2026ab).

Flow Intelligence. The operational implementation of Freedom Intelligence in physical space: the cyber-physical layer that senses, calibrates, averages, interprets, and guides flows.

Bathroom / WC Flow Intelligence. Flow Intelligence specialised to sanitary infrastructure: queue pressure, occupancy, waiting, and routing across WC clusters — this paper's system.

FLRP. The four-layer operating structure: Freedom (path topology) → Logic (cost filtering) → Relations (coordination) → Physics (persistent resistance). Generated forward, observed in reverse (Melo de Magalhães, 2026m, 2026z).

F-debt. The measurable hourly cost of suppressed freedom — time and value lost while paths remain imperceptible. The waiting tax of Section 5 is its festival form.

Counterfactual path. A path the agent would not have taken — often could not have seen — without the system: "not the visible queue; the better available cluster."

Operational Freedom. Conceptual venue-level quantity: Available Capacity × Perceived Alternatives / Distortion (Section 10). A design compass, not a fitted model.

Wave regime. One of four operational periods — Light, Growth, Peak, Post-concert — governing time-dependent weights we(t), wi(t), calibration posture, and staffing (Table 11).

Freedom collapse. The Peak-regime failure mode: a wave arrives, alternatives exist, nobody sees them in time; local F → 0 and a queue crystallises.

Entrance sensor (H1/M1/U1). Node at a section's approach; estimates queue pressure and detects wave fronts. Tighter sensing radius; weight rises at Peak/Post-concert.

Interior sensor (H2/M2/M3/U2). Node inside a section; validates actual occupancy and prevents passers-by from being booked as occupants.

Section. The unit of capacity and routing: male (H), female (M), or unisex (U). Section structure is operationally decisive — capacities and queue dynamics differ (Table 2).

Waiting allowance (Espera). Operational waiting/queue-pressure allowance per cluster in the official occupancy limits (Table 2); Total = capacity + Espera.

Sensing radius. How far a node listens for nearby devices, RSSI-gated, remotely calibratable 3–60 m (Table 5). Not the same as LoRa range.

LoRa communication range. How far a node's data reporting travels to a gateway — kilometre-scale under good conditions. Two orders of magnitude larger than sensing radius; conflating them is a deployment error.

Confidence C(s,t). The per-role trust score from freshness, completeness, entrance–interior agreement, stability, sensor health, geolocation confidence, and regime. Properties: multiplies fusion weights; publishes alongside every value; collapses gracefully to 'offline'. Excludes: retroactive repair of stale data; manufactured certainty.

Negative calibration. The field discipline of believing signals less: raising divisors, tightening radii, reducing weights, lowering confidence, offlining stale nodes, never republishing old values.

Five-minute rolling average (P5). The stability layer: all published and exported values are 5-minute means, never last-value snapshots; empty windows publish offline, not stale data.

Confidence C(s,t). Per-estimate score from freshness, completeness, entrance–interior agreement, stability, sensor health, geolocation confidence, and regime; governs publish/degrade/offline.

Queue pressure Q(s,t). max(0, P5 − CapacityThreshold): the smoothed excess over a section's threshold — the operational trigger for routing and alerts.

Canonical WC marker. One public marker per cluster, computed as the outlier-filtered median of sensor seed coordinates with manual override; sensor scatter never reaches the public map.

Anonymous aggregate data. Counts with no identifiers, computed at the edge — outside GDPR's material scope. Distinct from pseudonymous device data (stored identifiers, even hashed), which remains personal data.

Waiting tax. The friction paid in time by every attendee in an unmanaged queue system: queue time → lost experience → lost value, invisible on any balance sheet.

Pressure score. min(100, 100·P(s,t)): the published operational quantity — not a people count. Bands: 0–40 low, 40–70 medium, 70–90 high, 90+ critical.

E_norm / I_norm. Calibrated entrance and interior estimates normalised by the waiting-adjusted section limit and the physical section capacity respectively (Table 6) — the step that makes 80 people mean the right thing in both WC-03 and WC-06.

70/30 weighting. The entrance-dominant fusion typical of Growth, Peak, and Post-concert: ≈70% entrance pressure + ≈30% interior occupancy, regime-banded per Table 7 and confidence-reweighted. Rationale: entrance predicts the wave; interior validates the wave.

Occupancy vs pressure. Two different questions: how full is the WC (interior-led) vs is a queue forming (entrance-led). The system optimises for the second, because festivals fail at queue formation, not at fullness.

Worst-section rule. Public cluster pressure = 70% × worst section + 30% × average section — a saturated female section makes the cluster operationally problematic even if the male side is fine.

Wave Flow Control. The technology category this paper introduces (to our knowledge, new): sensing, calibration, prediction, routing, and privacy-preserving intelligence used to understand and coordinate crowd-flow waves in large physical spaces. Properties: wave-regime aware; capacity-normalised; confidence-scored; closed-loop to attendees; privacy by incapacity. Excludes: individual tracking; camera analytics; accuracy warranties; permanent infrastructure. First deployed module: WC / Bathroom Flow Intelligence, productised as Flow Intelligence for Festivals by Planta Smart Homes.

Flow. The physical phenomenon itself: people moving, waiting, converging, dispersing, occupying, entering, leaving, queuing. Properties: exists independently of measurement; carries momentum and waves; consumes capacity and time. Excludes: data, dashboards, and estimates — those are its shadows.

Smart data. The stabilised measurement layer of this system: the average of the last 60 seconds of raw values, computed per sensor, per gender section, per cluster, and at the all-clusters level, feeding the five-minute published series. Properties: flicker-free; wave-responsive; hierarchical (same rule at every level); the sole input to every public surface. Excludes: raw second-level detections; single-sensor spikes; any value older than its declared freshness.

Single-sensor mode. The compact-cluster configuration: one door-level node at tight radius covering queue, doorway, and interior where the section's depth sits within ~10–15 m. Properties: preserves occupancy estimation and ranking; cheap; fast to deploy. Excludes: entrance lead time (wave prediction); full confidence — its weight in venue-level decisions is capped by design.

Unisex pooling. Shared-cabin operation of a cluster (WC-05, WC-06), turning two queues into one multi-server queue. Properties: statistical multiplexing; sharply reduced female waiting (over six minutes to under 1.5 in the Ghent model); higher utilisation of the same cabins. Requires by design: full-enclosure cabins, lighting, attendants at Peak, tightened cleaning cycles.

Intelligence (intelligere). In this paper: the capacity to see the lightest path within or between distorted lines — the mechanism that makes distortion lighter (Thesis V).