Architectural Integration of the P-index Framework

Design-Stage Crowd Physics for Signature Projects

Architectural Integration White Paper

Yixuan (Kevin) Tao

Pneuma Theorem, Inc.

ORCID: 0009-0005-8136-3886

Framework developed with technical input from Dr G. Keith Still, co-author of the Kroll–Still–Neuman–Graham–Griffin (2017) biomechanical model of compression asphyxia.

Companion documents.

Flagship academic manuscript (B01): Zenodo DOI 10.5281/zenodo.21261698

Bipartite nine-cell framework (E01): Zenodo DOI 10.5281/zenodo.21266645

Sample deliverable (R02) available on request.

Public release.

Document series context

This white paper (W04) is the architectural-integration instrument in the Pneuma Theorem, Inc. document series. It presents the design-stage application of the P-index framework for architects, engineering consultancies, and landmark project owners. It does not restate the physical derivation; that is carried by the flagship manuscript.

Zenodo (B01) *"A Physical Quantification Framework for Crowd Compression Risk: The P-index System and Its Contact-Mechanical Foundation."* Flagship academic manuscript. Physical derivation and complete bibliography.
Zenodo DOI: 10.5281/zenodo.21261698.

Zenodo (E01) *"A Bipartite Nine-Cell Framework for Crowd Safety Analysis: Venue P-Report and Event DIM ICE as Complementary Instruments."*
Zenodo DOI: 10.5281/zenodo.21266645.
The bipartite structure described there has since been superseded by the three-axis matrix used in the present document; the Zenodo record is retained as the published version of record, and a revised version is in preparation.

W04 *(this document)* Architectural integration — design-stage crowd physics for signature projects.

Executive summary

This white paper (W04) sets out what the P-index framework establishes about built form, for architects, engineering consultancies, and landmark-project clients. Its subject is the one question a drawing can answer on its own: whether a configuration admits the mechanism by which a compressive load accumulates, or whether its shape forecloses it.

The design-stage argument is that physics-based crowd risk is decided in geometry, and geometry is fixed at design stage. Traditional crowd safety assessment enters after design freeze, by which point the cheapest mitigations — the geometric ones — are unavailable. The framework's inputs are available from RIBA Stage 2 onwards.

The central result is that two kinds of safety are equivalent in verdict and not otherwise. A configuration can fail to produce the mechanism because its shape does not admit one, or because a measure prevents one. Both satisfy the same criterion; they differ in that the second has an execution rate and the first does not. And a lapse in the second is not recovered by resuming it: forming a compacted region requires work, maintaining it requires none, so what forms during a lapse persists after the measure is restored. An operational measure can be specified at any point in a venue's life; a shape that requires no measure can only be chosen while the shape is still being chosen.

And the framework's own results say why the geometric stage is the decisive one. Whether a load accumulates along a route at all turns on how many directions of travel the route admits — a run bounded on both flanks supplies an alignment that an open one divides, and the difference is large enough to carry a crowd across the threshold at which accumulation begins, with no one exerting themselves any harder (Tao, 2026a). That property is set by the built form and by nothing else.

The consequence for what can be promised is sharper still. Of the channels through which any factor reaches a compressive load, the bounded ones — what a body can exert, what it can withstand — admit mitigation, while the unbounded one, the extent over which contact is continuous, is the only channel on which a guarantee can be constructed rather than a reduction (Tao, 2026b). Design is where that channel is set.

What a design-side assessment delivers is a statement about the derivative, not the integral. P quantifies exposure received and takes a positive value only in reconstruction; it is not a design-stage target. What a configuration can be shown to satisfy in advance is the *structural zero*: that the borne load does not grow with depth, over the whole occupied domain and every admissible arrangement of its occupants (Tao, 2026a). The claim is geometric impossibility rather than statistical rarity, and the distinction is the whole of it — a configuration that is rarely dangerous is one perturbation away from being dangerous.

The design variable that follows is a partition spacing rather than a headcount. None of the load relations contains a number of people: two hundred occupants at a given spacing and two hundred thousand at the same spacing carry the same guarantee, because what the relation bounds is the length of a connected run. Spacing is continuous, and the floor area it costs is small and calculable.

Findings land in a matrix of three columns and three rows. The columns are the conditions the mechanism requires — space, force, time — and they are necessary together, so a configuration is safe as soon as any one of them fails. The rows are a control loop: what the place physically is, whether anyone can tell what state it is in, and whether anyone acts on that. Each element of a venue is classified as closed by its shape, closable by a measure of stated dimension, or not closable within scope. What a review delivers is the count of elements still in the second state, together with the dimension each rests on. A missing input is recorded as missing; an absence of adverse findings is not a pass.

The framework states where it stops. Three quantities are missing — the per-contributor exertion, the residual a standing body absorbs, and the rate at which a region fills — and all three fall on the force and time axes. The space axis, the one a building controls and the one this document is about, is complete. Three of the missing quantities are properties of bodies rather than of any site: measured once, and thereafter available to everyone.

On what is available. Two assessment products are in service today, both closed-form: a pre-event geometric assessment and a retrospective backtrack analysis. Each answers what a geometry permits from the drawing and a declared exertion, by arithmetic a reader can check by hand. The design-stage engagement described in this paper is a capability of the framework rather than a service currently offered; the order of work is deliberate, and Section 6 sets out both, together with the four questions that lie outside the scope.

On thresholds. This paper adopts none. The framework establishes what load a configuration can produce and how that load accumulates; it does not establish what load injures a person, and it catalogues external values for the reader's selection rather than selecting among them (Tao, 2026a). A client adopting a different value obtains different figures from the same analysis, and the analysis is unchanged in form. Where this document reports a number against a level, the level is named at the point of use.

Framework foundations are documented in the flagship academic manuscript (Tao, 2026a) and the bipartite nine-cell paper (Tao, 2026d); commercial and methodological framing are in (Tao, 2026e,f); the framework itself was developed with technical input from Dr G. Keith Still.

Contents

1. The design-stage advantage

Physics-based crowd risk is decided in geometry, and geometry is fixed at design stage. Every operational mitigation applied later — stewarding patterns, ticketing controls, timed egress, information systems — works within constraints that the building has already imposed. When those constraints are unfavourable, operations absorb the residual risk indefinitely. When those constraints are favourable, operations can be lighter, cheaper, and more resilient across the venue's design life. The argument of this white paper is narrower than a claim of priority between the two. A configuration can fail to produce a compressive mechanism in either of two ways — because its shape does not admit one, or because a measure prevents one — and the two are equivalent in verdict and not otherwise. Only the first is available at design stage, and only the first has no execution rate.

1.1 Why physics-based crowd risk belongs in design, not operations

The distinction between design-stage and operational-stage crowd risk management is not a matter of timing alone; it is a matter of what can be changed. Once a concourse width, a stair pitch, or an aperture position is built, it is fixed for the venue's design life. Every subsequent operational decision — how many stewards to station where, how to phase egress, when to close a gate — works within the envelope that geometry has set.

Traditional crowd safety assessment is performed after design freeze, often by a consultant retained by the operator rather than the architect. The assessment identifies risks; the operator absorbs them as management overhead. This structure has two consequences. First, the cheapest risk mitigations — the geometric ones — are unavailable by the time the assessment happens. Second, the safety consultant and the design team never converge in the same room: the consultant's recommendations arrive as a report, not as a design input.

The P-index framework is designed to enter earlier. Its inputs (enclosure geometry, aperture topology, expected occupancy scenarios) are available from RIBA Stage 2 onwards; its outputs — which regions are closed by their shape, which are closable only by a measure, and what load a given partition spacing guarantees — are legible to architects and can be iterated against alongside daylighting, structural, and acoustic analyses. The framework does not replace the safety consultant's operational assessment; it precedes it, and reduces the residual risk that the operational assessment has to manage.

1.2 What an approval process can and cannot settle

A design that attracts concern is not usually stopped by concern. An objection raised in an approval process carries a burden: to withhold consent, the objecting party has to establish that the design fails against some criterion. Where a criterion exists, that is a tractable exercise. Where none does, it is not.

On the axis this framework addresses, none does. The applicable criteria are expressed in width, rate, density and time, and none of them is how many directions of travel a region admits. An objection resting on that property therefore arrives as a judgement rather than as a determination, and a judgement does not discharge a burden.

What follows is the part worth attention. A design team answering such an objection may offer a measure — a controlled admission rate, monitoring, additional stewarding. The burden then returns to the objector, who must now establish that the measure is insufficient. Whether a measure is sufficient is governed by the same absent criterion. The objection cannot be sustained, and the measure stands on the record as an answer to it.

The result is that a property no party can quantify passes through an assessment process as though it had been assessed. Nobody need have behaved improperly for this to happen: the design team offered a mitigation, the objector could not show it inadequate, and consent issued on the material available.

The same absence operates in the other direction. Where an outcome is examined afterwards, the question of what ought to have been foreseen returns to the same quantity, and a property that could not be established in advance is not readily established in retrospect either. The gap that lets a configuration through is the gap that makes it hard to say afterwards what should have been seen.

What a quantity changes, and for whom. Where the property can be read off the drawing, three things follow that do not follow otherwise. An objection becomes a determination rather than a judgement. A proposed measure can be tested against the quantity it is offered to control — and a measure acting on how many people are admitted does not act on a load bounded by the length of a connected run rather than by headcount. And what ought to have been foreseen becomes a question with an answer available before the event rather than only after it.

For a design team the consequence is narrower and more immediate. The approval process will not settle this axis, because no party to it holds the quantity that would settle it. A design that passes has not been cleared on this property; it has passed a process in which the property was not among the things assessed. Where that matters, it has to be asked separately — and it can only be asked cheaply while the geometry is still being drawn.

1.3 Two kinds of safety, equivalent in verdict and not otherwise

A configuration can fail to produce the mechanism in two ways, and the distinction is the one that separates two venues which both satisfy a criterion.

Geometrically. The region admits several directions of travel, the alignment factor divides among them, and the per-contributor exertion falls below what a receiving body absorbs. Nothing accumulates at any length, for any duration, and at any occupancy the exits permit (Tao, 2026a). No measure is required, so no measure can fail.

Operationally. The region admits one direction and the mechanism can operate, but a measure prevents the connected run from reaching a length at which the load matters. The configuration is safe while the measure holds.

Both satisfy the same criterion. They differ in that the second has an execution rate and the first does not.

And a lapse in the second is not recovered by resuming it. It might be supposed that a measure executed imperfectly yields proportionally less of the protection it was taken for. It does not. Forming a compacted region requires work, supplied by advance from behind; maintaining it requires none, since once the space between bodies has been closed it is held by geometry rather than by anything continuing to push. What forms during a lapse therefore persists after the measure is restored, and the framework's reversal bound establishes that a run beyond a modest depth cannot be opened from its downstream end at ordinary exertion (Tao, 2026a).

This is what the design stage buys. An operational measure can be specified at any point in a venue's life; a shape that requires no measure can only be chosen while the shape is still being chosen. The two are not alternative routes to the same outcome, because one of them carries a residual that is borne for the design life of the building and the other does not.

This positioning is central to the engagement structure described in Section 6. The audience addressed here — architects, engineering consultancies, and landmark-project clients — commissions the assessment as an input to a design process. The operational reporting stack is produced separately, at a later stage and in coordination with the venue operator; that is the subject of (Tao, 2026f).

1.4 The failure mode we address: SFM and DEM limitations at their operating regimes

The two established simulation families in crowd analysis — Social Force Models (SFM) and Discrete Element Methods (DEM) — have well-documented operating regimes and equally well-documented failure modes at the boundaries of those regimes (Helbing and Molnár, 1995; Helbing et al., 2000).

Social Force Models describe pedestrian motion as the response to attractive and repulsive social forces plus a desired-velocity term. They are calibrated on low-to-moderate density flows in benign conditions (transit stations, corridors, unhurried egress). In this regime they are accurate and useful. Their known failure mode is high-density adversarial motion: the desired-velocity term becomes ill-defined when individuals are not free to choose a velocity, and the pairwise force model, tuned to non-contact interactions, does not represent sustained mechanical contact (Helbing et al., 2007). This is exactly the regime in which fatal compression occurs.

Discrete Element Methods, imported from granular physics (Cundall and Strack, 1979; Langston et al., 2006; Majmudar and Behringer, 2005), represent each individual as a rigid particle with explicit contact mechanics. They handle the high-density contact regime well but require calibration inputs (friction coefficients, contact stiffnesses, damping parameters) that are hard to justify for human bodies in clothing, and the computational cost grows rapidly with population size. DEM is well suited to short-window, small-N scenarios; it does not scale to whole-venue design review.

The framework does not occupy the middle ground between them; it answers a different question. It does not predict where people go, and makes no attempt to. It bounds the load a geometry permits, given what is declared about the exertion, and that bound is available without a trajectory. The derivation is in (Tao, 2026a); the test of whether a given model may be used to characterise a real configuration at all is the subject of (Tao, 2026c).

1.5 Where the P-index integrates in the RIBA Plan of Work

The framework's entry points across the RIBA Plan of Work 2020 are:

- **Stage 2 (Concept Design)**: screening of massing options for which regions are closed by their shape and which admit a single direction of travel; identification of the scenarios that stress the design. This is the highest-leverage stage: geometric decisions taken here propagate through every subsequent stage.
- **Stage 3 (Developed Design)**: variant comparison across design options, read on the number of elements each variant leaves resting on a measure rather than on its shape; iterative refinement of concourse, stair, and aperture geometry. See Section 4.2.
- **Stage 4 (Technical Design)**: the partition schedule — each spacing with the load it guarantees and the premise it rests on — together with the quantification of any element that remains closable only by a measure. See Section 4.3.
- **Stage 5 (Handover)**: the list of elements whose safety rests on execution, each stated as a dimension that can be recorded, transferred to the operator's safety team as an operational input rather than a compliance document. Where the operator works to a DIM ICE grid, these are the entries that populate it. See Section 4.4.

The engagement structure in Section 6 maps these stages to two service tiers: BD1 (a Stage 2–3 review) and BD2 (a full Stage 2–5 engagement for landmark projects).

2. Framework foundations

The framework carries three quantities, and they answer different questions. The C-value is the input: what the scene supplies. The R-indicator is the process: what accumulates, and how fast. P is the conclusion: what was borne, and for how long (Tao, 2026a). What a design engagement delivers is none of the three — it is a certificate that the third is identically zero, and the subsections below build up to why that is the correct deliverable rather than a weaker one. Complete derivations are carried by the flagship manuscript (Tao, 2026a); readers already familiar with the framework may proceed directly to Section 4.

2.1 P: a conclusion about exposure, not a design target

P quantifies the sustained mechanical exposure a person received from the surrounding crowd. It is a *form* rather than a fixed quantity: it integrates a dose–response function the framework does not supply, normalised by a capacity the framework does not supply either, so its dimension and its scale are determined by those two choices (Tao, 2026a). What the framework fixes is the structure — which interval is integrated over, and the truncation that makes the result monotone in exposure. Loading is reversible and damage is not, and that asymmetry is written into the definition rather than attached as a caveat.

It is derived from contact mechanics rather than from statistical fits to historical incidents: given the geometric enclosure, a directional-intensity field (the C-value below), and the extent over which contact remains continuous, the load borne by a representative individual follows from contact mechanics and rigid-body packing arguments.

Which level, and whose choice it is. The framework establishes what load a configuration produces and how that load accumulates. It does not establish what load injures a person. External values are catalogued with their provenance and evidential character — an experiment on living subjects, an output of a biomechanical model, a retrospective estimate — and the selection among them is the reader's (Tao, 2026a). Two consequences follow for a design engagement.

The first is that a client adopting a different level obtains different figures from the same analysis, and the analysis is unchanged in form. A jurisdiction, an insurer, or an internal standard may each nominate its own; none requires a separate study.

The second is that this report names the level wherever it reports a figure against one. A number given without its level is not interpretable, and a level adopted without its stated population and provenance is a value used outside the setting that produced it.

Two designs at the same average occupancy differ in what they permit, and the source of the difference is geometric rather than behavioural: how many directions of travel each region admits, and over what extent contact remains continuous. The second of those is the one a design can be built to control, and the next subsection states what controlling it buys.

2.2 The design-side certificate

A certificate issued before an event cannot rest on an integral whose integrand is not the framework's to supply. It rests instead on the derivative: a configuration satisfies the *structural zero* if the borne load does not grow with depth, over the whole of the occupied domain and over every admissible arrangement of its occupants (Tao, 2026a).

Two independent routes establish it, and a configuration may qualify under either. The first bounds the chain: where partitions maintain a clear spacing, no coherent run exceeds that length, and the bound holds because the space a longer run requires is not present. The second bounds the growth: where each contributor delivers less than a receiving body absorbs, the load is the same at every depth, and a configuration certified this way remains certified however long its runs.

What the certificate says. That the accumulating mechanism is absent on grounds of geometric impossibility rather than statistical rarity. The distinction is the whole of the claim: a configuration that is rarely dangerous is one perturbation away from being dangerous, whereas one in which the mechanism does not operate does not become dangerous by being perturbed.

And what the word means here. It is used in the sense it carries in mathematics and in verification — an object that allows a claim to be checked — and not in the sense of an instrument that confers a status. Nothing in this document is issued by, or on behalf of, any body empowered to approve a configuration for use, and no such standing is claimed anywhere in it. Whether a configuration may lawfully be occupied is a question for whoever is empowered to answer it; a certificate of this kind is an input to that question and not an answer to it.

Premises are the substance. What distinguishes one certificate from another is not the bound alone but how difficult its premises are to violate. A bound resting on a partition that must be kept clear is weaker than the same bound resting on a wall, and a report that does not distinguish them has not reported the thing that matters. This is also the characteristic point at which as-built departs from design, and stating the premise is what allows a later reader to establish whether it held.

2.3 C-value: directional-intensity scaling of the contact load

Density alone does not determine risk. A crowd of 5 persons/m^2 moving uniformly in one direction produces low internal contact force; the same density with two counter-flowing streams converging on a shared axis is a different mechanical situation entirely, and the density figure does not distinguish the two (Tao, 2026a). The C-value is what does. In Tao (2026a) it is a *per-contributor* directional-intensity coefficient: the axial force one person delivers along the local compression axis, divided by a fixed scale $F_{\text{den}} = \mu W_0 \approx 480$ N — the maximum horizontal force a standing adult can exert against foot traction. It resolves into a horizontal factor C_h (directional alignment, on $[0, 1]$) and a vertical factor C_v (delivered magnitude), $C = C_\text{h}\, C_\text{v}$.

Because the scale is a physical limit, a single contributor's C-value is bounded, not open. Active pushing saturates at $C_{\text{active}} = 1.0$, the point at which the feet slip; a person leaning or pressed transfers body weight rather than a traction-limited push, reaching at most $C_{\text{passive}} = 1/\mu \approx 1.43$. No individual in a crowd exceeds this. High received loads therefore do not arise from a large C-value; they arise from *force-chain summation*. The load a thorax receives is $C \cdot F_{\text{den}} \cdot \Phi(L)$, where $\Phi(L)$ grows with the depth L of the chain of bodies transmitting force onto that individual (Tao, 2026a). A received load equivalent to two or three times F_{den} is the signature of a deep chain of individuals each at an ordinary C-value — not of any one person pushing implausibly hard. Risk escalates through two channels the design review must watch: directional convergence raising C_h, and geometry deepening the force chain L.

One premise travels with those ceilings. They bound what a person generates by pressing against the ground. They do not bound a body that has fallen, a barrier that has given way, or an agent that is not a person at all (Tao, 2026a). An assessment carrying the ceilings forward without declaring which sources of load it has admitted is asserting something stronger than it has established, and the gap will not be visible in the arithmetic. The declaration is short and belongs with the other premises: the sources treated, those excluded, and the reason for each exclusion.

The C-value scales the effective contact load. Its practical role in design review is to identify configurations where the density looks acceptable but the coherence structure is adversarial. Two families of configuration warrant attention, and they are not the same problem.

But first, the condition under which a bound on alignment applies at all. It is a phase property rather than a continuous scene parameter (Tao, 2026a). Where a blocking person has somewhere to move to, the cheapest way past them is a small sideways push, which is orthogonal to the axis along which a chain would transmit and contributes nothing to it. Where the blocking person has nowhere to move to, no lateral push accomplishes anything, and the effort is necessarily aligned with that axis. The practical reading is that alignment is low in an open crowd not because individuals exert little force, but because the force they exert is not oriented to accumulate — which is why ordinary crowds contain a great deal of visible pushing and no force chains of consequence.

The first is bounded travel, and the framework treats it directly. A region open on its flanks distributes its occupants among several directions of travel, and the alignment factor divides among them; a run bounded on both flanks admits one, and no such division occurs. The difference is a multiple large enough to carry a crowd across the threshold at which accumulation begins, with no one exerting themselves any harder (Tao, 2026a). In design review this means a walled approach, an under-croft, a single-sided concourse and a bounded ramp are each unfavourable *directly*, and independently of how many people are expected.

The second is intersecting streams, and the framework is explicit that its scalar does not resolve it. Where two directed flows meet at a shared axis, the situation is adversarial in a way the alignment factor as a scalar does not distinguish from a benign one: a scatter of weak multi-directional contacts and a collision of two coherent streams are both low-alignment states, and both sit below the single-direction bound (Tao, 2026b). They differ sharply in what they do to an occupant's ability to remain upright, and the framework records that as an open item rather than resolving it in the coefficient.

The practical instruction for design review is therefore asymmetric. Bounded runs are assessed by the framework and reported quantitatively. Intersections of directed flows — at concourse junctions between opposing seating blocks, at stair heads where ascending and descending flows meet, and wherever segregated groups share circulation space — are identified and flagged, and the assessment states that the quantitative treatment of that configuration is not settled. A design review that reported a figure there would be reporting one the framework does not support.

2.4 Partition spacing: the design variable, and why it is not a headcount

Traditional occupancy calculations derive a capacity number by multiplying usable floor area by a density limit (typically 2 or 4 persons/m^2 depending on jurisdiction). That produces a geometric ceiling, and it controls a variable this mechanism does not depend on: none of the load relations contains a number of people. Two hundred occupants at a given partition spacing and two hundred thousand at the same spacing produce the same guarantee, because what the relation bounds is the length of a connected run (Tao, 2026a).

What a designer chooses is therefore a spacing, and the choice is continuous rather than binary: each spacing corresponds to a guaranteed ceiling on the load a chain can deliver, because a partition does not block a chain — it denies the chain length. Consider a square occupied area twenty metres on a side. The spacing between successive load-bearing contacts follows from the geometric packing limit alone, at $1/\sqrt{7} \approx 0.378$ m, so the undivided area admits $\lceil 20/0.378 \rceil = 53$ bodies in depth. Divided at two metres, the figure is 6. The guarantee falls by a factor of nearly nine, and the capacity given up is the area the partitions themselves occupy — nine lines of a tenth of a metre across a twenty-metre span, or some four and a half per cent of the floor.

The conventional remedy is to reduce density, which purchases a lower load by giving up occupancy directly and in proportion. Partitioning purchases a larger reduction at a small fraction of that cost, and the difference is not a refinement of the same trade: the two act on different variables.

The guarantee is conditional on clearance. A partition that is kept clear divides the domain; a partition that is stood upon does not exist. The bound holds only while the clearance holds, and that premise belongs in the assessment rather than in a footnote.

So an assessment carries two findings rather than one. A load finding, determined by the partitioning and stated as a guarantee with its premises; and an occupancy finding, determined by refuge area and egress and drawn from the governing guidance. They are not substitutes, and a criterion combining them into a single admissible number obscures which of the two is binding.

2.5 DIM ICE integration: nine-cell mapping to design decisions

The Design–Information–Management by Ingress–Circulation–Egress grid, established by G. Keith Still, is the standard organising instrument for crowd safety planning, and a venue operator will in many cases already work to one (Still, 2000, 2014). It is not this framework's assessment instrument — that is the three-axis matrix described in the next section — but it is the interface through which a design-stage finding reaches an operations team.

For an architectural audience, the value of the DIM ICE grid at handover is that it forces every finding to terminate in an actionable cell. An element that this framework reports as closable only by a measure becomes either a Design–Circulation decision (change the geometry, and close it by shape instead), a Management–Circulation decision (who maintains the measure), or an Information–Circulation decision (signage, wayfinding, PA sequences). The grid prevents findings from remaining as narrative recommendations; each finding lands somewhere.

The two-layer stack — quantitative findings underneath, DIM ICE on top — is the subject of the methodology white paper (Tao, 2026f) and is applied throughout Section 4.

3. The assessment instrument

A finding has to land somewhere a client can act on it. The instrument this framework uses is a matrix of three rows and three columns: the columns are the conditions the mechanism requires, and the rows are the layers at which a venue can do something about them.

3.1 The columns: three conditions, all of them necessary

Space. Whether the geometry admits a connected run long enough that the load borne at its far end exceeds the level the reader has adopted. This is a question about lengths and boundaries, answerable from the drawing.

Force. Whether the scene generates enough exertion per contributor for a run to accumulate at all. Where each contributor delivers less than a receiving body absorbs, the load is the same at every depth and no length is dangerous (Tao, 2026a).

Time. Whether the filling process has run long enough for the state to form. A region that would be hazardous when filled is not hazardous while it is still filling.

The three are necessary together. A configuration is dangerous only where the geometry permits the run, the scene supplies the exertion, *and* the process has had time — which means safety has three routes rather than one, and a configuration is safe as soon as any single condition fails. The asymmetry is worth stating plainly because it runs the other way from intuition: the burden on a party asserting danger is heavier than the burden on one asserting safety, since the first must establish all three and the second need establish only one.

For a design team the first column is the one the building controls, and it is also the one this framework treats most completely. The other two are where its acknowledged gaps lie (Section 5).

3.2 The rows: a control loop

The rows are not three domains sitting side by side. They are the three elements of a control loop, and they are ordered by what it costs to change them.

Hardware — the plant. What the place physically is: lengths, widths, boundaries, partitions. Dearest to change and cheapest to rely on, because it does not have to be maintained.

Software — the sensor. Whether anyone can tell, while the event is running, what state the place is in. Not whether the space is adequate, and not what to do about it — only whether the quantity is observable.

Management — the controller. Whether anyone acts on what the sensor reports, and whether the action is available in the time it has.

The weakest row governs, and here that is a necessity rather than a convention. A generously proportioned venue with no sensor on the one space that matters is, for the purposes of this matrix, an uncontrolled system: nothing about the plant tells anyone when to act. The rows do not average, and a strong row does not compensate for a weak one.

3.3 What a cell contains

Each cell is written in three layers.

L0 — basis, supply, demand, determination. The criterion the cell applies; what the venue supplies against it; what the event demands of it; and the determination that follows. The venue and event sides are kept separate because they are commissioned separately.

L1 — measures. What could change the determination, and what could not. A measure that improves the record without changing what is known is named as such.

L2 — execution and traceability. Where as-built departs from design. Effective width consumed by fencing, concessions or stored material; gates shown on a drawing that are closed during admission; partitions that exist only while kept clear. This layer is where a premise stated at design stage is either confirmed or found not to hold.

The two-sided L0 is what makes three engagements possible. A venue assessment fills the venue side across all nine cells; an event assessment fills the event side; a venue promoting its own event fills both. **One limitation must be stated in either of the first two: the time axis cannot be closed from one side alone**, since the rate at which a space fills depends on the venue's geometry and the event's arrival profile together.

3.4 No data is not a pass

Two findings look similar in a report and are not the same.

Not established records a determination that was reached: the inputs sufficed and the criterion failed. *Undetermined* records that no determination was reached: the inputs did not suffice. The first is a finding about the configuration and calls for a change to it. The second is a finding about the assessment and calls for an input.

They are distinguishable only if the report distinguishes them, and a reader meeting an absence of adverse findings will otherwise read a pass. Where an input is missing, the correct report names the input, states what it would decide, and returns no verdict on that criterion. A blank cell is not available.

3.5 The outcome

The matrix supports a classification of the configuration, and the distinction it draws is the one a reader has to act on.

State	What it means	What safety rests on
Closed by geometry	The configuration keeps the load below the adopted level with no measure required.	the built form
Closable by a stated measure	It does not, but a measure of stated dimension does.	execution of that measure
Not closable within scope	No measure within the assessment's scope brings it below the adopted level.	—

The third state is a finding, not a refusal to assess. Where no measure within scope closes the gap, that is the assessment's result, and it is at least as informative as the other two — arguably more so, since it is the one that cannot be resolved by spending money on barriers.

And the classification is a statement about a computation. Given the level the reader adopts and the geometry supplied, the load either stays below that level or it does not. A reader adopting a different level obtains a different classification from the same arithmetic, and this framework has no preference between them. Nothing here permits a configuration to be occupied or withholds permission; whether a venue may lawfully operate is a question for whoever is empowered to answer it.

3.6 What the matrix is for, at design stage

Read against Section 1.3, the matrix makes the design task countable. Every element of a venue sits in one of the three states, and the work of design is to move elements from the second into the first — from safe because someone does something correctly, to safe because of the shape. What a design review delivers, at the end, is the number of elements still in the second state, together with the dimension each of them rests on.

That number is what an operator inherits, what an insurer is pricing, and what distinguishes two venues that satisfy the same criterion.

4. Application to the design workflow

This section describes how the framework enters an existing architectural process, stage by stage, without requiring the practice to change its workflow. It describes the methodology in the general case; Section 4.5 sets out what the resulting document contains.

4.1 Concept stage: screening the massing

At RIBA Stage 2, the geometry is still fluid. Massing options are being compared; the internal arrangement is indicative; the aperture positions and concourse widths are placeholders. The framework's role at this stage is not detailed analysis — there is no detail yet — but to screen candidate massings against the conditions that will stress the finished design.

The screening produces three outputs. First, a division of the massing into regions closed by their shape and regions admitting a single direction of travel; which of the two a region falls into is set by its boundaries, and not by how many people are expected to be in it. Second, an identification of the intersections and flow conflicts the massing implies — flagged rather than quantified, since the framework is explicit that its scalar does not resolve that configuration (Section 2.3). Third, for each region of the second kind, the run length a chain could occupy, which is the quantity a partition schedule will later act on.

The concept-stage output is a short technical note plus a design-team workshop. Deliverables are calibrated to the vocabulary and iteration cadence of the design team, not that of compliance documentation.

4.2 Developed design: variant comparison

By RIBA Stage 3, the massing is fixed and the design team is iterating on variants of the internal arrangement. Concourse widths are being tested; stair positions are being adjusted; aperture counts and locations are being negotiated against structure, servicing, and cost. The framework's role at this stage is comparative, and the quantity compared is not a score.

What is compared is how many elements each variant leaves resting on a measure. Every element of a variant is closed by its shape or it is not; where it is not, its safety rests on something someone does correctly, and that is a residual the operator carries for the design life of the building (Section 1.3). Counting those elements is a comparison the design team can act on, and it does not require either variant to be scored against a threshold.

The trades that follow are ordinary design trades and can be stated as such: this partition removes a bounded run at the cost of a four per cent reduction in seat count; that aperture position converts a single-direction approach into an open one at no cost in area but a cost in servicing. What the framework contributes is the first half of each sentence — that the run is bounded, and that the change removes it — with the second half priced by the team.

The Stage 3 output is a comparative report covering each variant, plus a design recommendation. The recommendation is expressed in the framework's own terms but is translated in the accompanying workshop into architectural language the design team can act on.

4.3 Technical design: the partition schedule

At RIBA Stage 4 the geometry is essentially fixed and the analysis moves from massing to detail. The principal deliverable is a partition schedule: for each region that admits a connected run, the spacing adopted, the load that spacing guarantees, and the premise the guarantee rests on — characteristically that the partition is kept clear (Section 2.4).

Any element identified at Stage 3 as resting on a measure is resolved here in one of two ways. Either a local geometric adjustment closes it — a widening, a repositioning, a partition — and it moves into the first state; or it does not, and it is carried to Stage 5 as an element whose safety rests on execution. Which of the two occurred is recorded for each element, because that record is what a later reader needs in order to establish whether the premise held.

4.4 Handover: what the operator inherits

At RIBA Stage 5 the design team's engagement with the framework terminates and the operator's begins. The Stage 5 deliverable is the list of elements whose safety rests on execution, each stated as a dimension rather than as an instruction. Where the operator works to a DIM ICE grid, these are the entries that populate it (Section 2.5); the framework does not require that they be organised that way, only that each element arrives with the dimension it depends on.

A dimension can be recorded; a method can only be described. A duty expressed as managing an approach appropriately cannot be checked afterwards, because nothing about it is falsifiable. A duty expressed as a maximum length and a minimum interval can be: it either was or was not met. The distinction matters at handover more than anywhere else, because it determines whether the operator's own records will later show what was done. A record of activity — staff deployed, briefings held, barriers positioned — records that something was done. It does not record that the dimension was met, and only the second is what the safety rests on.

The handover deliverable is legible to two audiences. For the design team, it documents which findings were closed by design — a record of design as an act of quantitative reasoning. For the operations team, it is the starting brief for the operational safety plan. The design team's obligations end at the handover; the operator's begin. Neither is asked to inherit the other's vocabulary.

4.5 What the report contains

The document produced at each stage has the same anatomy, and it is worth setting out because each part answers an objection the next part would otherwise attract.

A drawing that defines the scope. The segments shown are the segments analysed; anything not drawn is outside the scope. Where the actual space differs from the drawing, or is used in a manner the drawing does not cover, the conclusions do not apply to that situation. This is stated at the front rather than in a limitations section, because it governs everything after it.

The matrix. Nine cells, each written in the three layers of Section 3.3, with the determination stated for each and missing inputs recorded as missing rather than left blank.

A sensitivity, run against two criteria at once. This is the part that establishes the framework is measuring something the existing criteria do not. The demonstration is symmetrical: a change in the expected population can move a venue from compliant to non-compliant without changing a single figure in the framework's analysis, and removing a boundary can close the framework's finding without changing a single figure in the compliance calculation. Neither observation criticises the other instrument. They establish that the two bound different things.

The configuration that closes it. The finding expressed as a dimension — a spacing, or a maximum length and a minimum interval — together with what is given up if the dimension is not adopted. The report supplies the dimension and not the method: any particular scheme achieves the grouping under some conditions and fails under others, and a scheme carried out at seven parts in ten is not the scheme. How the dimension is met is for the venue, which knows its own constraints; a criterion that can be checked is worth more than a method that must be trusted.

Where the safety rests. Each element marked as closed by its shape or resting on a measure, with the classification of Section 3.5 applied to the configuration as a whole.

The gaps, and where their cost falls. Stated rather than worked around, and separated into those that are properties of bodies — measured once, and thereafter available to everyone — and those that are properties of the site (Section 5).

A worked instance of this anatomy is available as a sample deliverable under the engagement structure of Section 6: a pre-event assessment of a fictional venue, carrying every figure from two declared quantities and containing no claim that the venue is safe or unsafe.

5. Where the framework stops

A framework that states its gaps is easier to check than one that works around them, and the gaps here have a structure worth setting out — not least because it is a more favourable structure than it first appears.

5.1 Three quantities are missing, and they fall on two axes

The per-contributor exertion has no direct measurement. The C-value is declared rather than measured. No instrument reports it, and no experiment in the literature isolates it (Tao, 2026a). What follows from this is narrower than it sounds: because the guarantee is exactly linear in the declared value, an error of a given proportion produces an error of the same proportion in the spacing and no more, so a report can state the spacing required at each of several declared values and leave the declaration with the party who knows what the event is.

The residual a standing body absorbs is not quantified. This is the quantity the growth condition compares against, and it therefore decides whether a run accumulates at all. Earlier formulations carried it at zero on the conservative side, which is sound where it is used as a bound and unsound where it is used as a criterion: at zero the comparison is satisfied trivially (Tao, 2026a). It is a mechanical quantity, measurable under controlled loading, with no ethical obstacle.

The rate at which a region fills is not supplied. The framework declines to predict it, and the refusal is principled rather than provisional: an indicator must not depend on a quantity the framework does not possess (Tao, 2026a). A treatment of the filling rate is a separate undertaking of comparable size to the framework itself.

All three fall on the force and time axes. The space axis — the one a building controls, and the one this document has been about — is complete: it asks whether the geometry admits a connected run of a given length, and that question is answered from the drawing by arithmetic no part of which is in dispute here.

5.2 What each gap would close, and where its cost falls

Quantity	Obtained by	Would enable	Cost falls on
Per-contributor exertion for a class of scene	contact-force measurement under controlled loading; once per class	replacing a declared value with a measured one	**research, once**
Residual absorbed by a standing body	controlled measurement; general	the growth condition evaluated as a criterion rather than a bound	**research, once**
Area a recovery step requires	controlled measurement; general	the density at which a region enters continuous contact	**research, once**
Net inflow at a boundary	one counting line	whether accumulation is occurring, in sign	**the venue, per site**

Three of the four are not a venue's to pay for. They are properties of bodies rather than of any site, they are measured once, and once measured they are available to everyone. A venue is not being asked to fund them, and a document presenting them as site costs would be misrepresenting where the work lies.

The fourth is, and it is modest. A count at the upstream mouth of an approach is the same instrument already installed at a gate line, placed earlier. Occupancy behind a boundary is the difference between what entered and what left, so it needs two counts and a subtraction; where only one exists, the occupancy is not poorly known but undefined by the available data. What the framework needs from that pair is a sign rather than a magnitude — whether arrivals exceed admission — and a sign survives a common bias in both terms.

And how the missing values are obtained determines whether the result can be used. A physical experiment yields a method others can repeat and a figure a reader can check. A simulation yields a figure sooner, but where the simulation is not public it yields an input the reader cannot check, and on this framework's own rules such a figure cannot support a design-side certificate (Tao, 2026a).

5.3 A density instrument is not useless here; it is uncalibrated

It would be easy to read the foregoing as an argument against monitoring, and it is not one. Occupancy over time is computable: it rises at the difference between arrival and discharge, and both are schedulable quantities. A reading of occupancy is a real measurement of a real quantity and it locates a region on the path from free movement to continuous contact.

What is missing is the mark on that path — the density at which contact becomes continuous. Without it the computed curve cannot be told where it crosses into the state this framework is concerned with, and a density threshold adopted in its absence is a number compared against a convention. The instrument will report faithfully against a figure that has not been derived.

This is worth separating from the finding of Section 2.4, which is that the *load* does not respond to occupancy. Both hold. Occupancy is not what the mechanism responds to, and it is nonetheless the quantity that says which phase a region is in — once the mark is placed.

5.4 The choice this presents

The missing quantities are not an argument for waiting. They set out a choice, and it is the choice a venue is making whether or not it is stated.

One option is to obtain the information. Instrument the approach, record what it reports, and act on it. This is a real option, it becomes a better one as the research quantities above are measured, and nothing in this framework argues against it. Sensing and video analysis are advancing, and a determination not observable today may be observable in a few years.

The other is to configure so that the information is not needed. A region in which the mechanism cannot operate requires no reading, no threshold, and no response. It is safe in the sense this framework can state, at any occupancy the exits permit, whether or not anyone is watching.

The two differ in a respect that outlasts the instrumentation. An instrumented approach is safe while the instrument is read and acted on; a configured one is safe because of its shape. The first has an execution rate and the second does not, and Section 1.3 sets out why a lapse in the first is not recovered by resuming it.

5.5 One caution about undertakings

Where a permission rests on an undertaking to observe something, the undertaking becomes part of what makes the configuration acceptable. The observation is then not an optional refinement but the thing the permission was granted against.

An undertaking of that kind which is not performed leaves the configuration in a worse position than one that was never given, because the alternative measures that might have been required instead were not required — precisely because the undertaking was given.

This document makes no observation about any particular case. It notes the structure because it bears directly on the distinction of Section 1.3: an undertaking to observe is an operational measure, it carries an execution rate, and what it displaces may have carried none.

6. Engagement structure

This section states what is available, in what form, and what falls outside it. It is written as a scope rather than as a prospectus. The boundary between what the framework supports and what it does not is set by the framework's own gaps (Section 5), and a scope that omitted them would be describing a different instrument.

6.1 The dividing line: whether a model is run

Two kinds of work are distinguished, and the distinction is not one of depth. A closed-form assessment answers what a geometry permits, from the drawing and a declared exertion, by arithmetic a reader can check by hand. A modelled assessment computes a trajectory, and requires a model to be run. They differ in price and in delivery time by an order of magnitude, which is why the line is drawn here rather than left as a matter of scope.

Product	Basis	What it delivers	Instance
Pre-event geometric assessment	closed form; no model	load bounds, grouping and partition dimensions, three-state classification	sample deliverable available
Retrospective backtrack analysis	closed form; retrospective	whether a geometry as built and occupied suffices to produce a stated load	published instance
Venue P-report	a model is run	the trajectory and magnitude of P	none yet

The third row carries no instance, and the name is reserved rather than retired. P is not a design-stage indicator and takes a positive value only in reconstruction (Tao, 2026a); the first two products therefore do not compute it at any point. A product named for the index has to be the one that computes it, and that one requires a model.

Geometry input for the modelled product. Geometry is accepted in a documented, human-auditable TOML schema listing enclosure boundaries, aperture positions, stair geometries, and internal partitions. The schema is text-based rather than binary: legible in any editor, tracked in any version-control system, and adjustable live during a design-review session without a round trip through BIM authoring tools. Client BIM teams working in Revit, ArchiCAD, or Vectorworks export the relevant IFC entities and populate it. Native IFC ingestion is on the development roadmap. The closed-form products do not require the schema; their input is the drawing.

6.2 What a closed-form assessment delivers

Three findings are available, ordered here by how little each of them requires.

The mechanical image of an existing standard. A density, a width, or a timing requirement in the governing guidance implies a per-person load, and the conversion states which. This finding requires none of the missing quantities of Section 5.1 to be supplied, and it answers the question a venue operator asks first (Tao, 2026a): what load is the standard I already meet equivalent to? It compares two instruments; it does not determine compliance, which is not this framework's to determine.

A load guarantee at the geometric layer. For each region admitting a connected run, the partition spacing adopted, the load that spacing guarantees, and the premise the guarantee rests on — characteristically that the partition is kept clear (Section 2.4). The figure follows from the drawing and the packing limit; apart from the declared exertion, nothing else enters, and the sensitivity of the result to that declaration is reported with it.

The three-state element list. Each element classified as closed by its shape, closable by a measure of stated dimension, or not closable within the assessment's scope (Section 3.5). Its commercial form is a count — how many elements still rest on someone doing something correctly — and that count is what an operator inherits and an insurer is pricing (Section 3.6).

6.3 What lies outside the scope, stated as scope

Four questions are not answered. They are set out here, and in the engagement's terms, rather than in a limitations paragraph at the back, because a scope that a client discovers late is a scope that was not agreed.

When it will happen. The rate at which a region fills is declined, on the principled ground that an indicator must not depend on a quantity the framework does not possess (Tao, 2026a). What is supplied is the state once filled, not the hour at which it arrives.

The density at which to act. The density at which contact becomes continuous has not been derived, so a density instrument here reports faithfully against a mark that has not been placed (Section 5.3). A threshold adopted in its absence is a number compared against a convention.

The quantitative treatment of intersecting streams. Where two directed flows meet on a shared axis, the framework is explicit that its scalar does not distinguish that state from a benign one (Section 2.3). Such configurations are identified and flagged; no figure is reported for them, because a figure reported there would be one the framework does not support.

Whether the venue is safe. The framework carries no quantity on which such a determination could rest. What is delivered is a load against a level the reader adopts, and the classification that follows from it. Nothing in an assessment permits a configuration to be occupied or withholds permission.

6.4 Which side of the matrix an engagement fills

The three-axis matrix of Section 3 has a two-sided first layer — what the venue supplies, and what the event demands — and an engagement fills one side or both. A venue assessment fills the venue side across all nine cells; an event assessment fills the event side; a venue promoting its own event fills both.

The first two are not halves of the third. One limitation must be stated in either of them: the time axis cannot be closed from one side alone, since the rate at which a space fills depends on the venue's geometry and the event's arrival profile together. A single-sided assessment is structurally incomplete rather than proportionally smaller, and it is reported as such.

6.5 What is required from the client

For a closed-form assessment: the drawing of the region to be assessed, with the boundaries and any partitions as they will exist on the day rather than as designed; the operational description sufficient to say which regions are occupied and in what manner; and one nominated point of contact for the engagement's duration.

Two values are declared rather than measured, and the declaration belongs with the party who knows what the event is. The first is the per-contributor exertion, for which the assessment reports the result at each of several declared values rather than selecting one (Section 5.1). The second is the level against which the load is read, which the client, its insurer, or its jurisdiction nominates; the analysis is unchanged in form whichever is chosen (Section 2.1).

6.6 Design-stage engagement

The design-stage form set out in Section 4 is described in this document as a capability of the framework. It is not currently offered as a service. The order of work is deliberate: the assessments above require no qualification to perform and carry the lower risk, and a design-stage engagement — which enters a live design process and is relied upon within it — is taken up after that work has established a record, not before.

Scope and fee. Both are set per configuration, against the extent of the region assessed and the number of scenarios carried. Neither is quoted here, because the quantities that determine them are the ones the assessment itself establishes.

7. References and further reading

The physical derivation, validation, and complete bibliography underlying this white paper are carried by the open-access source manuscripts:

- B01 — flagship manuscript, Zenodo DOI `10.5281/zenodo.21261698` (Tao, 2026a).
- E01 — bipartite nine-cell framework, Zenodo DOI `10.5281/zenodo.21266645` (Tao, 2026d).
- Kroll, Still, Neuman, Graham and Griffin (2017) — biomechanical model of compression asphyxia (Kroll et al., 2017).
- Helbing et al. (1995, 2000, 2007) — Social Force Model and its operating regime (Helbing and Molnár, 1995; Helbing et al., 2000, 2007).

References

Peter A. Cundall and Otto D. L. Strack. A discrete numerical model for granular assemblies. *Géotechnique*, 29(1):47–65, 1979. doi: 10.1680/geot.1979.29.1.47.

Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, 1995.

Dirk Helbing, Illés Farkas, and Tamás Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(6803):487–490, 2000.

Dirk Helbing, Anders Johansson, and Habib Zein Al-Abideen. Dynamics of crowd disasters: An empirical study. *Physical Review E*, 75(4):046109, 2007.

Mark W. Kroll, G. Keith Still, Tom S. Neuman, Michael A. Graham, and Larry V. Griffin. Acute forces required for fatal compression asphyxia: A biomechanical model and historical comparisons. *Medicine, Science and the Law*, 57(2):61–68, 2017. doi: 10.1177/0025802417695711.

Paul A. Langston, Robert Masling, and Basel N. Asmar. Crowd dynamics discrete element multi-circle model. *Safety Science*, 44(5):395–417, 2006. doi: 10.1016/j.ssci.2005.11.007.

Trushant S. Majmudar and Robert P. Behringer. Contact force measurements and stress-induced anisotropy in granular materials. *Nature*, 435(7045):1079–1082, 2005.

G. Keith Still. *Crowd Dynamics*. PhD thesis, University of Warwick, 2000.

G. Keith Still. *Introduction to Crowd Science*. CRC Press, 2014.

Yixuan Tao. A physical quantification framework for crowd compression risk: The p-index system and its contact-mechanical foundation. Zenodo, 2026a. URL `https://doi.org/10.5281/zenodo.21261698`. Concept DOI (all versions); current release v3.

Yixuan Tao. Event C-value quantification: Completeness of the C-value for external crowd-safety factors. PneumaTheorem, Inc. Technical Report, 2026b. Preprint. Zenodo DOI pending.

Yixuan Tao. Anchor-based similarity analysis for forensic and regulatory crowd risk assessment. PneumaTheorem, Inc. Technical Report, 2026c. Preprint. Zenodo DOI pending.

Yixuan Tao. A bipartite nine-cell framework for crowd safety analysis: Venue p-report and event dim ice as complementary instruments. Zenodo, 2026d. URL `https://doi.org/10.5281/zenodo.21266645`. Concept DOI (all versions); current release v1.

Yixuan Tao. The p-index framework: A physics-based system for crowd compression risk quantification — commercial white paper. PneumaTheorem, Inc., 2026e. Version 1.

Yixuan Tao. Closing the quantification gap in crowd safety reporting — logic before data: A side-by-side reference implementation. PneumaTheorem, Inc., 2026f. Methodology white paper. Available on request.

Author: Yixuan Tao
Correspondence: kevin@pneumatheorem.org
ORCID: 0009-0005-8136-3886
PneumaTheorem, Inc. · Delaware, United States

© 2026 PneumaTheorem, Inc. Some rights reserved.

This work is licensed under a Creative Commons Attribution–NoDerivatives 4.0 International License (CC BY–ND 4.0). You are free to share it, in any medium or format, for any purpose, provided that it is passed on unchanged and in whole and that appropriate credit is given.

This white paper is not deposited on Zenodo and carries no DOI. The flagship manuscript remains the canonical citation reference.