Reconstructing the Fluid–Structure Interaction of an Earthquake-Triggered LP-Gas Deflagration and Progressive Collapse in a Two-Storey Shopping Mall in Kumamoto, Japan, on 28 July 2026

Jun Ishimoto

Institute of Fluid Science, Tohoku University, 2-1-1 Katahira, Aoba-ku, Sendai 980-8577, Japan

Corresponding author: ishimoto@alba.ifs.tohoku.ac.jp — ORCID iD: 0000-0002-5674-643X

Preprint, version 2, 23 September 2026 (supersedes version 1 of 5 August 2026)

Notice of correction (version 2)

This is version 2 of the preprint first posted on engrXiv on 6 August 2026 (DOI 10.31224/7880). It corrects version 1; it is not a revised paper. The text, structure, tables and figures of version 1 are kept except where they were wrong, and every change is listed in the Appendix, "Changes from version 1". A substantially revised version of this work has been submitted for journal publication.

Two principal claims of version 1 are withdrawn. (1) Version 1 stated that "with intact member capacities the same combustion produced markedly less damage". No simulation with intact member capacities was performed. The statement is replaced by what the computation does support: a static comparison of the computed pressure record with the intact capacities, which is not a coupled simulation. (2) Version 1 presented, as an unadjusted independent validation, an impulse duration of 22 ms back-calculated from the observed debris and agreeing within 12% with the 25 ms assumed a priori. That back-calculation used roof-tile failure pressures of "$\Delta p \approx 5.6$–24 kPa with median ≈ 10 kPa", which match no computed record; the recorded values are 4.1–10.8 kPa with a median of 4.8 kPa, which would give about 46 ms. The agreement is withdrawn.

Further statements are withdrawn or narrowed; the Appendix lists all of them. Among them: version 1 listed the collapse morphology among the independently reproduced features, but the collapse mechanisms, their region and their timing were prescribed to represent it; of the roof breach, only the position agrees with the observation, the simulated extent being about an order of magnitude too large; version 1 argued that, because the observed damage could be explained only with members already weakened by the mainshock, the loading was not of detonation order, an argument that is circular because the loads come from a deflagration model; and version 1 attributed all panel capacities to Japanese design references, whereas only the composite slab and its connections follow them and the other capacities are the author's own derivations from typical details (the values are unchanged).

After version 1 was posted, a pressure probe lying on the boundary of a removed solid region was found to return an undefined value, which opened the panel it monitored at t = 0.02 s although that panel had not failed: one partition in each of the two 1.0 m runs (central and south-edge ignition) and three interior panels on the 0.5 m grid. All runs were repeated with a guard that keeps such a panel closed, and every computed value in version 2 is from those corrected runs; the only exception is the rendering in Figure 2(j–l), which still shows the earlier south-edge computation, as its caption states. The relocated-ignition run, called the corridor run in version 1, is called the south-edge run, since the model has no corridor at that position. On the 1.0 m central-ignition run, the effect on the damage was confined to the spuriously opened partition (panels failed by direct pressure 79 → 78) and shifts in some failure times, but the ignition-kernel peak moved from 34 to 43 kPa, and it is now reported without quantitative interpretation. On the 0.5 m grid, the defect changed the failure pattern (panels failed 100 → 95). The south-edge result also changed (panels failed 60 → 62, roof tiles 11 → 8, contiguous both-storey stripping 30 → 20 m), so that this run now fails fewer roof tiles than the central run rather than more. Errors of version 1 unrelated to the defect are also corrected, among them the burning velocity of the lean ceiling layer (about 8 → 3.6 m/s) and the load of the slab check (16.2 → 12.15 kPa, peak fibre stress 30.6 MPa).

Unchanged: the contiguous facade-stripping length of 60 m, one of the two calibration targets of the single calibrated parameter (the effective burning velocity), is reproduced by the corrected central-ignition runs on both grids; with the ignition point moved to the south edge, the second-storey band is again 60 m, but the both-storey length is 20 m (Section 4.1). The qualitative conclusions—damage concentrated on the south side, a roof breach above the ignition zone, interior loads of a few kilopascals acting on seismically weakened members, and a fast deflagration rather than a cloud detonation—stand.

Abstract

On 28 July 2026, approximately 80 minutes after the Mj 7.1 Kumamoto earthquake (maximum JMA seismic intensity 7), a large gas explosion destroyed part of a two-storey steel-frame shopping mall in Kashima Town, Kumamoto, Japan. Seven fatalities had been reported as of 30 July 2026. Public reporting indicates that the earthquake damaged the supply piping from an outdoor bulk LP-gas vessel and that leaked propane accumulated in a second-floor retail zone before igniting. This paper reconstructs the event using a three-stage fluid–structure interaction (FSI) chain built entirely from open-source tools: (i) a 3-D premixed-deflagration CFD model of an approximately 330 kg stratified propane cloud in a 160 m section of the mall (constructed with OpenFOAM XiFoam using Weller b–Xi flame-wrinkling closure and k–omega SST) ; (ii) a quasi-two-way venting FSI scheme in which 255 breakable envelope and floor panels are monitored during the run and opened through a deterministic mesh-rebuild restart when their failure pressure is exceeded, with capacities derived from Japanese design references or typical details and degraded by an earthquake damage factor; and (iii) an event-driven progressive-collapse and debris/dust dynamics model that distinguishes a column-first storey-crush (pancake) core from a hinge-type slab-unzipping periphery. The effective turbulent burning velocity is calibrated using two observables: the contiguous facade-stripping extent and the near-frame-one envelope failure, which are therefore calibration targets rather than validation metrics. With that single calibration, the reconstruction independently reproduces the absence of a sustained external fireball, the radial debris scatter on the tank-side facade, the position of the roof breach above the ignition zone, and a dust-front speed of 8.9 m/s against a measured 13 m/s (band 8–17 m/s). The debris model is further tested, without any adjustment, against the ballistics recorded in the footage: it reproduces the observed apex band. However, it does not reproduce the longest observed throw distances, a deficiency reported here rather than removed by tuning. A uniform 0.5 m grid-refinement study shows that the facade-stripping extent and the sustained compartment overpressure are essentially unchanged, while the external blast spike and the local failure counts are not convergent and are reported with that caveat. An ignition-location sensitivity run, in which the kernel is moved to the south-edge position suggested by post-event drone imagery, leaves the headline conclusions—south-side damage concentration, a roof breach above the epicentre, and few-kilopascal interior loads acting on seismically weakened members—unchanged. A numerical pitfall is that the subsetMesh utility of OpenFOAM-v2606 assigns the constraint type “empty” to newly exposed patches, which silently turns off momentum solution on 3-D meshes. All findings are hypotheses for the official investigation and imply no attribution of responsibility. Based on damage magnitude, impulse, and debris velocity, the reconstruction supports a fast deflagration rather than a cloud detonation, though a transient, localized transition cannot be excluded.

Keywords: gas explosion; vented deflagration; vapor cloud explosion; Natech; fluid–structure interaction; progressive collapse; CFD; forensic engineering; OpenFOAM; earthquake-triggered accident; LP gas; combustion regime; deflagration-to-detonation transition

1. Introduction

Earthquake-triggered gas releases are a recognized cascading hazard (Natech): ground shaking damages fuel-gas infrastructure, and delayed ignition of the accumulated cloud can produce structural damage far exceeding that of the shaking itself. The catastrophic explosion at a shopping mall in Kumamoto Prefecture on 28 July 2026 is a severe example: dashcam recordings carry an on-screen timestamp of 17:46 for the explosion (a device clock reading not independently calibrated), about 80 minutes after the mainshock, and the operator's public release reports the blast location as "the rear of the second floor". Aerial footage shows roughly 55 m of the front facade stripped to the frame, a local roof breach, and a partial collapse of the second floor in which floor plates dropped between largely standing columns, with a heavily crushed core zone.

Forensic reconstructions of building gas explosions have historically informed both regulation and structural engineering, from the Ronan Point town-gas explosion of 1968 [9] to modern vented-deflagration test programs [6] and code provisions for internal explosions [8]. Computational reconstructions, however, rarely close the gap between the gas-dynamic loading and the progressive structural failure that governs both the venting history and the final damage state: the opening of windows, cladding, and floor plates feeds back strongly on the achievable overpressure.

This study contributes: (1) quantitative evidence extracted from public video (the cloud kinematics, debris ballistics, damage extents, and timeline); (2) an open-source, reproducible three-stage fluid–structure interaction (FSI) chain coupling a premixed deflagration solver to a panel-resolved venting model with seismically degraded capacities, and to an event-driven progressive-collapse model that represents both storey-crushing and slab-draping morphologies; (3) a comparison against observables that were *not* used for calibration — the dust-front kinematics, the absence of a sustained external fireball, the rear-facade debris scatter and, in particular, the debris ballistics, which test the fragment model without any parameter being fitted to the footage; (4) a grid-sensitivity study on a uniformly refined 0.5 m mesh that separates the conclusions that survive refinement from those which do not; and (5) documentation of the silent solver-crippling behaviour of a standard OpenFOAM mesh utility that is directly relevant to any explosion-venting workflow using cell-removal geometry. The reconstruction is based exclusively on public information; building dimensions and inventories are typified where unknown.

2. Incident Evidence from Public Sources

2.1 Timeline and mechanism as publicly reported

The mainshock occurred at 16:27 JST on 28 July 2026 with Mj 7.1 and a maximum seismic intensity of 7 [14], where Mj denotes the magnitude scale of the Japan Meteorological Agency (JMA), and the JMA seismic intensity is a closed scale running from 0 to 7 that is distinct from the Modified Mercalli scale. According to the operator's press release of 29 July, customers and staff were evacuated outdoors by about 17:00, and the post-quake inspection confirmed ceiling falls, partial exterior wall damage, and sprinkler releases. The explosion occurred at "the rear of the 2nd floor" at 17:50 as reported by the operator, while dashcam timestamps read 17:46:24 [15]. The two times differ by about four minutes. Because the dashcam clock is not independently calibrated and the operator's time may be a report-receipt time, no attempt is made here to resolve the discrepancy, and the accumulation interval is treated as approximately 80 minutes in either case. A television documentary aired on 3 August 2026 [15] reports, based on drone surveys of the site, the operator's statement that the explosion was "concentrated in one corner of the south side"; identifies a staff-only back-of-house corridor running behind the second-floor shops as the suspected accumulation space; and notes an expert assessment that rooftop cooling units directly above that corridor were thrown upward "by pressure from below", marking the estimated epicentre. The same programme also reports that a gas odour had been noticed inside the store before the explosion.

Broadcast commentary raised the question of whether seismic shutoff provisions on the bulk LP-gas installation operated as intended; this point is reportedly under official investigation and is not addressed by the present study [15]. What is used here as a modelling premise, and only as a premise, is the publicly reported outcome: that a sustained release from the damaged supply line accumulated in a second-floor fashion-retail zone (not a restaurant zone) over the interval between the mainshock and the explosion. As of 30 July 2026, seven fatalities had been reported; this figure was still subject to change at the time of writing.

2.2 Quantitative video analysis

Frame-by-frame analysis of two dashcam clips (30 fps) and extended aerial footage was performed using utility poles (with a visible height of 10.5–12.5 m) and vehicles as length references at the cloud distance; the absolute scale uncertainty is estimated at ±35–40%. Table 1 summarizes the observables used as reconstruction targets. Notably, the very first post-onset frame (33 ms) already shows a ~27 m-wide eruption above the roofline, indicating near-simultaneous envelope failure over tens of meters. It should be emphasized that this observation constrains the internal *pressure field* rather than the flame speed: pressure equalization within a compartment proceeds on the acoustic timescale (~340 m/s, i.e. about 0.15 s across the cloud length, short compared with a burn of order 1 s), so the failure of panels tens of metres apart within a few frames does not require the flame front to have traversed that distance. No sustained fireball appears in any clip. A close-range recording aired on 3 August [15] shows only brief, localized orange luminosity at the base of the erupting cloud, lasting approximately one second; this is consistent with transient flame jets at the vent openings and with an inside-burning lean-to-stoichiometric LPG deflagration venting combustion products and entraining structural dust.

Table 1. Reconstruction targets and contextual evidence from public sources (the ±35–40% scale uncertainty applies to footage-derived lengths and speeds). The last three rows are contextual evidence used qualitatively in Sections 4.1 and 5, not simulation targets.

Observable	Value	Basis
Explosion onset	17:46:24 JST (dashcam clock); 17:50 (operator release)	timestamp overlay
Initial dust-front speed	13 m/s (range 8–17)	first 0–0.6 s, single-clip differencing
Late front speed	4 m/s (range 3–7)	t = 4–7 s
Roof-debris launch speed	30–40 m/s vertical; apex ~24 m (band 15–33 m)	ballistic fit, gravity-consistent (no slow-motion)
Facade fully stripped	~55 m contiguous (±15 m), stripped to the frame	aerial, parking-stall/vehicle scale
Facade partially damaged	~95 m total	aerial
Roof breach	local, 10–15 m across, deck folded inward	aerial
Collapse morphology	floor plates dropped between standing columns; heavily crushed core	aerial close-ups
External fireball	no sustained fireball; brief base flashes in one close view	both angles, 0–14 s; 3 Aug broadcast close view [15]
Debris throw	steel sheeting to ~100 m (lower bound)	aerial

Observable	Value	Basis
Vehicle damage	one truck driving past the mall wrecked by the blast; dashcams in the nearby queue recorded continuously	broadcast reports [15]
Rooftop equipment	cooling units above the epicentre thrown upward; roof holed (pre/post drone comparison)	Drone survey, 3 Aug broadcast [15]
Estimated epicentre	2F back-of-house staff corridor, south edge	Expert assessment based on drone imagery [15]

3. Reconstruction Method

3.1 Geometry and gas-accumulation scenario

A 160-m-long section of the two-storey steel-frame mall was modelled on an integer-metre lattice (Fig. 1) and consisted of a 10 m column grid, second floor at $z = 6$ m, roof deck at $z = 11$ m, parapet 12 m high, south tenant strip (30 m deep), 10 m central corridor, north strip containing a statically walled cinema box and food-court zone, and 1 × 2 × 3 m merchandising obstacles in the tenant strip (congestion). The baseline computational domain was 200 × 110 × 26 m, discretized with 1.0 m hexahedral cells (0.57 M cells). A uniformly refined 0.5 m mesh was used for the sensitivity study of Section 4.3. The stratified propane cloud occupied the tenant strip over $75 < x < 125$ m, with fuel mass-fraction layers in the order 0.058 / 0.045 / 0.030 from floor to ceiling, corresponding to equivalence ratios of approximately 0.96, 0.74, and 0.48 for a stoichiometric propane mass fraction of 0.060 (near-stoichiometric to lean with a propane vapor density 1.55 times that of air). The cloud integrates to 325.5 kg on the 1.0 m grid and 331.0 kg on the 0.5 m grid—a 1.7% difference arising from cell-centre snapping at the $z = 8.5$ m layer boundary—and is quoted as approximately 330 kg throughout. This order of magnitude is consistent with tens of minutes of release from a damaged bulk-vessel supply line. Ignition was a 2.5 m spherical kernel at floor level in the cloud centre. The 80-minute accumulation itself was not simulated. The cloud was an initial condition, and its extent was a hypothesis.

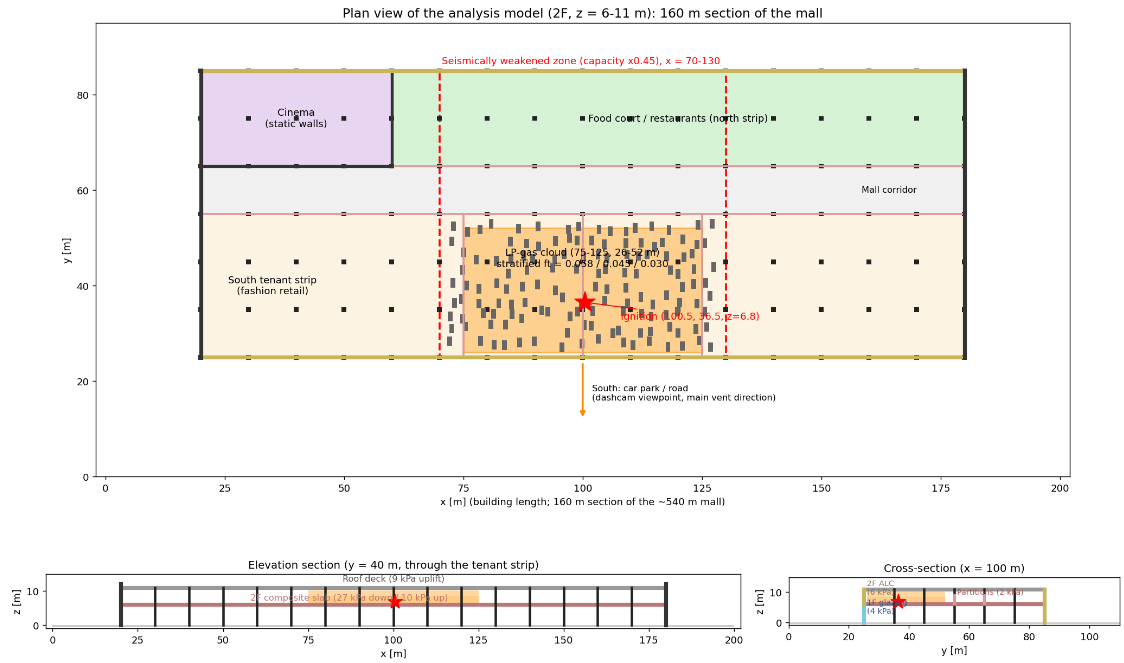

Figure 1. Analysis model: plan (top), elevation y = 40 m (bottom left), cross-section x = 100 m (bottom right). Orange: stratified LP-gas cloud; red dashes: seismically weakened zone (capacity × 0.45); star: ignition.

3.2 Deflagration CFD

The premixed deflagration was computed with the XiFoam solver of OpenFOAM-v2606 [4]. The computation used a Favre-averaged compressible flow, regression variable b with Weller flame-wrinkling closure [1], unstrained laminar flame speed from the Gülder correlation for propane [2], heheuPsiThermo two-state thermodynamics with an inhomogeneous fuel/oxidant/product mixture with JANAF C3H8 properties and a mixture fraction capped just below stoichiometric to remain within the model's lean-side validity, and the k–omega shear-stress transport (SST) turbulence model [3] (k–epsilon was found to be non-robust for impulsive venting jets at this resolution, with runaway turbulent viscosity). The boundary conditions were a wave-transmissive far field and wall functions on all solid surfaces. The time step was limited to 0.25 ms on the baseline grid (acoustic Courant–Friedrichs–Lewy number ~0.09) and to 0.125 ms on the refined grid.

At 1 m resolution, a Reynolds-averaged Navier–Stokes (RANS) model cannot spontaneously reproduce the obstacle-driven flame acceleration observed in a congested retail interior. Following established practice in congestion-calibrated effective burning velocities in industrial VCE tools [6, 13], we prescribed the wrinkling factor as Xi = 50. Because the cloud is stratified, the unstrained laminar flame speed S_L differs between layers, so the effective turbulent burning velocity S_T = Xi x S_L is not a single value: it ranges from approximately 3.6 m/s in the lean ceiling layer (ft = 0.030) to approximately 20.9 m/s in the near-stoichiometric floor layer (ft = 0.058), slightly wider than the 4–20 m/s range reported for highly congested spaces.

This wrinkling factor is the single significant calibration parameter of the gas-dynamic stage. It was selected to reproduce two observables — the contiguous facade-stripping extent and the near-frame-one envelope failure. Those two quantities are consequently **calibration targets rather than independent validations,** and they are labelled as such in Table 5. Prescribed (assumed but not calibrated) parameters of the collapse and debris stage are listed separately in Table 3, so that the reader can distinguish computed quantities from modelling choices throughout.

3.3 Vent-panel FSI (quasi-two-way coupling)

The building envelope and floors were decomposed into 255 failure panels (Table 2), each embedded in the mesh as a zero-thickness baffle. During the run, the pressure difference across each panel was monitored with probe pairs. When the two most recent samples exceeded panel capacity, the run was checkpointed, and the panel was opened; the solver excluded any panel whose probe it could not locate (returning the sentinel value −1e300), and that panel remained shut. The opening was implemented with a deterministic mesh rebuild (blockMesh → removal of obstacles/columns from the cells • baffle creation for the surviving panels only). The cell ordering of this pipeline was verified to be bit-reproducible: the maximum cell-centre displacement between rebuilds was less than 1e−13 m on the 1.0 m grid and 8.5e−14 m on the 0.5 m grid. Consequently, the volume fields were carried over unchanged, and only face-based fluxes were re-initialized. Venting therefore feeds back on the interior pressure, with a resolution of 1 leg at 10–20 ms. Capacities followed Japanese design references (composite-deck slabs, AIJ composite-structure studs, and puddle-weld capacities [10]); the author derived the glazing, ALC rocking-clip connection, and LGS gypsum partition capacities from typical details. Inside the blast zone, they were reduced by a factor of η = 0.45, which is within the damage-class-III residual-capacity range of the Japanese post-earthquake damage-classification standard [11] and consistent with the operator-reported pre-explosion ceiling/wall damage.

Table 2. Vent-panel classes and differential-pressure capacities (Δp) for each panel. Derivations are from the author's section analysis, AIJ-based for the composite slab and its connections.

Panel class	Unit size	Failure Δp (intact)	Failure Δp (η = 0.45 zone)
1F storefront glazing	10 × 6 m	4 kPa	1.8 kPa
ALC/metal facade (both storeys)	10 × 5–6 m	6 kPa	2.7 kPa

Panel class	Unit size	Failure Δp (intact)	Failure Δp (η = 0.45 zone)
Gypsum/LGS partitions (demising walls 30 × 3 m)	10 × 3 m	2 kPa	0.9 kPa
Roof deck tiles (uplift)	10 × 15 m	9 kPa	4.05 kPa
2F composite slab tiles (down/up)	10 × 10 m	27 / 10 kPa	12.2 / 4.5 kPa

3.4 Progressive collapse and debris/dust dynamics

The third stage converted the CFD failure record into building-destruction kinematics, resulting in 1,358 rigid fragments. CFD-failed panels fragmented and flew with an impulse $v0 = \min(|\Delta p|, 40\ kPa) \times 25\ ms \times A/m$ (capped at 55 m/s, then scaled by a random factor of 0.7–1.3 and perturbed by Gaussian scatter; roof tiles received an additional upward kick of 3–10 m/s, and slab fragments instead dropped at 0.5–2 m/s) and were advected by the time-resolved CFD velocity field with quadratic drag (with a drag coefficient of 1.2). In the fragment model, the roof tiles lumped the deck together with rooftop equipment and were assigned an areal density of 50 kg m^{-2}, whereas the debris identifiable in the aerial footage was steel skin roughly 0.8 mm thick, corresponding to an areal density of approximately 6.3 kg m^{-2}. The two classes were therefore kept distinct in the ballistic comparison of Section 4.2, where the identical impulse law was applied to sheeting of the lower areal density released from the same failed tiles.

Structural collapse was modelled event-wise using two mechanisms that reflected the observed morphology. The first was a storey-crush (pancake) core at x = 90–110 m, where blast impulse plus seismic degradation was assumed to cause the columns to fail first (t = 0.35–0.7 s): each column split at the knee, the lower segment buckled outward, and the bays above — second-floor slab and roof tiles — dropped near-horizontally as units, stacking at rubble heights of ~1 m (slab) and ~2.2 m (roof). The second was a peripheral unzipping band within the weakened zone where support loss propagated tile-to-tile (0.35–0.65 s per tile), with each tile hinge-rotating about its last supported edge before detaching. These floor plates were draped between standing columns. Columns that lost three or more adjacent slab tiles within a prescribed plan region were set to topple; this rule did not activate in any reported run. Failed panels emitted dust for 0.8 s after failure, and ground impacts of heavy fragments released further dust bursts. The dust, capped at 260,000 particles, was advected by the CFD velocity field (which faded within 0.5 s after the 1.5 s CFD record ended), with settling and stochastic dispersion. Rendering was done in ParaView in batch mode.

The quantities in Table 3 are prescribed modelling choices, not results. They are stated explicitly because they govern the timing and visual character of the collapse. They did not enter the gas-dynamic solution of Section 3.2, which is one-way coupled to this stage. The firstof them —the 25 ms impulse duration—was subsequently tested against the footage in Section 4.2 and was not adjusted there.

Table 3. Prescribed parameters of the collapse and debris/dust stage (assumed values, not calibrated against the video record).

Parameter	Value	Basis/status
Fragment impulse duration	25 ms	assumed a priori; independently tested in Section 4.2
Impulse pressure cap	40 kPa	assumed; limits the near-kernel transient
Fragment speed cap	55 m/s	assumed
Drag coefficient/drag law	1.2/quadratic	assumed
Areal density, roof-tile class	50 kg m^{-2}	deck plus lumped rooftop equipment
Areal density, sheet class	6.3 kg m^{-2}	0.8 mm steel skin matched to the debris visible in the footage
Column failure window (crush core)	0.35–0.7 s	prescribed to reproduce the observed crush morphology
Slab unzipping delay	0.35–0.65 s per tile	prescribed
Column toppling criterion	loss of ≥ 3 adjacent slab tiles	assumed
Rubble stack heights	~1 m (slab); ~2.2 m (roof)	assumed from plate thickness and bay area

3.5 Verification and a numerical pitfall

The chain was verified stepwise:

1. Sealed-volume combustion reached the theoretical constant-volume pressure ratio (~9).

2. A vented pilot enclosure reproduced the canonical vented-deflagration signature (a pressure rise during burning outpaced venting, peaking, and the negative phase) [5, 7].

3. The simulation software solids4foam [12] provided an independent linear-elastic solution for a 10 m composite-slab strip under the degraded failure load (12.15 kPa), yielding a peak bending stress of 30.6 MPa. This is approximately 12 times the

cracked-RC tensile capacity, consistent with the tabulated connection-governed capacities, not gross-section strength, controlling (Fig. 6).

4. Shock/acoustic transmission through opened panels was verified in minimal cases.

5. The ballistic integrator was checked against the closed-form quadratic-drag rise height $h = (\mu/2k) \ln(1 + k v0^2/\mu g)$, where $k = ½ \rho_{air} C_d$, with agreement within 0.15% for every fragment tested. Section 4.3 assesses grid sensitivity.

During this work, a silent failure mode of OpenFOAM-v2606 was isolated that practitioners of explosion-venting simulations should be aware of: subsetMesh assigns the constraint type "empty" (intended for 2-D cases) to the patch that receives newly exposed faces. On a 3-D mesh, empty faces corrupt solution-direction detection, and the momentum equation is never solved: pressure then spreads only diffusively, velocities remain at the single-step value $u \approx -\Delta t \nabla p/\rho$, and unphysical static high-pressure domes persist in open air; checkMesh reports such meshes as healthy. The fix is to retype the exposed patch as a wall before use. The archived dataset includes the symptoms, diagnosis, and minimal reproduction.

4. Results

Unless stated otherwise, the results of Sections 4.1 and 4.2 are those of the 1.0 m baseline grid. Section 4.3 gives the corresponding 0.5-m-grid values and the sensitivity of each conclusion to refinement.

4.1 Failure sequence and pressures

Table 4 lists the simulated sequence, Fig. 2 shows representative frames, and Fig. 3 shows the overpressure histories. The ignition kernel generates a local +43 kPa transient that fails the roof tile and slab tile directly above and below the kernel and the nearest demising wall within 40 ms. As the flame expands at the prescribed S_T, the reference probes in the cloud column (cloud centre, adjacent mall corridor and first floor below the slab) reach instantaneous overpressures of 3.9–4.6 kPa, and the second-storey probes 50 m to either side 5.1–5.6 kPa (with a mean of 2.26 kPa at the cloud centre over 0.05–0.30 s), sufficient to strip the degraded ALC facade panel-by-panel along a 60 m contiguous run of the south facade centred on the cloud, and to cause ten roof tiles clustered above the cloud to fail. An external blast of ~+2.8 kPa reaches the dashcam side of the road. Burning completes in ~1.5 s, with 78 of 255 panels failing owing to direct pressure. The subsequent progressive collapse (Section 3.4) adds the storey-crush core (9 columns; ~12 crushed bays) and a 12-tile unzipping band, for a total collapse footprint of ~60 × 40 m.

Ignition-location sensitivity. After the model was frozen, the 3 August broadcast [15] placed the estimated epicentre in the back-of-house corridor along the south edge, rather than in the tenant-zone interior. To test whether the conclusions depend on the assumed ignition point, the baseline run was repeated as a matched pair on one machine, with every parameter unchanged except the kernel centre, which was moved 9 m south from (100.5, 36.5, 6.8) to (100.5, 27.5, 6.8) — inside the south tenant strip (the model contains no back-of-house corridor), at the position suggested by the drone imagery. The qualitative pattern is unchanged: damage remains concentrated on the south side, the roof is breached directly above the ignition point (8 roof tiles versus 10), a single slab tile fails, and the demising walls nearest the kernel fail within the first venting legs (the first tens of milliseconds). The south-first sequence becomes even more pronounced: the facade panel directly outside the kernel now fails at the first sampled instant (0.02 s, rather than 0.04 s). The aerial comparison in Fig. 2(g–l) shows the same signature: the south-edge run vents along the south facade from the first frames, and at 2–4 s its dust plume remains concentrated toward the south edge, with weaker northward spill than the central-ignition baseline. Totals shift downward—62 of 255 panels fail, versus 78—because with the kernel adjacent to the south wall, the earliest vent openings form immediately in the closest facade panels and bleed the interior sooner; north-facade failures drop from 15 to 12. The sustained overpressure at the tenant-zone reference probe rises from 2.2 to 3.1 kPa—with edge ignition, the burn reaches the cloud centre later, extending pressure production into the 0.05–0.30 s averaging window—but remains in the same few-kilopascal order. The second-storey stripping band on the south facade is identical in the two runs (a contiguous 60 m from x = 70 to 130 m), whereas ground-storey stripping shrinks, so the both-storey contiguous measure used in Table 5 drops from 60 m to 20 m. Under the no-tuning principle of Section 4.2, both results are reported as computed: the south-edge ignition matches the newly reported epicentre position and preserves the upper-storey stripping band and roof breach, while the baseline central-aisle ignition reproduces the observed both-storey stripping extent more closely. The headline conclusions—south-side concentration, a roof breach above the epicentre, and interior loads of a few kilopascals acting on seismically weakened members—do not hinge on the exact ignition location.

This paper defines contiguous facade stripping as the longest run of adjacent 10 m facade slots in which both facade layers (fs1 and fs2) have failed, corresponding to the “stripped to the frame” condition identified in the aerial record; the same definition applies to both grids. A single-layer criterion (fs2 only), used in an earlier draft of this work, gives 60 m on the baseline grid.

Table 4. Simulated failure sequence (1.0 m baseline grid).

Time	Simulated event
0.02 s	Kernel pressure pulse (+43 kPa locally): roof tile above and slab tile below ignition fail; nearest demising wall fails

Time	Simulated event
0.04–0.12 s	South facade ALC begins stripping (+2.8–6.9 kPa); 1F glazing; north facade follows; partition band destroyed
0.1–0.5 s	Main burn: 3.9–4.6 kPa peak at the cloud-column probes, 5.1–5.6 kPa at the 2F probes 50 m to either side; 40 m contiguous facade loss; 10 roof tiles; external blast at +2.8 kPa
0.35–0.7 s	Storey-crush core: 9 columns fail first; bays above pancake to rubble heights of 1–2.2 m
0.5–3 s	Peripheral slab unzipping (12 tiles with durations of 0.35–0.65 s per tile); the column-toppling rule did not activate
~1.5 s	Combustion complete; 78/255 panels fail under direct gas pressure

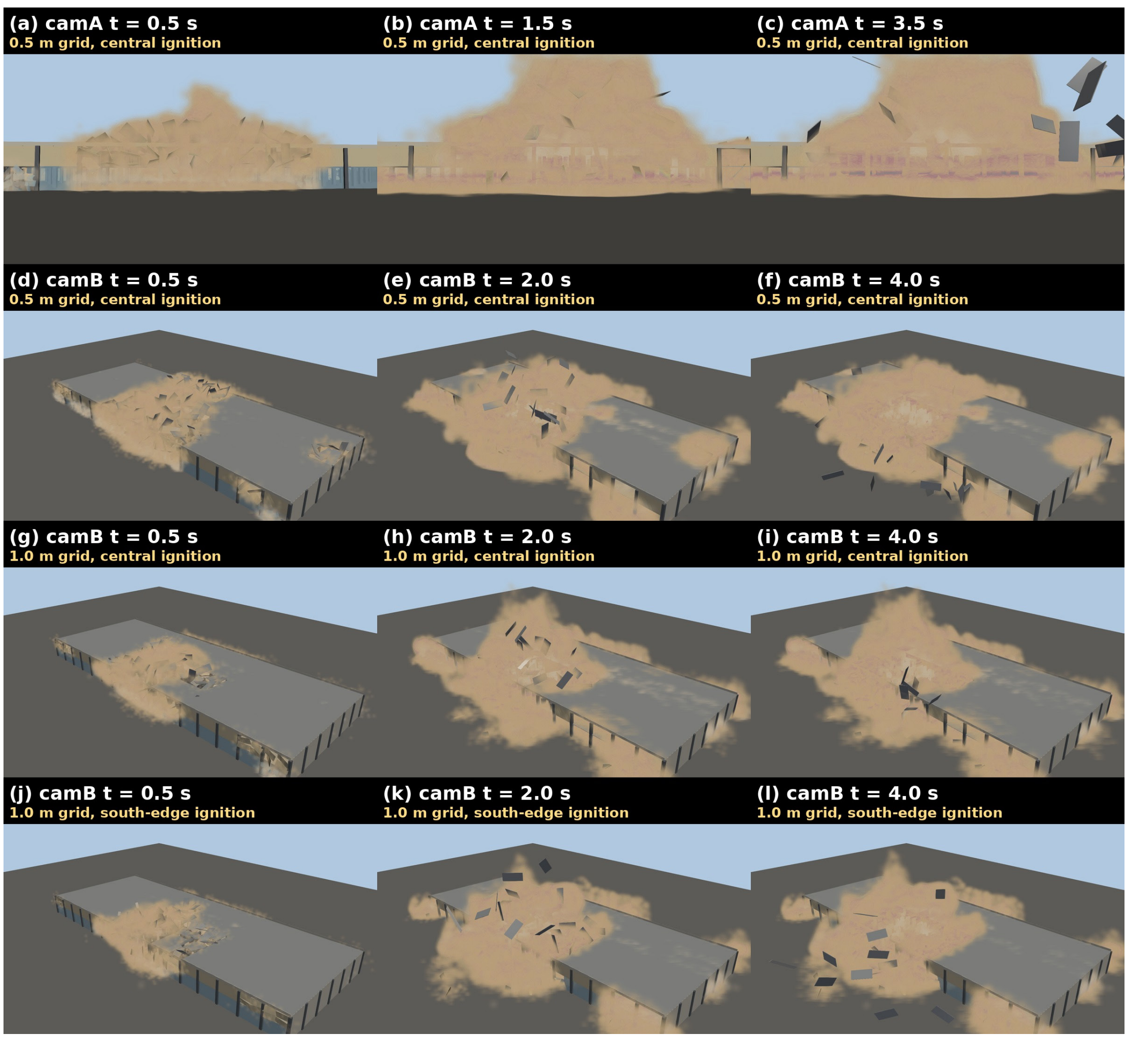

Figure 2. Simulation frames. (a–c) Road-level view corresponding to the dashcam angle at t = 0.5, 1.5, and 3.5 s, and (d–f) aerial view at t = 0.5, 2.0, and 4.0 s, both from the refined 0.5 m central-ignition run; (g–i) the 1.0 m central-ignition baseline and (j–l) an earlier computation of the 1.0 m south-edge-ignition sensitivity run of Section 4.1, in which one partition was opened spuriously at t = 0.02 s by an out-of-mesh pressure probe, aerial views at the same instants. Comparing (g–i) with (j–l) isolates the effect of ignition location on the same grid, apart from that spurious opening: the south-edge run vents along the south facade from the outset (j), whereas the central run ejects more material through the roof breach in the first half-second (g); at 2–4 s, the south-edge plume remains concentrated toward the south edge, with visibly weaker northward spill (k, l versus h, i), consistent with the failure counts in Section 4.1, while the roof breach above the ignition point and the crushed core develop similarly in both runs. Tan particles: structural dust; grey boxes: panel fragments. The crushed core with stacked floor/roof plates is discernible in (f) through the dust. The refined run resolves the vent jets and the debris release more finely, and its panel-failure counts are those in Table 6; those of the two 1.0 m runs are given in Table 4 and Section 4.1, the latter from a recomputation of the run shown in (j–l). All frames are renderings of the present simulation; no broadcast imagery is reproduced in this paper.

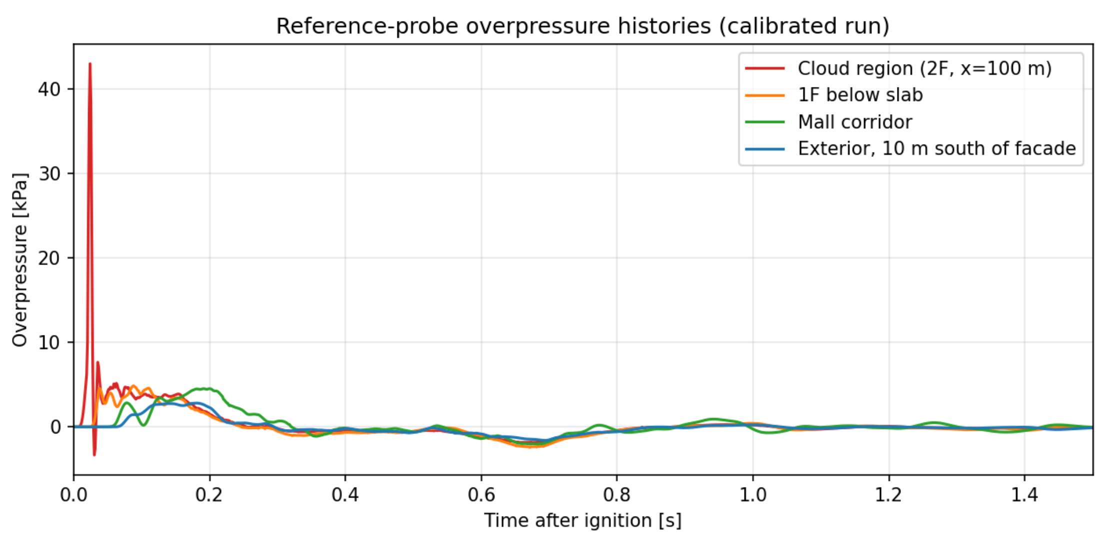

Figure 3. Overpressure histories at reference probes at the cloud centre, first floor (1F), corridor, and exterior.

4.2 Comparison with observations

Table 5 compares the footage-derived reconstruction targets of Table 1 with the simulation. The “role” column is essential to interpreting this table: the first two rows were used to select the burning velocity in Section 3.2 and therefore cannot be counted as validations, whereas the remaining rows were not used in any calibration.

Table 5. Reconstruction targets versus simulation, with the role of each quantity in the study. Ballistic quantities were never used in calibration.

Quantity	Observed	Simulated	Role
Contiguous facade stripping	~55 m (±15 m)	60 m (both grids)	**calibration target**
Near-frame-one envelope failure	~27 m wide eruption at 33 ms	envelope failure in time range 0.02–0.04 s over tens of metres	**calibration target**
Total facade damage span	~95 m	scattered failures to ~150 m (over-predicted fringe)	independent
Rear (tank-side) damage	radial sheet scatter	16–19 north-facade panels ✓	independent
Roof damage zone	local breach 10–15 m	tiles clustered above cloud (same zone, but over 60 × 45 m on the 1.0 m grid: extent about an order of magnitude larger than observed; count grid-sensitive)	independent
Collapse morphology	~7 bays; crushed core + draped plates	crush core ~20 × 20 m + unzipping band; footprint ~60 × 40 m	prescribed input — not independent
External fireball	no sustained fireball (brief vent flashes)	none sustained — transient vent flames are below grid resolution	independent
Initial dust-front speed	13 m/s (band 8–17)	8.9 m/s (1.0 m); 8.5 m/s (0.5 m)	independent — within band on both grids
Late front speed	4 m/s (3–7)	0.10 m/s (1.0 m), 0.01 m/s (0.5 m) over 4–6 s (under-predicted after CFD record ends)	independent

Quantity	Observed	Simulated	Role
Roof-debris launch speed and apex	30–40 m/s; apex ~24 m (band 15–33 m)	sheets (6.3 kg m^{-2}): v_z median 20.9 m/s (IQR 15.5–29.5, P90 43.6), apex median 21.8 m (IQR 18.6–32.6), n = 200; tiles (50 kg m^{-2}): 9.4 m/s, apex 18.1 m, n = 60	independent — apex ✓, launch speed partially
Debris throw	steel sheeting to ~100 m (lower bound)	sheets P95 8.9 m, max 20.4 m (n = 200)	independent **— not reproduced** (see text)

The last two rows constitute the most demanding independent test in this study, and the outcome is mixed and instructive. The observed launch speed and apex height are not independent observables: for thin steel skin of areal density ~6.3 kg m^{-2}, quadratic drag limits the still-air rise to 11–13 m for launch speeds of 30–40 m/s, so the observed pair (30–40 m/s and apex ≈ 24 m above ground from an 11 m roof) jointly constrains the release impulse and the drag treatment (Fig. 7b). Applying the unchanged 25 ms impulse law of Section 3.4 to sheeting of that areal density at the recorded roof-tile failure pressures (Δp ≈ 4.1–10.8 kPa with median ≈ 4.8 kPa) yields apex heights with an interquartile range (IQR) of 18.6–32.6 m, which covers the observed band well, and launch speeds whose IQR is 15.5–29.5 m/s with P90 = 43.6 m/s, so that the observed 30–40 m/s band corresponds to roughly the 75^{th} to 90^{th} percentile of the simulated sheet population rather than to its centre. The agreement in apex is therefore stronger than the agreement in launch speed. The CFD vent jets lift part of the population above the still-air curves, and the heavier lumped tiles reach the same apex band at lower launch speeds, as the analytic curves in Fig. 7b make explicit.

The throw distance is **not** reproduced: simulated sheets land within 20 m of the building (P95 = 8.9 m) against an observed lower bound of ~100 m. Vertically launched roof debris cannot traverse such distances under quadratic drag. The long-range sheeting in the aerial record most plausibly originated as near-horizontally launched facade skin or roof-edge flashing carried by the lateral vent jets — a debris class that the ALC-typified facade panels of this model do not represent. This deficiency is deliberately reported as a finding rather than removed by adjustment so that the ballistic observables remain an independent test.

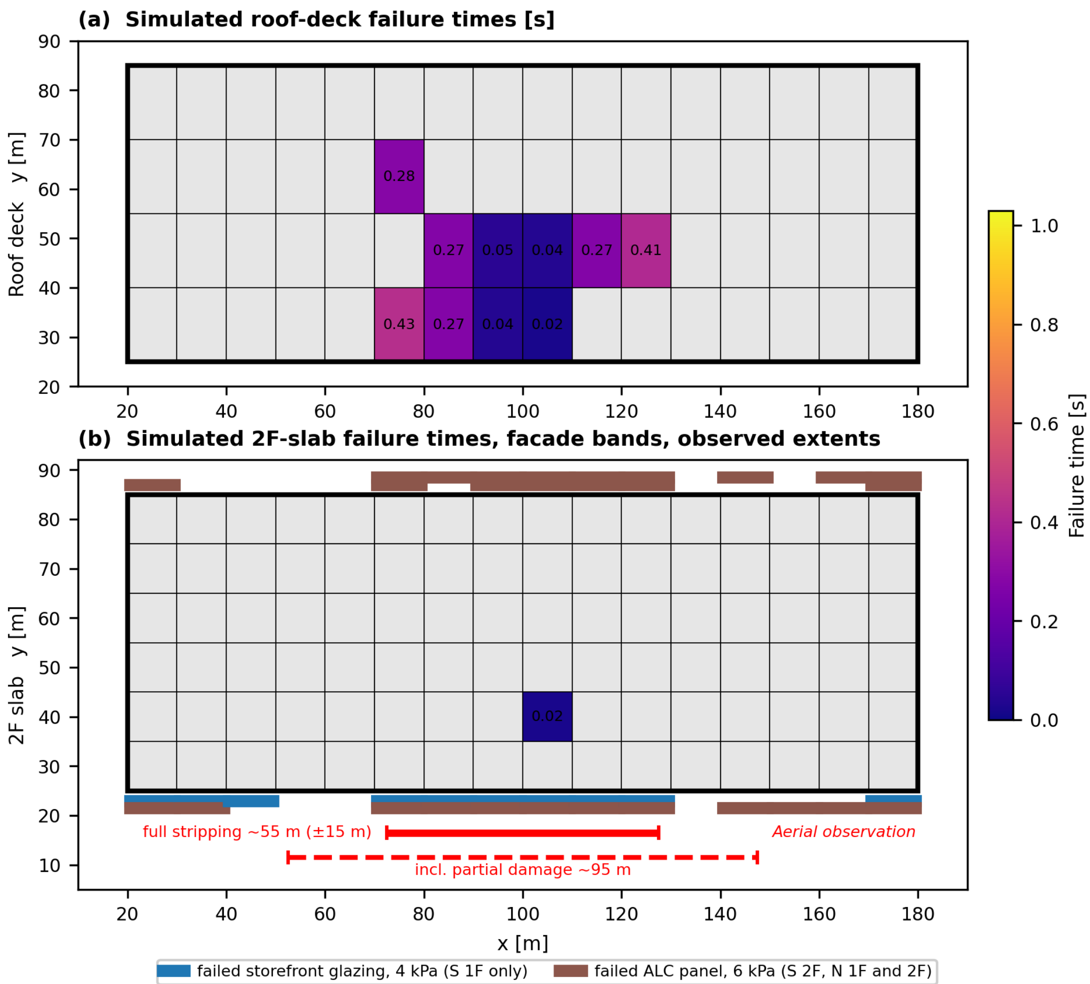

Figure 4. Simulated damage map (failure times) for the roof deck (top) and second-floor (2F) slab with facade bands (bottom), compared with the aerially observed facade-stripping extents (red lines). The extents are transcribed from public aerial footage, and no imagery is reproduced.

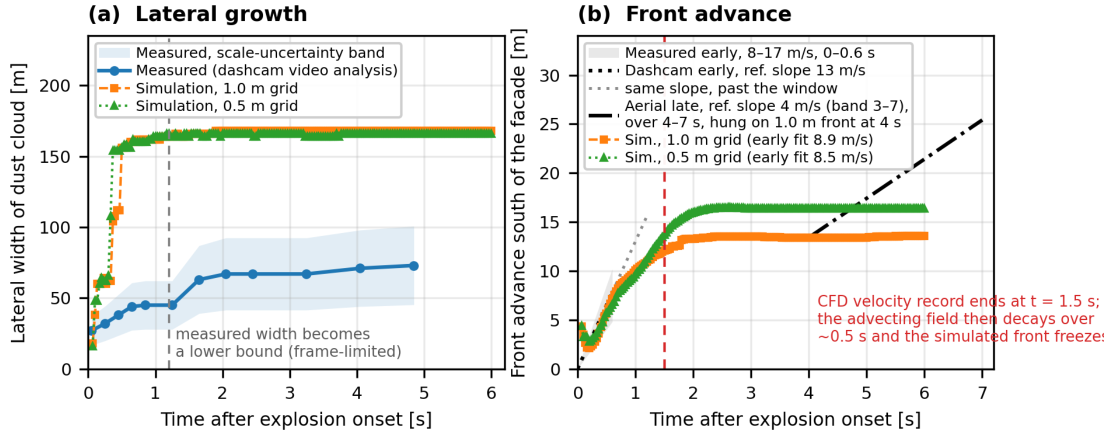

Figure 5. Dust-cloud kinematics versus dashcam measurements: lateral growth (left) and front advance (right), shown for both grids. The front speeds quoted in the text and in Tables 5 and 6 — 8.9 m/s on the 1.0 m grid and 8.5 m/s on the 0.5 m grid — are linear fits over 0.3–1.2 s. The measured value is 13 m/s with an 8–17 m/s scale-uncertainty band.

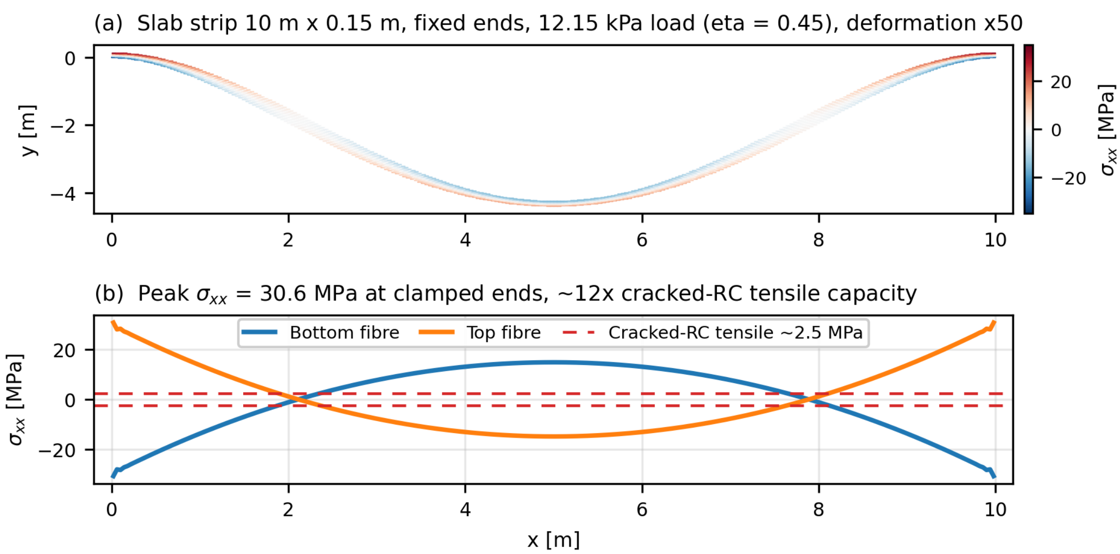

Figure 6. Independent solids4foam check of the slab failure criterion.

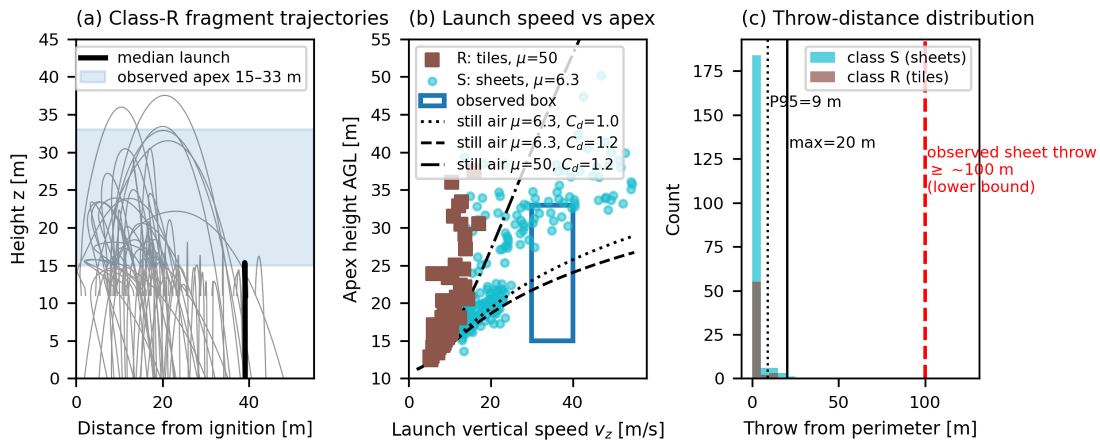

Figure 7. Independent test of the debris model against public-footage ballistics. (a) Simulated trajectories of roof-deck fragments, where the shaded band represents the observed apex with video-scale uncertainty. (b) Launch vertical speed versus apex height for both fragment classes, with the observed box and the analytic quadratic-drag relations for areal densities of 6.3 kg m^{-2} (two drag coefficients) and 50 kg m^{-2}; points above the still-air curves are lifted by the CFD vent jets. (c) Throw-distance distributions, where the observed ~100 m is a lower bound set by the debris identifiable in the footage and is not reproduced by roof-normal launches (see the main text). None of these quantities were used in any calibration.

4.3 Grid sensitivity

The reconstruction was repeated on a uniformly refined 0.5 m mesh (4.56 M fluid cells and the time step halved to 0.125 ms) with all other settings — the prescribed wrinkling factor, the panel capacities of Table 2, the degradation factor, the cloud definition, and the ignition kernel — kept identical. The propane inventory integrates to 325.5 and 331.0 kg on the two grids (a 1.7% lattice-snapping difference), and the deterministic mesh-rebuild machinery reproduces cell centres to within 8.5e−14 m on the fine grid. The 0.5 m spacing preserves the 10 m panel lattice exactly and resolves the 1 m merchandising obstacles with two cells rather than one.

Because the wrinkling factor is prescribed rather than emergent, the burning rate is by construction largely independent of the mesh. This test therefore isolates the discretization of the pressure field, the venting flow, and the geometry, which is precisely the part of the chain on which the damage-extent conclusions depend.

Table 6 and Fig. 8 summarize the comparison. Interpret the outcome as partial rather than complete convergence. Near-field and impulsive quantities behave as expected for single- or double-cell vent jets: the external blast spike at the road probe more than doubles, and the tiles that fail rise from 10 to 16 (roof) and from 1 to 8 (slab), carrying the total panels failing under direct pressure from 78 to 95, an increase of

22%. **The 1-m-mesh results are therefore not convergent for local peaks or for failure counts, and are not conservative with respect to the latter.** Two quantities that carry the conclusions of this paper, by contrast, are essentially unchanged: the contiguous facade-stripping length is 60 m on both grids, and the sustained tenant-strip overpressure changes by 9%. The initial dust-front speed changes by 5%, from 8.9 m/s to 8.5 m/s. Both values lie inside the observed 8–17 m/s band, so the cloud kinematics are reproduced at the order-of-magnitude level on both grids; two grids cannot establish convergence. Overall, the contiguous facade-stripping length and gross cloud kinematics survive halving the cell size, whereas absolute local pressures and failure counts do not. The latter are reported here rather than suppressed, and any use of the present overpressures as design quantities would be unwarranted.

Table 6. Grid-sensitivity metrics. This paper's conclusions rest on the facade-stripping length, the sustained compartment overpressure, and the dust-front speed (bolded rows).

Metric	1.0 m	0.5 m	Δ
Cells/maximum time step	0.57 M/0.25 ms	4.56 M/0.125 ms	—
Peak overpressure, ignition-kernel probe (not interpreted)	43.0 kPa	28.3 kPa	−34%
Sustained overpressure, tenant strip (0.05–0.30 s mean)	**2.26 kPa**	**2.47 kPa**	**+9%**
External blast peak, road probe	2.82 kPa	6.15 kPa	+118%
Panels failed under direct pressure (out of 255)	78	95	+22%
Contiguous facade stripping length	**60 m**	**60 m**	**0**
Roof tiles failed	10	16	+6 tiles
Slab tiles failed	1	8	+7 tiles
Cloud-centre decay below 1 kPa	0.218 s	0.198 s	−9%
Initial dust-front speed (0.3–1.2 s fit)	**8.9 m/s**	**8.5 m/s**	**−5%**
Maximum cell-centre displacement across rebuilds	< 1e−13 m	8.5e−14 m	—

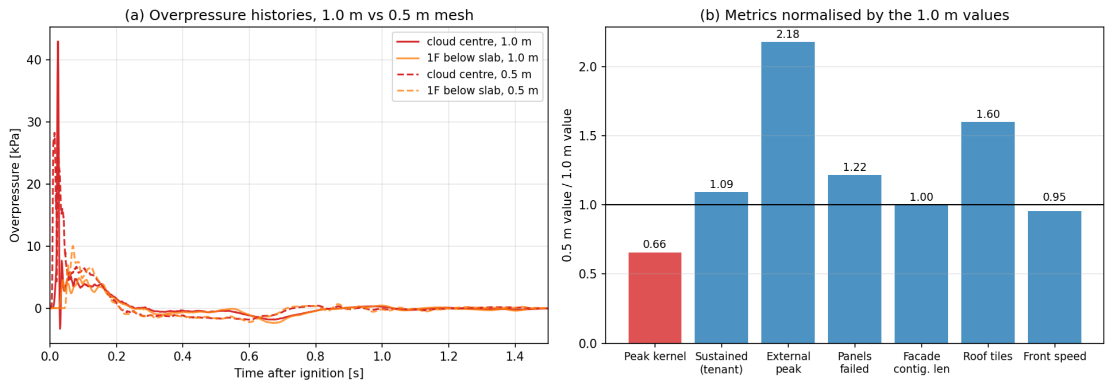

Figure 8. Grid sensitivity. (a) Overpressure histories at the cloud-centre and 1F probes for the 1.0 m (solid) and 0.5 m (dashed) meshes. (b) Integrated metrics normalized by their 1.0 m values. The external blast spike is strongly grid-dependent, as expected for single- or double-cell vent jets; the facade-stripping length and the sustained compartment overpressure are not. The near-kernel transient is reported without quantitative interpretation.

5. Discussion

Calibration and its meaning. With realistic momentum dynamics and venting, uncalibrated RANS burning (with initial cloud turbulence u′ ≈ 0.8–2 m/s) produced only ~1–2 kPa, which is sufficient for partitions and little else, but incompatible with the observed damage. The evidence was reproduced with congested-space burning velocities at the upper end of the reported range together with damage-class-III seismic degradation (0.45) of member capacities. Read inversely, the observed damage pattern itself constrains the event physics: either the effective burning velocity reached the top of the congested-space range (strong obstacle-driven acceleration, possibly aided by ceiling-plenum channelling that our geometry omits), or member capacities after intensity-7 shaking were substantially below intact values, or (most plausibly) both. This coupling of Natech pre-damage with explosion loading is, we believe, the principal engineering lesson: post-earthquake gas safety assessments should not assume intact structural resistance when judging acceptable leak and ignition risk. If the official investigation establishes that gas isolation did not occur as intended, the reliability of seismic shutoff provisions for bulk LP installations serving large public buildings warrants review. That determination, however, is beyond the scope of this study.

What the ballistic test does and does not establish. The debris comparison in Section 4.2 is the only place in this work where a prescribed parameter in Table 3 is confronted with an observable it did not help set. The apex heights agree. The launch-speed distribution is centred lower than the observed band, and the throw distance fails outright. The failure is informative: it localizes a missing debris class—near-horizontally launched thin skin—rather than casting doubt on the impulse law itself, and it identifies the model's most valuable single extension.

Overpressure in context. For scale context only, the town-gas explosion at Ronan Point [9] is commonly associated with internal overpressures of approximately 20 kPa. The 4–6 kPa peak overpressures at the interior reference probes in the present reconstruction, with a 43 kPa local kernel transient on the baseline grid, are of the same order of magnitude. The comparison should not be read as validation: the enclosure volume, venting area, fuel, and structural systems all differ substantially; no overpressure was measured in the present event; and the local peak itself is reported without quantitative interpretation (Section 4.3).

Deflagration or detonation? The combustion regime of this event has been debated publicly, and this work addresses that question asymmetrically. It does not test the detonation hypothesis: the solver is a deflagration model with a prescribed turbulent burning velocity and contains no deflagration-to-detonation-transition (DDT) mechanism, and the 1.0 m and 0.5 m meshes are orders of magnitude too coarse to resolve one. The reconstruction establishes sufficiency: a fast, congestion-driven deflagration at the upper end of the reported burning-velocity range, acting on seismically degraded members, reproduces the observed facade-stripping extent, the early dust-front kinematics, and the absence of a sustained fireball without invoking detonation-level loads. The brief flame luminosity at the cloud base in the close-range recording aired on 3 August [15], interpreted in Section 2.2 as transient vent-jet flames, is itself a classic signature of a vented deflagration (the external combustion of vented unburned mixture), and the expert assessment that rooftop units were lifted "by pressure from below" [15] is likewise consistent with sustained quasi-static internal pressure rather than with a propagating shock.

Conversely, a fully developed Chapman–Jouguet (CJ) detonation of the cloud can be excluded on damage-magnitude grounds. For near-stoichiometric propane–air, the CJ pressure is about 1.9 MPa, and reflected pressures are typically about 2.5 times higher [16]. Against the connection-governed capacities of Table 2 — 12.2–27 kPa for the composite slab tiles and 2.7–6 kPa for the ALC facade — the static excess is two orders of magnitude. The natural objection, that detonation loading is impulsive and should be judged on impulse rather than peak pressure, does not change the conclusion. Taking the post-Taylor-wave average of 0.5 MPa over a 10–30 ms positive phase yields a surface impulse of roughly 5–15 kPa·s. The corresponding figure obtained from the reflected peak and a triangular decay is three to five times higher, so the estimate used here is deliberately conservative. For floor plates with a fundamental period of order 0.1–0.3 s, the impulsive asymptote of the pressure–impulse curve corresponding to a 12–27 kPa static resistance is of order 1–3 kPa·s [17]; the detonation demand therefore exceeds the tabulated capacities by a factor of about 2 to more than an order of magnitude, depending on which end of each range is taken.

Nor can the building's exceptional venting reconcile a detonation with the observations: venting caps the pressure of a deflagration because the burn (of order 1 s) is slow compared with the response of the failing panels, whereas a detonation would traverse the 50 m cloud in approximately 28 ms, outrunning vent formation entirely, so the structure would receive near-CJ loading regardless of how well the envelope vented. Such loading would leave neither standing columns nor floor plates draped between them.

The debris ballistics of Section 4.2 provide a direct, observational bound on the load actually delivered. The sheet debris tracked in the footage (with areal density ~6.3 kg m^{-2}) was launched at 30–40 m/s, corresponding to a surface impulse of only ~0.2 kPa·s. The detonation-scale impulse estimated above would impart velocities of order 10^3 m/s to the same sheeting — a loading at which the skin would tear and fragment rather than accelerate coherently. That the aerial record instead shows intact sheets following gravity-consistent parabolas at a few tens of metres per second is, by itself, incompatible with detonation loading of the surfaces that released them. The same reasoning disposes of the argument that the frame-to-frame simultaneity of envelope failure implies supersonic propagation: as noted in Section 2.2, pressure equalizes on the acoustic timescale (11 m per video frame), so panels tens of metres apart can fail within a few frames under a subsonic deflagration — and the debris they shed in those same frames moves at deflagration, not detonation, speeds.

The computed damage was obtained with both the upper-bound burning velocity and the damage-class-III capacity reduction of 0.45. No intact-capacity run was performed; applying the intact thresholds to the computed pressure record removes the contiguous both-storey stripping in the weakened zone (60 m to none), although 43 of the 78 failures also reach their intact capacity. That, on the computed loads, the observed stripping appears only with members already weakened by the mainshock is not evidence about the combustion regime, since the loads were computed with a deflagration model; nor does a static comparison of this kind show that the weakening was necessary, because intact members would have changed the vent-opening history and hence the loads themselves. The capacity reduction therefore remains an assumption of this reconstruction. The vehicle damage reported in the public record (Table 1) is also informative: a single truck driving past the mall was wrecked [15], an outcome that an external blast of a few kPa accompanied by sheet and glass debris can readily cause, while multiple dashcams in the nearby queue recorded continuously — a record that a detonation-scale external blast at tens of metres would not have left intact.

Two caveats are necessary. First, the calibrated burning velocity sits marginally above the top of the 4–20 m/s range reported for highly congested spaces, so the inferred regime is a fast deflagration rather than a slow one, and regimes in which a DDT becomes conceivable are not far above it. Second, this work cannot exclude a transient, localized transition to detonation within a particularly congested sub-region that decays upon leaving it. Mixture detonability confines where such a transition could have occurred: the detonation cell size of propane–air is of order 5 cm near stoichiometry and grows steeply on the lean side [16], so the lean ceiling layer (ft = 0.030 with an equivalence ratio of approximately 0.48) lies outside the reported lean detonability limit of propane–air, which sits near an equivalence ratio of 0.7 depending on configuration, and is effectively non-detonable. The intermediate layer (0.74) is marginal, and only the near-stoichiometric floor layer within the congested racking is a plausible host. Addressing that possibility would require sub-centimetre resolution and detailed chemical analysis. The unmodelled ceiling plenum remains the most plausible run-up path to examine first, with the same detonability caveat applying to its leaner mixture.

Collapse morphology. The two observed morphologies are a crushed core that "lost its original form" and surrounding bays with plates draped between standing columns. These morphologies follow from the two implemented mechanisms (column-first crush versus slab-first unzipping), which were prescribed to represent them, and are consistent with the interpretation that near-field blast impulse on quake-damaged columns initiated a storey collapse. At the same time, the wider floor loss propagated as a gravity-driven cascade of connection unzipping. The timing parameters of both mechanisms are prescribed (Table 3), so the agreement concerns morphology rather than collapse kinetics.

Limitations. The limitations of this study are as follows. Geometry is typified from public information (without drawings); the ceiling plenum, HVAC, fire compartmentation, and real merchandising are absent; the cloud mass and extent are a hypothesis consistent with, but not proven by, public data; S_T is prescribed rather than emergent, and the facade-stripping extent used to select it is consequently not an independent check; the collapse-stage parameters of Table 3 are likewise prescribed, including the impulse duration tested in Section 4.2; slab "opening" in the CFD approximates collapse venting; the collapse dynamics are one-way (no back-coupling of collapsing plates on the gas); the model contains no near-horizontally launched thin-skin debris class and consequently under-predicts long-range throwing by an order of magnitude; absolute video scales carry ±35–40% uncertainty, and measured cloud widths beyond ~1.2 s are frame-limited lower bounds; and the mesh under-resolves local blast interactions near the ignition kernel and the vent jets, as quantified by the 0.5-m-mesh comparison of Section 4.3 (the external blast peak is at +118%, and failure counts are +22%). Convergence was assessed on two grids only, and no Richardson extrapolation or grid-convergence index is claimed. None of these affects the qualitative sequence, but quantitative local overpressures should be interpreted as order-of-magnitude estimates. The solver contains no DDT mechanism, and the question of a localized DDT is outside the scope of the present grids and chemistry (as discussed in the subsection "Deflagration or detonation?" in Section 5).

6. Conclusions

An open-source, three-stage FSI reconstruction of the 2026 Kumamoto mall explosion was calibrated on a single gas-dynamic parameter, and fixed using the observed facade-stripping extent and the near-frame-one envelope failure. With that calibration in place, the model independently reproduces the absence of a sustained external fireball, the radial debris scatter on the rear facade, the position of the roof breach, and a dust-front speed of 8.9 m/s against a measured 13 m/s within an 8–17 m/s band. This is done using a stratified propane cloud with a mass of approximately 330 kg, congested-space burning velocities ($S_T \approx 3.6$–20.9 m/s across the stratified layers), and damage-class-III seismically degraded member capacities.

The debris model was tested—not calibrated—against the ballistics recorded in the footage. It reproduces the observed apex band. It does not reproduce the longest observed throws, implying a missing class of near-horizontally launched thin skin, and this is reported as a finding rather than removed by adjustment.

A uniform mesh refinement to 0.5 m leaves the contiguous facade-stripping length unchanged and the sustained compartment overpressure within about 9%, while the external blast spike and local failure counts change by tens of percent and do not converge. The facade-stripping conclusions therefore survive refinement whereas the absolute local pressures do not, and the latter should not be used as design quantities.

Based on damage magnitude, impulse, and debris velocity, the reconstruction supports a fast, congestion-driven deflagration rather than a cloud detonation. The current grids and chemistry cannot rule out a transient, localized transition confined to the near-stoichiometric floor layer.

The staged panel-venting scheme (deterministic mesh-rebuild restarts) provides a practical route to pressure–failure feedback in explosion CFD without moving-mesh machinery, and the documented subsetMesh empty-patch pitfall is of immediate practical value to OpenFOAM users. The results are consistent with a Natech interpretation in which earthquake pre-damage of both the gas supply system and the structure contributed to the severity of the outcome. Whether that interpretation is correct, and what regulatory consequences follow, rest with the official investigations.

Ethics and disclaimer

This reconstruction is based exclusively on publicly available information (broadcast footage, operator and government releases) and typical engineering data. It does not use, nor did its authors have access to, any non-public investigation material. No broadcast or third-party imagery is reproduced in this paper. All figures are the authors' own renderings and plots, and quantities derived from footage are reported as measurements only. All quantitative findings are engineering hypotheses intended to support scientific discussion; official investigations (fire service, METI, police) determine the cause of the accident. No attribution of legal responsibility is made or implied. Casualty figures cited reflect public reporting as of 30 July 2026 and may be revised. The authors express their condolences to the victims and their families.

Use of AI tools

In accordance with the engrXiv AI policy, the author discloses the following use of AI/LLM tools in preparing this work: assistance with English language editing and document preparation of author-written text, assistance in identifying candidate publication venues, critical review of author-written drafts leading to revisions of framing and presentation, and assistance in specifying the design and reporting format of the independent verification tests presented in Sections 4.2 and 4.3. The author produced and checked all numerical models, the simulation set-up, the video measurements, the execution of all simulations and analyses, the results, the verification cases, and the scientific interpretation. No text was accepted verbatim without author review, AI tools generated no data or figures, and no AI system is listed as an author.

Data and code availability

The complete case-generation scripts, staged-venting driver, collapse and dust dynamics code, minimal reproductions of the subsetMesh pitfall, and rendered animations are available from the corresponding author. They are intended for deposit in a public repository alongside this preprint, together with the 0.5 m grid-sensitivity case, the per-fragment ballistic trajectories (fragments_roof.csv, fragments_sheet.csv, ballistics_raw.json), and the extraction and plotting scripts. Fragment and particle counts quoted in Section 3.4 refer to the archived production run.

Acknowledgements

The simulations were performed on a workstation at the Institute of Fluid Science, Tohoku University. Some numerical results in this research were obtained using supercomputing resources at the Cyberscience Center, Tohoku University. The numerical simulations were partially performed using FUJITSU AFI-NITY scalar parallel computers at the Institute of Fluid Science, Tohoku University. No external funding was received for this study. Mark R. Kurban from Edanz (https://www.jp.edanz.com/ac) edited a draft of this manuscript.

References

[1] Weller HG, Tabor G, Gosman AD, Fureby C. Application of a flame-wrinkling LES combustion model to a turbulent mixing layer. Proc Combust Inst. 1998;27:899–907. DOI: 10.1016/S0082-0784(98)80487-6

[2] Gülder ÖL. Correlations of laminar combustion data for alternative S.I. engine fuels. SAE Technical Paper 841000; 1984. DOI: 10.4271/841000

[3] Menter FR. Two-equation eddy-viscosity turbulence models for engineering applications. AIAA J. 1994;32(8):1598–1605. DOI: 10.2514/3.12149

[4] OpenCFD Ltd. OpenFOAM-v2606 User Guide; 2026. Available at https://www.openfoam.com; User Guide (PDF): https://dl.openfoam.com/source/v2606/UserGuide.pdf

[5] Harris RJ. The Investigation and Control of Gas Explosions in Buildings and Heating Plant. E & FN Spon; 1983. Record: Google Books

[6] Bauwens CR, Chaffee J, Dorofeev S. Effect of ignition location, vent size, and obstacles on vented explosion overpressures in propane–air mixtures. Combust Sci Technol. 2010;182(11–12):1915–1932. DOI: 10.1080/00102202.2010.497415

[7] NFPA 68. Standard on Explosion Protection by Deflagration Venting. National Fire Protection Association; 2018. Record: NFPA LiNK (2018 edition)

[8] EN 1991-1-7:2006. Eurocode 1: Actions on structures — Part 1-7: General actions — Accidental actions. CEN; 2006. Record: European Commission JRC

[9] Griffiths H, Pugsley A, Saunders O. Report of the Inquiry into the Collapse of Flats at Ronan Point, Canning Town. HMSO; 1968. Record: Google Books

[10] Architectural Institute of Japan. Design Recommendations for Composite Constructions. AIJ; 2010 (in Japanese). Record: NDL Search

[11] Japan Building Disaster Prevention Association. Standard for Damage Level Classification and Guidelines for Restoration Techniques of Earthquake-Damaged Buildings. JBDPA; 2015 (in Japanese). Record: JBDPA publication page

[12] Cardiff P, Karač A, De Jaeger P, Jasak H, Nagy J, Ivanković A, Tuković Ž. An open-source finite volume toolbox for solid mechanics and fluid–solid interaction simulations. arXiv:1808.10736; 2018. DOI: 10.48550/arXiv.1808.10736

[13] Hansen OR, Hinze P, Engel D, Davis S. Using computational fluid dynamics (CFD) for blast wave predictions. J Loss Prev Process Ind. 2010;23(6):885–906. DOI: 10.1016/j.jlp.2010.07.005

[14] Japan Meteorological Agency. Information on the earthquake of 28 July 2026 in Kumamoto Prefecture (Mj 7.1, maximum seismic intensity 7). JMA press releases; July 2026 (in Japanese). Record: JMA event record

[15] Operator and broadcast sources on the 28 July 2026 mall explosion: AEON Co., Ltd. press release (29 July 2026); NHK, NTV/news zero, FNN and KAB broadcasts (28–30 July 2026), including dashcam and aerial footage; dashcam footage provided to the Mainichi Shimbun (published 30 July 2026); and the NTV “Shinso Hodo Bankisha!” documentary (aired 3 August 2026), including drone surveys of the site, employee interviews, and a close-range recording of the explosion (in Japanese). Records: AEON news archive.

[16] Kaneshige M, Shepherd JE. Detonation database. GALCIT Report FM97-8, California Institute of Technology; 1997 (with updates). Record: Caltech Detonation Database

[17] US Department of Defense. Structures to Resist the Effects of Accidental Explosions. UFC 3-340-02; 2008. Record: WBDG Unified Facilities Criteria

Appendix. Changes from version 1

Each row gives the location in version 1, the words in version 1, the words in version 2, and the reason. “Corrected run” means that the value is taken from the runs repeated with the probe guard described in the notice of correction; “error” means that the version 1 value was wrong independently of that defect; “withdrawn” marks a claim that is no longer made; “consequential” marks a change forced by another change. Figures 3 and 8, and rows (a)–(i) of Figure 2, keep the layout of version 1 and are redrawn from the corrected runs; row (j)–(l) of Figure 2 still shows the earlier south-edge computation (see its caption); Corrected figures replace figures 4 to 7; Figure 1 is unchanged. Where version 1 typesetting split a table row across a page break (Table 4) or repeated a table header inside the body of a table (Tables 1 and 5), the table is shown as one; this is not a content change and is not listed.

Location	Version 1	Version 2	Reason
Title block	Preprint, 5 August 2026	Preprint, version 2, 23 September 2026 (supersedes version 1 of 5 August 2026)	Version line.
Abstract, para. 1	capacities derived from Japanese design references	capacities derived from Japanese design references or typical details	Consequential. Section 3.3 now states that the glazing, ALC rocking-clip, and LGS partition capacities are the author's own derivations from typical details, not from design references.

Location	Version 1	Version 2	Reason
Abstract, para. 1	the localized roof breach above the ignition zone	the position of the roof breach above the ignition zone	Withdrawn. The breach position agrees with the observation, but its extent is over-predicted by about an order of magnitude (60 x 45 m simulated against 10-15 m observed).
Abstract, para. 1	the mixed crush/draped collapse morphology,	(deleted)	Withdrawn. The collapse mechanisms, their region and their timing were prescribed to represent the observed morphology, so the model does not reproduce it independently.
Abstract, para. 1	8.8 m/s	8.9 m/s	Corrected run. The early dust-front speed recomputed from the guarded 1.0 m run is 8.86 m/s, which still lies inside the observed 8-17 m/s band.
Abstract, para. 1	it reproduces the observed launch-speed/apex pair	it reproduces the observed apex band	Withdrawn. Only the withdrawn impulse inversion supported agreement with both values of the pair; the sheet population matches the apex band, but the observed 30-40 m/s launch speed lies at only its 75th-90th percentile.
Abstract, para. 1	, and the back-calculated impulse duration of 22 ms agrees with the 25 ms assumed a priori	(deleted)	Withdrawn. The 22 ms figure used a roof-tile failure-pressure median of about 10 kPa that matches no run; the true median of 4.8 kPa gives about 46 ms, so the two durations do not agree.
Abstract, para. 1	the near-kernel peak, the external blast spike, and the local failure counts	the external blast spike and the local failure counts	Consequential. Adding the probe guard alone moved the near-kernel peak by 26% on the same grid, so the peak is not interpreted as grid behaviour.
Abstract, para. 1	moved to the back-of-house corridor position	Moved to the south-edge position	Error. The relocated kernel lies 2.5 m inside the south facade, in the south tenant strip; the model contains no back-of-house corridor.
1. Introduction, para. 3	that reproduces both storey-crushing and slab-draping morphologies	that represents both storey-crushing and slab-draping morphologies	Withdrawn. The collapse mechanisms, region, and timing were prescribed to represent the observed morphologies, so the model does not reproduce them.
2.2 Quantitative video analysis, para. 1	i.e. well under 0.1 s across the cloud length	i.e. about 0.15 s across the cloud length, short compared with a burn of order 1 s	Error. The cloud spans 75-125 m (Section 3.1), so sound takes about 50/340 = 0.15 s to cross it, not well under 0.1 s; the argument holds only because this is short compared with the roughly 1 s burn.

Location	Version 1	Version 2	Reason
2.2 Quantitative video analysis, para. 1	within a single frame	within a few frames	Error. At 340 m/s, sound travels only about 11 m in one 33 ms frame (about 80 ms, or two to three frames, across 27 m), so acoustic equalization cannot make panels tens of metres apart fail within a single frame.
3.1 Geometry and gas-accumulation scenario, para. 1	1 × 1 × 2–3 m merchandising obstacles	1 × 2 × 3 m merchandising obstacles	Error. The geometry used in every simulation includes 1 × 2 × 3 m merchandising obstacles, so version 1 described their size incorrectly. No result changes.
3.1 Geometry and gas-accumulation scenario, para. 1	approximately 0.97, 0.74, and 0.48	approximately 0.96, 0.74, and 0.48	Error. For a floor-layer mass fraction of 0.058 and a stoichiometric mass fraction of 0.060, the equivalence ratio is 0.96; version 1 rounded it incorrectly.
3.2 Deflagration CFD, para. 2	approximately 8 m/s in the lean ceiling layer	approximately 3.6 m/s in the lean ceiling layer	Error. The Gülder correlation with the case's own coefficients gives S_T = 3.6 m/s at ft = 0.030; version 1's 8 m/s was too high by a factor of 2.2.
3.2 Deflagration CFD, para. 2	approximately 20 m/s in the near-stoichiometric floor layer	approximately 20.9 m/s in the near-stoichiometric floor layer	Error. The same evaluation gives 20.9 m/s at ft = 0.058, which is 4.5% above the 20 m/s upper bound of the quoted range.
3.2 Deflagration CFD, para. 2	spanning the 4–20 m/s range	slightly wider than the 4–20 m/s range	Consequential. With the corrected values of 3.6 and 20.9 m/s, the model's burning velocities fall slightly below and slightly above the literature range, so they no longer lie within it.
3.3 Vent-panel FSI (quasi-two-way coupling), para. 1	the run was checkpointed, and the panel was opened.	the run was checkpointed, and the panel was opened; a panel whose probe the solver could not locate (returning the sentinel value −1e300) was excluded from this test and remained shut.	Corrected run. All version-2 values come from runs with a probe guard, in which a panel whose probe lies outside the mesh is no longer opened spuriously by the sentinel value but is kept shut.
3.3 Vent-panel FSI (quasi-two-way coupling), para. 1	Capacities followed Japanese design references (glazing, ALC rocking-clip connections, LGS gypsum partitions, composite-deck slabs, AIJ composite-structure studs, and puddle-weld capacities [10]).	Capacities followed Japanese design references (composite-deck slabs, AIJ composite-structure studs, and puddle-weld capacities [10]); the glazing, ALC rocking-clip connection and LGS gypsum partition capacities were the author's own derivations from typical details.	Error. Reference [10] covers only the composite slab, stud and puddle-weld capacities; the glazing, ALC and partition capacities are the author's own derivations, and the capacity values are unchanged.

Location	Version 1	Version 2	Reason
Table 2 caption	Derivations are from JIS/AIJ-based section analysis.	Derivations are from the author's section analysis, AIJ-based for the composite slab and its connections.	Withdrawn. The caption said every panel capacity came from JIS/AIJ, but only the composite slab and its stud and puddle-weld connections follow AIJ; the other classes are the author's own derivations, and no JIS source is used.
Table 2, row "Gypsum/LGS partitions", column "Panel class"	Gypsum/LGS partitions	Gypsum/LGS partitions (demising walls 30 × 3 m)	Error. The three tenant demising walls in the partition class are 30 × 3 m (y = 25–55 m, z = 6–9 m) in every run; the other partitions are 10 × 3 m as stated.
3.4 Progressive collapse and debris/dust dynamics, para. 1	(capped at 55 m/s)	(capped at 55 m/s, then scaled by a random factor of 0.7–1.3 and perturbed by Gaussian scatter; roof tiles received an additional upward kick of 3–10 m/s, and slab fragments instead dropped at 0.5–2 m/s)	Error. Version 1 described the launch law as deterministic, but the fragment code (destruction_dynamics.py) always applied a random factor and Gaussian scatter after the cap, gave roof tiles an extra upward kick of 3–10 m/s, and released slab fragments at 0.5–2 m/s downward instead of using the impulse law.
3.4 Progressive collapse and debris/dust dynamics, para. 2	caused the columns to fail first	was assumed to cause the columns to fail first	Withdrawn. The crush-core column failures, their region and their timing are prescribed inputs; no column response was computed.
3.4 Progressive collapse and debris/dust dynamics, para. 2	Columns that lost three or more adjacent slab tiles subsequently toppled.	Columns that lost three or more adjacent slab tiles within a prescribed plan region were set to topple; this rule did not activate in any reported run.	Error. The toppling rule applies only within a prescribed plan region, and no column toppled under it in any run.
3.4 Progressive collapse and debris/dust dynamics, para. 2	Ground impacts of heavy fragments released dust bursts. These bursts consisted of 260,000 particles of dust advected	Failed panels emitted dust for 0.8 s after failure, and ground impacts of heavy fragments released further dust bursts. The dust, capped at 260,000 particles, was advected	Error. 260,000 is the particle-pool cap, not a count, and most of the dust comes from continuous emission by failed panels rather than from impact bursts.
3.5 Verification and a numerical pitfall, para. 4	(16.2 kPa)	(12.15 kPa)	Error. All runs used η = 0.45, which gives a degraded failure load of 12.15 kPa; 16.2 kPa came from a reduction factor that was later replaced.
3.5 Verification and a numerical pitfall, para. 4	peak bending stress of 34.7 MPa	peak bending stress of 30.6 MPa	Error. At 12.15 kPa, the peak fibre (bending) stress σxx is 30.6 MPa; the 34.7 MPa in v1 was a von Mises value at the old load.
3.5 Verification and a numerical pitfall, para. 4	approximately 14 times	approximately 12 times	Error. 30.6 MPa is about 12 times the 2.5 MPa cracked-RC tensile strength.

Location	Version 1	Version 2	Reason
3.5 Verification and a numerical pitfall, para. 4	confirming that the tabulated connection-governed capacities, not gross-section strength, control	consistent with the tabulated connection-governed capacities, not gross-section strength, controlling	Withdrawn. A linear-elastic, unreinforced strip calculation cannot show which mechanism governs, so it is consistent with connection-governed capacities but does not confirm them.
4.1 Failure sequence and pressures, para. 1	+34 kPa	+43 kPa	Corrected run. The kernel peak in the adopted guarded 1.0 m run is 42.97 kPa; 34 kPa came from the unguarded run. The value is reported without quantitative interpretation.
4.1 Failure sequence and pressures, para. 1	the nearest partitions	the nearest demising wall	Corrected run. In the guarded run, the only partition-type member that fails within 40 ms is demising wall dw_100; the plural came from a spurious partition opening in the unguarded run.
4.1 Failure sequence and pressures, para. 1	the tenant strip sustains instantaneous overpressures of 2.7–4.6 kPa (with a mean of 2.24 kPa over 0.05–0.30 s)	The reference probes in the cloud column (cloud centre, adjacent mall corridor and first floor below the slab) reach instantaneous overpressures of 3.9–4.6 kPa, and the second-storey probes 50 m to either side 5.1–5.6 kPa (with a mean of 2.26 kPa at the cloud centre over 0.05–0.30 s)	Error. No probe, panel set or averaging window of the computed record gives version 1's 2.7–4.6 kPa. Over the main burn (0.1–0.5 s) the three reference probes in the cloud column (cloud centre, mall corridor, first floor below the slab) peak at 3.9, 4.5 and 4.6 kPa, and the second-storey probes 50 m east and west at 5.6 and 5.1 kPa; the statement now names the probes it refers to.
4.1 Failure sequence and pressures, para. 1	mean of 2.24 kPa	mean of 2.26 kPa	Corrected run. The 0.05–0.30 s mean at the cloud-centre probe in the adopted guarded run is 2.26 kPa; 2.24 kPa came from the unguarded run.
4.1 Failure sequence and pressures, para. 1	~+2.9 kPa	~+2.8 kPa	Corrected run. The external blast peak at the road probe in the adopted guarded run is 2.82 kPa; 2.87 kPa came from the unguarded run.
4.1 Failure sequence and pressures, para. 1	79 of 255	78 of 255	Corrected run. The count of 79 included one partition opened spuriously by an out-of-mesh probe; the guarded run fails 78 panels.
4.1 Failure sequence and pressures, para. 1	a 48-tile unzipping band	a 12-tile unzipping band	Error. The unzipping band is 12 slab tiles, released as 48 quarter-tile fragments; version 1 reported the fragment count as a tile count.
4.1 Failure sequence and pressures, para. 1	~50 × 40 m	~60 × 40 m	Error. All dropped slab plates together span x = 70–130 m and y = 25–65 m, which is 60 × 40 m.

Location	Version 1	Version 2	Reason
4.1 Failure sequence and pressures, para. 2	the baseline run was repeated with every parameter unchanged	the baseline run was repeated as a matched pair on one machine, with every parameter unchanged	Corrected run. The corrected comparison uses a guarded matched pair of runs made on one machine, and its central run differs slightly from the adopted baseline (15 vs 16 north-facade panels).
4.1 Failure sequence and pressures, para. 2	inside the corridor strip	inside the south tenant strip (the model contains no back-of-house corridor)	Error. The relocated kernel at y = 27.5 m lies in the model's south tenant strip, and the model has no back-of-house corridor.
4.1 Failure sequence and pressures, para. 2	(11 roof tiles versus 10)	(8 roof tiles versus 10)	Corrected run. The guarded south-edge run fails 8 roof tiles against 10 for central ignition, so the sign of the difference reverses.
4.1 Failure sequence and pressures, para. 2	the corridor doors fail	the demising walls nearest the kernel fail	Consequential. The model has no corridor doors; in the guarded south-edge run, the demising walls nearest the kernel fail early, at 0.04-0.09 s.
4.1 Failure sequence and pressures, para. 2	60 of 255 panels fail, versus 79	62 of 255 panels fail, versus 78	Corrected run. The guarded matched pair gives 62 failed panels for south-edge ignition against 78 for central ignition.
4.1 Failure sequence and pressures, para. 2	drop from 16 to 9	drop from 15 to 12	Corrected run. The guarded matched pair gives 15 north-facade failures for central ignition and 12 for south-edge ignition.
4.1 Failure sequence and pressures, para. 2	drops from 60 m to 30 m	drops from 60 m to 20 m	Corrected run. The guarded south-edge run has 20 m of contiguous both-storey stripping.
4.1 Failure sequence and pressures, para. 2	the corridor run vents along the south facade	the south-edge run vents along the south facade	Consequential. The model has no corridor at the relocated kernel (100.5, 27.5, 6.8), which lies 2.5 m inside the south facade in the tenant strip, so the run is named after its position, as elsewhere in version 2.
4.1 Failure sequence and pressures, para. 2	the corridor ignition matches the newly reported epicentre position	the south-edge ignition matches the newly reported epicentre position	Consequential. The model has no corridor at the relocated kernel (100.5, 27.5, 6.8), which lies 2.5 m inside the south facade in the tenant strip, so the run is named after its position, as elsewhere in version 2.
4.1 Failure sequence and pressures, para. 3	gives 50 m on the baseline grid.	gives 60 m on the baseline grid.	Error. Recomputed from the failed-panel records, the fs2-only contiguous run on the 1.0 m baseline grid is 60 m (x = 70-130 m), the same as the both-layer measure; 50 m matches no run.

Location	Version 1	Version 2	Reason
4.1 Failure sequence and pressures, para. 3	Readers comparing with that earlier figure should note the change in definition rather than a change in result.	(deleted)	Consequential. With the single-layer value corrected to 60 m, the two definitions give the same length, so there is no 50/60 m difference left for this sentence to explain.
Table 4, row "0.02 s", column "Simulated event"	+34 kPa locally	+43 kPa locally	Corrected run. The 34 kPa kernel peak came from the unguarded run; the guarded 1.0 m run gives 42.97 kPa, which is reported without quantitative interpretation.
Table 4, row "0.04–0.12 s", column "Simulated event"	+3.5–6.9 kPa	+2.8–6.9 kPa	Error. The six south 2F ALC panels that fail between 0.04 and 0.12 s do so at 2.76–6.90 kPa, so the lower bound was wrong in v1's own run as well.
Table 4, row "0.1–0.5 s", column "Simulated event"	2.7–4.6 kPa sustained	3.9–4.6 kPa peak at the cloud-column probes, 5.1–5.6 kPa at the 2F probes 50 m to either side	Error. No probe, panel set, or averaging window of the computed record gives version 1's 2.7–4.6 kPa. Over the main burn (0.1–0.5 s), the three reference probes in the cloud column (cloud centre, mall corridor, first floor below the slab) peak at 3.9, 4.5 and 4.6 kPa, and the second-storey probes 50 m east and west at 5.6 and 5.1 kPa; the statement now names the probes it refers to.
Table 4, row "0.1–0.5 s", column "Simulated event"	60 m contiguous facade loss	40 m contiguous facade loss	Error. In the paper's both-layer contiguous-stripping measure, the run reaches 40 m by 0.5 s; the final 60 m is reached only at 0.65 s.
Table 4, row "0.1–0.5 s", column "Simulated event"	+2.9 kPa	+2.8 kPa	Corrected run. The external blast peak at the road probe is 2.82 kPa in the guarded run (2.87 kPa in the unguarded run).
Table 4, row "0.5–3 s", column "Simulated event"	48 tiles	12 tiles	Error. The unzipping band has 12 slab tiles; 48 is the number of quarter-tile fragments released.
Table 4, row "0.5–3 s", column "Simulated event"	columns losing all slab support topple; roof follows down	the column-toppling rule did not activate	Error. The support-loss toppling rule, and the roof follow-down that depends on it, fired in no run: the qualifying columns were either already removed by the crush core or outside the rule's region.
Table 4, row "~1.5 s", column "Simulated event"	79/255	78/255	Corrected run. The count of 79 included one partition opened spuriously by the out-of-mesh probe; the guarded run fails 78 of 255 panels.

Location	Version 1	Version 2	Reason
Figure 2 (image)	rows (a)–(i) rendered from the superseded runs	rows (a)–(i) re-rendered from the corrected runs; same layout, instants and labels; row (j)–(l) unchanged; the labels of row (j)–(l) read “south-edge ignition” instead of “corridor ignition”	Figure. Rows (a)–(i) showed runs that contained the probe defect.
Figure 2 caption	(j–l) the 1.0 m corridor-ignition sensitivity run of Section 4.1,	(j–l) an earlier computation of the 1.0 m corridor-ignition sensitivity run of Section 4.1, in which one partition was opened spuriously at t = 0.02 s by an out-of-mesh pressure probe,	Corrected run. Rows (a–i) are now drawn from the corrected runs, but no corrected rendering of the corridor-ignition run exists, so row (j–l) still shows the earlier computation containing the probe defect.
Figure 2 caption	isolates the effect of ignition location on the same grid	isolates the effect of ignition location on the same grid, apart from that spurious opening	Consequential. Row (j–l) of Figure 2 is the earlier run with one spurious opening, while (g–i) now shows the corrected baseline, so the two rows differ in that opening as well as in ignition location.
Figure 2 caption	given in Table 4 and Section 4.1.	given in Table 4 and Section 4.1, the latter from a recomputation of the run shown in (j–l).	Consequential. The Section 4.1 failure counts come from the corrected recomputation, not from the computation rendered in (j–l).
Figure 2 caption	the 1.0 m corridor-ignition sensitivity run	the 1.0 m south-edge-ignition sensitivity run	Consequential. The model has no corridor at the relocated kernel (100.5, 27.5, 6.8), which lies 2.5 m inside the south facade in the tenant strip, so the run is named after its position, as elsewhere in version 2.
Figure 2 caption	the corridor run vents along the south facade	the south-edge run vents along the south facade	Consequential. The model has no corridor at the relocated kernel (100.5, 27.5, 6.8), which lies 2.5 m inside the south facade in the tenant strip, so the run is named after its position, as elsewhere in version 2.
Figure 2 caption	the corridor plume remains	the south-edge plume remains	Consequential. The model has no corridor at the relocated kernel (100.5, 27.5, 6.8), which lies 2.5 m inside the south facade in the tenant strip, so the run is named after its position, as elsewhere in version 2.
Figure 3 (image)	probe histories of the superseded 1.0 m run (kernel peak about 34 kPa)	same plot from the corrected 1.0 m run (kernel peak about 43 kPa)	Figure. Superseded run.
Table 5, row "Rear (tank-side) damage", column "Simulated"	9–11	16–19	Error. The adopted central-ignition runs fail 16 (1.0 m) and 19 (0.5 m) north-facade panels. The value 9 belongs to the south-edge ignition run.

Location	Version 1	Version 2	Reason
Table 5, row "Roof damage zone", column "Simulated"	(same zone; count grid-sensitive)	(same zone, but over 60 × 45 m on the 1.0 m grid: extent about an order of magnitude larger than observed; count grid-sensitive)	Withdrawn. Only the roof-damage position agrees. The simulated extent is about an order of magnitude larger than the observed 10–15 m breach.
Table 5, row "Collapse morphology", column "Simulated"	footprint ~50 × 40 m ✓	footprint ~60 × 40 m	Error. All dropped slab plates span x = 70–130 m and y = 25–65 m, which is 60 × 40 m. The agreement tick is removed because the collapse mechanisms were prescribed.
Table 5, row "Collapse morphology", column "Role"	independent	prescribed input — not independent	Withdrawn. The collapse mechanisms, region, and timing were prescribed to represent the observed morphology, so this is not an independent reproduction.
Table 5, row "Initial dust-front speed", column "Simulated"	8.8 m/s (1.0 m); 7.6 m/s (0.5 m)	8.9 m/s (1.0 m); 8.5 m/s (0.5 m)	Corrected run. The early dust-front speeds were recomputed from the adopted guarded runs: 8.86 and 8.46 m/s, rounded to one decimal as in v1.
Table 5, row "Initial dust-front speed", column "Role"	within band on the baseline grid, marginally below on the refined grid	within band on both grids	Consequential. With the corrected 8.5 m/s, the refined-grid value lies inside the observed 8–17 m/s band.
Table 5, row "Late front speed", column "Simulated"	~2 m/s	0.10 m/s (1.0 m), 0.01 m/s (0.5 m) over 4–6 s	Error. The ~2 m/s was the 1.2–2.5 s speed. Over the 4–6 s window that overlaps the measurement, the simulated front is almost stationary.
Table 5, row "Roof-debris launch speed and apex", column "Simulated"	v_z median 21.0 m/s (IQR 15.6–29.5, P90 43.6), apex median 21.6 m (IQR 18.5–32.6)	v_z median 20.9 m/s (IQR 15.5–29.5, P90 43.6), apex median 21.8 m (IQR 18.6–32.6)	Corrected run. The sheet ballistics were regenerated from the failure record of the adopted guarded 1.0 m run.
Table 5, row "Roof-debris launch speed and apex", column "Simulated"	apex 18.0 m	apex 18.1 m	Corrected run. The roof-tile fragment apex median from the guarded 1.0 m run is 18.06 m.
Table 5, row "Debris throw", column "Simulated"	P95 8.7 m, max 20.2 m	P95 8.9 m, max 20.4 m	Corrected run. The sheet throw distances were regenerated from the adopted guarded 1.0 m run.
4.2 Comparison with observations, para. 2	Δp ≈ 5.6–24 kPa with median ≈ 10 kPa	Δp ≈ 4.1–10.8 kPa with median ≈ 4.8 kPa	Error. The quoted roof-tile failure pressures match no run; the ten roof tiles of the adopted 1.0 m run fail at 4.1–10.8 kPa, median 4.8 kPa.
4.2 Comparison with observations, para. 2	18.5–32.6 m	18.6–32.6 m	Corrected run. The sheet apex interquartile range recomputed from the guarded run is 18.6–32.6 m.
4.2 Comparison with observations, para. 2	15.6–29.5 m/s	15.5–29.5 m/s	Corrected run. The sheet launch-speed interquartile range recomputed from the guarded run is 15.5–29.5 m/s.

Location	Version 1	Version 2	Reason
4.2 Comparison with observations, para. 2	at proportionally lower launch speeds	at lower launch speeds	Consequential. The roof-tile launch speeds are set mainly by the added upward kick now disclosed in Section 3.4 (tile median 9.4 m/s), so they are not proportionally lower than the sheet speeds.
4.2 Comparison with observations, para. 3	A closed-form check makes the significance of this test explicit. The impulse law is inverted for the observed pair to obtain $\tau = v0\ \mu / \Delta p \approx 35 \times 6.3 / 10{,}000 \approx 22$ ms. Compared with the 25 ms assumed a priori in Table 3, this is a 12% agreement obtained from an observable that played no part in setting the parameter.	(paragraph deleted)	Withdrawn. The back-calculated 22 ms and its 12% agreement with the assumed 25 ms are withdrawn: the inversion used a median failure pressure (10 kPa) that no run produces, and with the recorded 4.8 kPa the same formula gives about 46 ms.
4.2 Comparison with observations, para. 4	P95 = 8.7 m	P95 = 8.9 m	Corrected run. The simulated sheet-throw P95 from the adopted guarded 1.0 m run is 8.9 m; version 1's 8.7 m came from the superseded unguarded run.
Figure 4 (image)	damage map with the "~95 m" reference line drawn 80 m long; north ground-storey bands labelled as glazing	corrected damage map: 95 m line drawn to scale; bands coloured by capacity class (glazing only on the south ground storey)	Figure. Drawing errors in the version 1 figure; the failure data shown are unchanged.
Figure 5 (image)	dust kinematics of the superseded runs (early fits 8.8 / 7.6 m/s; 0.5 m record ended at 4 s; late reference slope drawn outside its 4–7 s window)	dust kinematics of the corrected runs (early fits 8.9 / 8.5 m/s; both records to 6 s; late slope drawn over its window)	Figure. Superseded runs and a drawing error.
Figure 5 caption	8.8 m/s on the 1.0 m grid and 7.6 m/s on the 0.5 m grid	8.9 m/s on the 1.0 m grid and 8.5 m/s on the 0.5 m grid	Corrected run. The quoted early dust-front speeds came from the superseded runs, which lacked the probe guard; the guarded runs give 8.86 and 8.46 m/s over the same 0.3–1.2 s fit window, and these round to 8.9 and 8.5 m/s at version 1's one-decimal precision.
Figure 6 (image)	slab check at the superseded load of 16.2 kPa ("~14x")	slab check at 12.15 kPa (eta = 0.45): peak fibre stress 30.6 MPa, about 12 times the cracked-RC tensile strength	Figure. Corrected load.
Figure 7 (image)	ballistics of the superseded run; panel (a) showed the 100 m throw line as a maximum from the ignition point	ballistics of the corrected run; the 100 m observed throw is shown only in (c), as a lower bound from the building perimeter	Figure. Superseded run and a mislabelled observation.

Location	Version 1	Version 2	Reason
Figure 7 caption	with video-scale uncertainty, and the dashed line the observed maximum sheet-metal throw.	with video-scale uncertainty.	Figure. Figure 7 is redrawn from the adopted guarded run. Panel (a) no longer carries the ~100 m throw line, which v1 had plotted against distance from the ignition point and labelled a maximum, so the throw comparison now appears only in panel (c), as a lower bound.
4.3 Grid sensitivity, para. 3	the kernel peak decreases by 17%, the external blast spike	the external blast spike	Corrected run. The 17% decrease came from the unguarded runs, and the corrected kernel peak (43.0 to 28.3 kPa) is reported without quantitative interpretation, so it can no longer be listed as expected grid behaviour.
4.3 Grid sensitivity, para. 3	behave as expected for a marginally resolved ignition kernel and single- or double-cell vent jets	behave as expected for single- or double-cell vent jets	Consequential. With the kernel peak removed from the list, the reference to the ignition kernel would tie the tile counts to the kernel, and some of the extra tiles are far from it; this matches the Fig. 8 caption edit.
4.3 Grid sensitivity, para. 3	the tiles that fail because of kernel-driven jets rise	the tiles that fail rise	Consequential. In the corrected runs, four of the fifteen additional failed tiles lie far from the ignition kernel, so the rise cannot be put down to kernel-driven jets.
4.3 Grid sensitivity, para. 3	from 10 to 17 (roof) and from 1 to 6 (slab)	from 10 to 16 (roof) and from 1 to 8 (slab)	Corrected run. The 0.5 m tile counts came from the unguarded run; the corrected run gives 16 roof tiles and 8 slab tiles.
4.3 Grid sensitivity, para. 3	from 79 to 100, an increase of 27%	from 78 to 95, an increase of 22%	Corrected run. Both panel counts came from the unguarded runs, which include spurious openings; the corrected runs give 78 and 95 panels, an increase of 22%.
4.3 Grid sensitivity, para. 3	for near-kernel failure counts	for failure counts	Consequential. The failure counts that differ between the grids include tiles far from the kernel, so the non-convergence is not limited to near-kernel counts.
4.3 Grid sensitivity, para. 3	overpressure changes by 8%	overpressure changes by 9%	Corrected run. The corrected runs give a sustained overpressure of 2.26 kPa and 2.47 kPa, a change of 9%.
4.3 Grid sensitivity, para. 3	changes by 14%, from 8.8 m/s to 7.6 m/s	changes by 5%, from 8.9 m/s to 8.5 m/s	Corrected run. The corrected early dust-front speeds are 8.86 m/s and 8.46 m/s, a change of 4.5%; they are given at the same precision as version 1.

Location	Version 1	Version 2	Reason
4.3 Grid sensitivity, para. 3	The baseline value lies inside the observed 8–17 m/s band while the refined value falls marginally below its lower bound, so	Both values lie inside the observed 8–17 m/s band, so the cloud kinematics are reproduced at the order-of-magnitude level on both grids; two grids cannot establish convergence.	Consequential. Both corrected values lie inside the band. Version 1's "not themselves grid-convergent" rested on a 14% change and on the refined value falling below the band; neither holds, and two grids cannot establish convergence either way.
4.3 Grid sensitivity, para. 3	the extent of the damage and the gross cloud kinematics	the contiguous facade-stripping length and the gross cloud kinematics	Error. The roof-failure extent grows from 60 x 45 m to 110 x 60 m under refinement, so only the contiguous facade-stripping length (60 m on both grids) remains unchanged.
Table 6, row "Peak overpressure, ignition-kernel probe", column "1.0 m"	34.0 kPa	43.0 kPa	Corrected run. The 1.0 m kernel peak comes from the adopted guarded run (42.97 kPa), not the unguarded run (33.99 kPa).
Table 6, row "Peak overpressure, ignition-kernel probe", column "0.5 m"	28.4 kPa	28.3 kPa	Corrected run. The 0.5 m kernel peak comes from the adopted guarded run (28.26 kPa).
Table 6, row "Peak overpressure, ignition-kernel probe", column "Δ"	−17%	−34%	Corrected run. The guarded runs give a change of −34%, and because a guard-only re-run moved the kernel peak by 26% it is reported without quantitative interpretation.
Table 6, row "Sustained overpressure, tenant strip (0.05–0.30 s mean)", column "1.0 m"	2.24 kPa	2.26 kPa	Corrected run. This is the sustained overpressure of the guarded 1.0 m run.
Table 6, row "Sustained overpressure, tenant strip (0.05–0.30 s mean)", column "0.5 m"	2.43 kPa	2.47 kPa	Corrected run. This is the sustained overpressure of the guarded 0.5 m run.
Table 6, row "Sustained overpressure, tenant strip (0.05–0.30 s mean)", column "Δ"	+8%	+9%	Consequential. The change is recomputed from the corrected values (2.26 to 2.47 kPa).
Table 6, row "External blast peak, road probe", column "1.0 m"	2.87 kPa	2.82 kPa	Corrected run. This is the external blast peak of the guarded 1.0 m run.
Table 6, row "External blast peak, road probe", column "0.5 m"	6.16 kPa	6.15 kPa	Corrected run. This is the external blast peak of the guarded 0.5 m run.

Location	Version 1	Version 2	Reason
Table 6, row "External blast peak, road probe", column "Δ"	+114%	+118%	Consequential. The change is recomputed from the corrected values (2.82 to 6.15 kPa).
Table 6, row "Panels failed under direct pressure (out of 255)", column "1.0 m"	79	78	Corrected run. The v1 count included one partition that the out-of-mesh probe opened spuriously.
Table 6, row "Panels failed under direct pressure (out of 255)", column "0.5 m"	100	95	Corrected run. The guarded 0.5 m run fails 95 panels under direct pressure, and the v1 count included spurious openings.
Table 6, row "Panels failed under direct pressure (out of 255)", column "Δ"	+27%	+22%	Consequential. The change is recomputed from the corrected counts (78 to 95).
Table 6, row "Roof tiles failed", column "0.5 m"	17	16	Corrected run. The guarded 0.5 m run fails 16 roof tiles.
Table 6, row "Roof tiles failed", column "Δ"	+7 tiles	+6 tiles	Consequential. The difference is recomputed from 10 and 16 tiles.
Table 6, row "Slab tiles failed", column "0.5 m"	6	8	Corrected run. The guarded 0.5 m run fails 8 slab tiles.
Table 6, row "Slab tiles failed", column "Δ"	+5 tiles	+7 tiles	Consequential. The difference is recomputed from 1 and 8 tiles.
Table 6, row "Cloud-centre decay below 1 kPa", column "1.0 m"	0.213 s	0.218 s	Corrected run. This is the cloud-centre decay time of the guarded 1.0 m run.
Table 6, row "Cloud-centre decay below 1 kPa", column "0.5 m"	0.195 s	0.198 s	Corrected run. This is the cloud-centre decay time of the guarded 0.5 m run.
Table 6, row "Cloud-centre decay below 1 kPa", column "Δ"	−8%	−9%	Consequential. The change is recomputed from the corrected times (0.218 to 0.198 s).
Table 6, row "Initial dust-front speed (0.3–1.2 s fit)", column "1.0 m"	8.8 m/s	8.9 m/s	Corrected run. The guarded 1.0 m run gives 8.86 m/s, written to v1's one decimal.
Table 6, row "Initial dust-front speed (0.3–1.2 s fit)", column "0.5 m"	7.6 m/s	8.5 m/s	Corrected run. The guarded 0.5 m run, integrated to 6 s, gives 8.46 m/s, written to v1's one decimal.
Table 6, row "Initial dust-front speed (0.3–1.2 s fit)", column "Δ"	−14%	−5%	Consequential. The change is recomputed from the corrected speeds (−4.5%) and written as an integer percentage, like the rest of the column.

Location	Version 1	Version 2	Reason
Table 6, row "Peak overpressure, ignition-kernel probe", column "Metric"	Peak overpressure, ignition-kernel probe	Peak overpressure, ignition-kernel probe (not interpreted)	Consequential. Between the corrected and superseded 1.0 m runs, the kernel peak moved by 26%, so it is reported without quantitative interpretation; its grid change is not read as grid behaviour.
Figure 8 (image)	grid comparison of the superseded runs; the "Front speed" bar plotted the decay-time ratio	same seven-bar layout from the corrected runs; the "Front speed" bar now plots the dust-front speed ratio	Figure. Superseded runs and a plotting error.
Figure 8 caption	The near-kernel transient and the external blast spike are strongly grid-dependent, as expected for a 2.5 m ignition kernel and single- or double-cell vent jets;	The external blast spike is strongly grid-dependent, as expected for single- or double-cell vent jets; ... are not. The near-kernel transient is reported without quantitative interpretation.	Consequential. The kernel peak moved by 26% between two runs on the same 1.0 m grid, so its change between grids (-34%) can no longer be put down to grid resolution and is not interpreted.
5. Discussion, para. 1	Reproducing the evidence required congested-space burning velocities	The evidence was reproduced with congested-space burning velocities	Withdrawn. The capacity reduction was 0.45 in every run, and no intact-capacity run was performed, so the combination is shown to be sufficient but not required.
5. Discussion, para. 2	, and the impulse duration back-calculated from the observed pair (22 ms) agrees with the assumed value (25 ms) to 12%	(deleted)	Withdrawn. The 22 ms back-calculation used a median roof-tile failure pressure of about 10 kPa, which matches no run; with the correct median of 4.8 kPa, the same inversion gives 46 ms, so we withdraw the claimed 12% agreement.
5. Discussion, para. 3	with a 34 kPa local kernel transient	with a 43 kPa local kernel transient	Corrected run. 34 kPa was the kernel peak of the unguarded 1.0 m run; the adopted guarded 1.0 m run gives 42.97 kPa, written to whole kPa as in version 1.
5. Discussion, para. 3	The broad 3–5 kPa sustained	The 4–6 kPa peak overpressures at the interior reference probes	Error. Version 1’s 3–5 kPa was the rounded 2.7–4.6 kPa “sustained” range, which no computed record reproduces. The five interior reference probes peak at 3.9–5.6 kPa over the main burn, i.e. 4–6 kPa at this precision; “sustained” is kept for the 0.05–0.30 s mean.
5. Discussion, para. 3	Section 4.3 shows that the local peak itself is grid-dependent	the local peak itself is reported without quantitative interpretation (Section 4.3)	Consequential. The kernel peak moved 26% between the guarded and unguarded runs on the same grid, so it is no longer read as grid behaviour, and Section 4.3 and the Fig. 8 caption no longer say that it is.

Location	Version 1	Version 2	Reason
5. Discussion, para. 4	the observed damage extent	the observed facade-stripping extent	Withdrawn. The model over-predicts the roof-breach extent by about an order of magnitude, so only the facade-stripping extent is reproduced.
5. Discussion, para. 4	the mixed crush/draped collapse morphology,	(deleted)	Withdrawn. The collapse mechanisms, region and timing were prescribed to represent the observed morphology, so it matches by construction and is not evidence of sufficiency.
5. Discussion, para. 4	the dust-front kinematics	the early dust-front kinematics	Consequential. In the corrected runs only the early front speed (8.86 and 8.46 m/s) lies in the observed 8-17 m/s band; after about 2.5 s the front is stationary, while the observed late speed is 4 m/s.
5. Discussion, para. 5	by a factor of several to more than an order of magnitude	by a factor of about 2 to more than an order of magnitude	Error. The section's own ranges, 5–15 kPa·s of demand against 1–3 kPa·s of capacity, give a factor of about 1.7 to 15, so the low end is about 2, not several.
5. Discussion, para. 5	the response of the 255 failing panels	the response of the failing panels	Error. 255 is the total number of breakable panels, of which 78 fail on the baseline grid, so calling all 255 'failing panels' gives a wrong count.
5. Discussion, para. 6	— consistent with the ~22 ms deflagration impulse inferred independently in Section 4.2	(deleted)	Withdrawn. The 22 ms impulse duration was back-calculated from a wrong median failure pressure and is withdrawn as an independent inference; the ~0.2 kPa·s sheet impulse is unchanged.
5. Discussion, para. 6	can fail within one frame under a subsonic deflagration — and the debris they shed in that same frame moves	can fail within a few frames under a subsonic deflagration — and the debris they shed in those same frames moves	Error. At 11 m of sound travel per video frame, pressure needs about 80 ms (two to three frames) to cross the observed 27 m span, not one frame.
5. Discussion, para. 7	Reproducing the observed damage required both the upper-bound burning velocity and the damage-class-III capacity reduction of 0.45; with intact member capacities the same combustion produced markedly less damage.	The computed damage was obtained with both the upper-bound burning velocity and the damage-class-III capacity reduction of 0.45. No intact-capacity run was performed; applying the intact thresholds to the computed pressure record removes the contiguous both-storey stripping in the weakened zone (60 m to none), although 43 of the 78 failures also reach their intact capacity.	Withdrawn. No run with intact member capacities was ever performed, so the reported result is withdrawn and replaced by the static threshold comparison on the same computed pressures.

Location	Version 1	Version 2	Reason
5. Discussion, para. 7	Th at the observed damage is explicable only for members already weakened by the mainshock is itself evidence that the loading was not of detonation order.	That, on the computed loads, the observed stripping appears only with members already weakened by the mainshock is not evidence about the combustion regime, since the loads were computed with a deflagration model; nor does a static comparison of this kind show that the weakening was necessary	Withdrawn. The inference was circular: the loads come from a deflagration model, so they cannot discriminate the combustion regime; and a static threshold comparison against a fixed pressure history cannot show that the weakening was necessary, because intact members would have changed the vent openings and hence the loads.
5. Discussion, para. 7	The strongest structural indication is indirect. … A detonation would have required no capacity reduction at all.	(both sentences deleted)	Withdrawn. Both sentences introduced the capacity-reduction argument as evidence about the combustion regime, which it cannot be, because the loads were computed with a deflagration model.
5. Discussion, para. 7	— and it is precisely this coupling that gives the event its Natech character.	The capacity reduction therefore remains an assumption of this reconstruction.	Withdrawn. Without an intact-capacity simulation, the necessity of the weakening is not established; the paragraph now states that it is an assumption.
5. Discussion, para. 7	is consistent with the same picture:	is also informative:	Consequential. "The same picture" referred to the withdrawn inference that the loading was not of detonation order; the sentence now stands on its own evidence.
5. Discussion, para. 8	sits at the top of the 4–20 m/s range	sits marginally above the top of the 4–20 m/s range	Error. Recomputed with the case's Gülder coefficients, the floor-layer burning velocity is 20.9 m/s, which is 4.5% above the 20 m/s upper bound of the quoted range, not at its top (v1 used about 20 m/s).
5. Discussion, para. 9	These morphologies emerge naturally from the two implemented mechanisms (column-first crush versus slab-first unzipping), supporting the interpretation.	These morphologies follow from the two implemented mechanisms (column-first crush versus slab-first unzipping), which were prescribed to represent them, and are consistent with the interpretation.	Withdrawn. The collapse mechanisms, the crush-core region and their timing were prescribed to represent the observed morphologies, so the morphologies do not emerge independently from the model and can only be consistent with the collapse interpretation, not support it.
5. Discussion, para. 10	except for the impulse duration tested in Section 4.2	including the impulse duration tested in Section 4.2	Withdrawn. The 25 ms impulse duration is assumed a priori like the other Table 3 parameters, and it was only set apart by the withdrawn 22 ms back-calculation, so it is no longer an exception.

Location	Version 1	Version 2	Reason
5. Discussion, para. 10	(the local peaks are at −17% and +114%, and near-kernel failure counts are +27%)	(the external blast peak is at +118%, and failure counts are +22%)	Corrected run. The v1 deltas came from the unguarded runs; the guarded runs give +118% for the external blast, and +22% for the total panel-failure count, and the kernel peak is reported without quantitative interpretation, so it no longer quantifies under-resolution.
6. Conclusions, para. 1	the localized roof breach	the position of the roof breach	Withdrawn. The model reproduces where the roof breach is but not its size: it predicts about 60 × 45 m against an observed breach 10–15 m across, roughly an order of magnitude too large.
6. Conclusions, para. 1	, the mixed pancake/draped partial collapse of the second floor	(deleted)	Withdrawn. The collapse mechanisms, region and timing were prescribed to represent the observed morphology, so the model does not reproduce it independently.
6. Conclusions, para. 1	8.8 m/s	8.9 m/s	Corrected run. The guarded 1.0 m re-run gives an early dust-front speed of 8.86 m/s (0.3–1.2 s linear fit), which is 8.9 m/s to v1's one decimal and still inside the 8–17 m/s band.
6. Conclusions, para. 1	$S_T \approx 8–20$ m/s	$S_T \approx 3.6–20.9$ m/s	Error. With the case coefficients, the Gülder correlation gives 3.6 m/s in the lean ceiling layer and 20.9 m/s in the floor layer; v1 overstated the lean-layer value by a factor of 2.2.
6. Conclusions, para. 2	, and the back-calculated impulse duration of 22 ms agrees to within 12% with the 25 ms assumed a priori	(deleted)	Withdrawn. The 22 ms back-calculation used a failure-pressure median of about 10 kPa that matches no run; with the correct 4.8 kPa median it gives 46 ms, so the claimed 12% agreement is withdrawn.
6. Conclusions, para. 3	within 8%	within about 9%	Corrected run. The adopted guarded runs give a sustained overpressure of 2.26 kPa on the 1.0 m grid and 2.47 kPa on the 0.5 m grid (+9%; the exact figure is 9.2%), replacing the +8% from the unguarded runs.
6. Conclusions, para. 3	the near-kernel peak, the external blast spike, and the local failure counts	the external blast spike and the local failure counts	Consequential. A guard-only re-run on the same 1.0 m grid moved the kernel peak by 26%, so its change between the grids is reported without quantitative interpretation and is no longer classed as grid non-convergence.

Location	Version 1	Version 2	Reason
6. Conclusions, para. 3	The damage-extent conclusions	The facade-stripping conclusions	Consequential. On refinement, the roof-failure footprint grows from 60 x 45 m to 110 x 60 m, so only the facade-stripping length (60 m on both grids) survives refinement.