



GEM4D DOCUMENTATION

BEM Stress Analysis

Theory and Verification

PRODUCT	GEM4D
DOCUMENT	BEM Stress Analysis — Theory and Verification
VERSION	1.9.1.4
DATE	6 October 2026

Contents

1. The boundary element formulation	3
1.1 The indirect boundary element method	3
1.2 Fictitious-force (FF) elements	3
1.3 Kernels.....	3
1.4 Displacement-discontinuity (DD) elements	4
1.5 Combined FF and DD models	4
1.6 In-situ stress.....	4
1.7 What a solution holds.....	4
2. Field-point evaluation and classification	4
2.1 Evaluating a point	4
2.2 Classification.....	5
2.3 Interior values	5
2.4 The wall probe	5
3. How GEM4D prepares meshes	5
3.1 FF meshes	5
3.2 Watertightness	6
3.3 Mixed normals and seams	6
3.4 Body diagnostics.....	6
3.5 FF / DD classification	6
3.6 DD surfaces	7
3.7 Thin closed solids: the mid-surface.....	7
3.8 Glued and stepped blocks	8
3.9 Several surfaces and several steps	8
3.10 Refusals	9
4. Solver paths and memory	9
4.1 Direct and iterative solves	9
4.2 The H-matrix	9
4.3 The cache budget	9
4.4 The near-field threshold	10
4.5 What is logged	10
4.6 Energy evaluations	10
5. Energy formulations	10
5.1 Stored strain-energy density	10
5.2 Released energy on FF boundaries.....	10
5.3 Released energy on DD elements	11
5.4 Path independence	11
5.5 Area normalisations	11
6. Excess shear stress and Kirsch tangential stress	12
6.1 Plane-specific ESS	12
6.2 Coulomb ESS	12
6.3 Kirsch tangential stress.....	12
7. Accuracy and validation	12
7.1 Thin bodies and element size	12
7.2 Triangle shape	13
7.3 Coarsening a thin body	13
7.4 Gaps and steps between DD pieces.....	13
7.5 Released energy: the pressurised sphere	13
7.6 DD elements: the penny crack and curved reefs	13
8. Result changes from BEM library updates	14
8.1 BasRock.BEM 1.2.1.0 (October 2026).....	14
8.2 BasRock.BEM 1.2.0.0 (October 2026).....	15

Appendix A — Worked example: checking BEM_Tang_MPa by hand	15
A.1 The example row	15
A.2 Step 1 — the tunnel axis	16
A.3 Step 2 — the two cross-section axes	16
A.4 Step 3 — project the stress tensor onto the cross-section	17
A.5 Step 4 — the in-plane stress extremes	17
A.6 Step 5 — the Kirsch tangential stress	17
A.7 The same thing in a spreadsheet	17
A.8 Notes on the fine print	18
Appendix B — Worked example: checking a released-energy row by hand	18
B.1 The example	18
B.2 Step 1 — the traction that acted on the face	19
B.3 Step 2 — how far the wall moved	19
B.4 Step 3 — the release density	19
B.5 Step 4 — the element's contribution	19
B.6 The check	20
B.7 Checking a whole column in a spreadsheet	20
Appendix C — Worked example: checking a tabular panel by hand	20
C.1 The case	20
C.2 Step 1 — the traction that acted on the reef before mining	21
C.3 Step 2 — how far the reef closed	21
C.4 Step 3 — the release density	21
C.5 Step 4 — the panel's contribution, and the reef total	21
C.6 The check	22
C.7 On a dipping reef	22

This document is the companion to the *GEM4D — BEM Stress Analysis: User Guide*. The User Guide tells you how to build a model and read its results. This document explains how the module works inside, how it was checked against known solutions, and how three of its results can be reproduced by hand. It is written for engineers who want to know what is behind a number before they rely on it.

Scope and engineering responsibility. The same scope applies as in the User Guide: the BEM module is a linear-elastic tool, and its results are only as good as the geometry, rock properties and in-situ stress supplied. The accuracy figures in this document were measured on test geometries; they show how the method behaves, not the accuracy of any particular site model.

1. The boundary element formulation

GEM4D's stress solver is the BasRock BEM library. This chapter describes the formulation it implements. The library has been audited independently; its kernels are validated against the Lamé sphere, the Kirsch cylinder, the Boussinesq half-space, Sneddon's penny-shaped crack and plane-strain crack solutions.

1.1 The indirect boundary element method

The rock is a linear-elastic, homogeneous, isotropic solid with Young's modulus E and Poisson's ratio ν . The excavations are holes in it. Only their surfaces are discretised: each triangle of a mesh is one **element**. The rock around them is represented exactly by an elasticity solution — the **kernel** — so there is no outer boundary and no volume mesh.

GEM4D uses the **indirect** method. Each element carries an unknown, fictitious source of constant strength. The solver chooses the source strengths so that the boundary conditions hold at one point on each element, its **collocation point**. Once the strengths are known, the stress and displacement anywhere in the rock are sums of each element's analytic influence (Chapter 2).

Every reported stress is the **total** stress: the in-situ stress plus the induced change. The boundary condition on a mined surface is that the total traction on it is zero.

1.2 Fictitious-force (FF) elements

An FF element carries a constant **force density** — three unknowns per triangle. The solver finds the force densities that make every excavation surface free of traction, which is exactly what an open void does to the stress field.

The force density sits on a sheet, so the stress jumps across the element: on the element itself, the traction is the average of its two sides, and the half-jump appears on the diagonal of the system. This is why the energy maps evaluate just inside the rock (§2.4).

1.3 Kernels

Kernel	Domain	Notes
Kelvin	Infinite, full space	Elements close to the evaluation point, and an element's influence on itself, are integrated in closed form. Exact up to the wall.
Mindlin	Half-space with a traction-free plane surface at the ground surface Z	The Kelvin part is integrated in closed form; the free-surface image term by adaptive quadrature. Points above the surface are outside the model.

The solver works in **internal coordinates**, with $Z = 0$ at the ground surface and depth $= -Z$. GEM4D converts to and from mine coordinates at every exchange.

1.4 Displacement-discontinuity (DD) elements

A DD element is a surface across which the displacement may jump. It carries three unknowns: the **normal discontinuity** D_n (opening-positive; closure is $-D_n$) and two in-plane **ride** components, along the element's strike (D_s) and dip (D_t) axes. Mining an element means requiring the total traction across it to vanish.

DD elements are **triangles**. The displacement field of each is the analytic, artefact-free triangular dislocation of Nikkhoo & Walter (2015), built from three angular dislocations. The stress is obtained from Hooke's law on a high-order finite-difference gradient of that displacement. An element's influence on itself is the limit of the two-sided average as the evaluation point approaches the element.

Each element is given its reef's strike as a hint, so that the ride components of a faceted reef share one frame. Older models carry rectangular DD elements, which still load.

1.5 Combined FF and DD models

A model with both kinds of element is solved as **one coupled system**: the excavation walls feel the reef and the reef feels the walls. DD influence columns are of a different magnitude from FF columns (of order $2\mu/L$ against order 1), so the library scales the columns of the system before solving and rescales the unknowns afterwards. This keeps the iterative solver's accuracy uniform across both kinds.

The combined formulation uses the infinite-domain kernel only. A DD row with the Mindlin kernel is refused.

1.6 In-situ stress

The in-situ stress has one vertical and two horizontal principal components:

- **Uniform:** σ_v as entered; $\sigma_H = k_H \cdot \sigma_v$ along the trend; $\sigma_h = k_h \cdot \sigma_v$ at right angles to it.
- **Gravitational:** $\sigma_v = \gamma \cdot \text{depth below the ground surface } Z$, with σ_H and σ_h in the same ratios at every depth.

Compression is positive. Above the ground surface, the gravitational field becomes tensile, which is why GEM4D checks for meshes above it (*User Guide* §6.6).

1.7 What a solution holds

A solved model is the set of element geometries and their solved unknowns: force densities on FF elements, discontinuities on DD elements, for every mining step. Everything GEM4D reports afterwards is evaluated from them. A `.bem4d` file stores exactly this, in mine coordinates, together with the Setup, so reloading restores the solution without solving again.

2. Field-point evaluation and classification

2.1 Evaluating a point

Once the system is solved, the stress and displacement at any point in the rock follow by summing every element's influence on that point, then adding the in-situ stress. Each element is a **constant-strength** element: one force density (FF) or one displacement discontinuity (DD) over the whole triangle.

Elements close to the point are integrated analytically, so the Kelvin result is exact right up to the wall. The Mindlin kernel adds a free-surface image term, integrated by adaptive quadrature; its error falls below 1 % of the in-situ stress beyond about two element edges from a wall. With either kernel,

values within about one element of a wall are less accurate than values deeper in the rock, because a constant-strength element cannot represent the stress gradient across its own width.

2.2 Classification

Every evaluated point is classified:

1. **Inside test.** A point inside a closed excavation surface — or, with the Mindlin kernel, above the ground surface — is Interior.
2. **Distance.** The distance to the nearest wall is calculated and reported as `DistToSurface_m`, negative inside.
3. **Near-surface band.** A point outside the excavations but within the near-surface band (Setup, in multiples of the mean element edge) is NearSurface; everything else is Far. The default band is 0 with the Kelvin kernel — its integration is exact to the wall — and 2 with Mindlin.

The inside test relies on closed surfaces. Points near holes or non-manifold edges may be classified on the wrong side (§3.2).

NearSurface is a flag only: the values written to tables and files are the evaluated values.

2.3 Interior values

By default every map writes NaN for interior points. With *Iso-surface treats interior as in-situ* ticked in Setup, all three maps — Grid, Octree and datagrid, single models and sequences, combined files and per-step files — write the in-situ stress instead, with the in-situ values of every derived column and zero displacement, so that an iso-surface interpolated from the points relaxes smoothly into the void. The in-situ tensor is evaluated at the point itself, so a gravitational field is back-filled with the stress for that depth. The Status column still reads Interior.

On a sequence each step is classified on its own geometry.

2.4 The wall probe

An FF element is a sheet of force. Exactly on the sheet, the stress is the **average** of its two sides — the traction jumps by half the force density across it. A point placed exactly on a wall therefore does not return the wall stress on the rock side.

The energy maps (**Stored on surfaces, Released on surfaces**) need the rock-side value. They evaluate each element at a point **1 % of the element size** ($0.01 \cdot \sqrt{A}$) into the rock, and report the result at the wall position. This keeps the evaluation on the analytic near-field integrator, so the probe converges to the rock-side limit and is meaningful at the wall despite the general near-wall caution of §2.1.

3. How GEM4D prepares meshes

This chapter describes what **Add FF and DD objects** does to a mesh between the mesh table and the solver (*User Guide*, Chapter 6). It covers the FF ingest, the quality checks, the FF/DD classification and the preparation of DD reefs.

3.1 FF meshes

Each ticked FF mesh goes through five stages:

1. **Triangulation.** Quads and triangle strips become triangles; lines and points are dropped.
2. **Welding.** Vertices that nearly coincide are merged, restoring the shared edges that DXF round-trips often break.

3. **Orientation.** Each connected, manifold region is walked from its own seed so that all its triangles face the same way. The solver's sign convention is then applied internally, so the user never has to fix the winding.
4. **Degenerate triangles.** Triangles of zero area are skipped and counted.
5. **Elements.** Every remaining triangle becomes one constant-strength fictitious-force element.

3.2 Watertightness

GEM4D counts the **open** edges (used by one triangle) and the **non-manifold** edges (used by three or more) of the combined surface. A few such edges at floating-point-noise scale have no effect on the solver, so they are only logged. The threshold is the larger of 10 edges and $1/2\,000$ of the estimated edge count. Above it, the completion summary reports *"Not watertight"*: field points near the defects may then be classified on the wrong side of a wall.

3.3 Mixed normals and seams

The orientation walk in §3.1 repairs a flipped face anywhere on a manifold region. It cannot cross a **non-manifold seam** — an internal wall or a T-junction where three or more triangles share an edge — because it cannot tell which sheet is the excavation wall. Sheets on either side of a seam therefore keep the winding the file gave them.

GEM4D measures the winding of the source file before any repair. A mesh is flagged with **mixed normals** when both of these hold:

- some edges are traversed the same way by both their triangles (one of the pair is wound backwards), and
- the mesh has non-manifold seams above the noise threshold of §3.2.

A flipped face on a seam-free mesh is repaired and only noted (*"re-oriented automatically"*).

A wrongly facing element carries the wrong boundary traction, so results near it are wrong. It does not, on its own, stop the solve converging: a test model with about 2 500 reversed elements converged normally and was silently wrong. Meshes with mixed normals do, however, usually also carry coincident sheets and slivers, which do stop it. The mesh this check was built on ran 2 000 GMRES iterations without converging; after **Remesh for BEM analysis** it converged in 59.

3.4 Body diagnostics

GEM4D splits each mesh into its connected closed bodies and checks them for three defects that the watertightness test cannot see:

Defect	Test	Why it matters
Zero-volume body	Characteristic thickness $2V/A$ below 1 cm	Coincident elements with opposite normals give linearly dependent equations.
Interpenetrating pair	Sample points of one closed body lie strictly inside another	Collocation points of one body sit inside the other's void.
Sliver triangle	Longest edge ÷ height above 100	Near-collinear triangles ill-condition the system without failing the zero-area test.

A mesh with more than 1 500 bodies skips the pairwise overlap test. Bodies that merely touch are not reported, because exact face contact is ambiguous.

3.5 FF / DD classification

When the mesh table opens, every mesh is classified, whatever its size. The verdict is cached until the mesh changes.

Open or closed. A mesh is **open** when any of these holds:

- its open edges exceed the noise threshold of §3.2;
- it encloses no volume and has any open edge at all;
- its open edges are more than 5 % of all its edges (this catches the rim of a small, folded open surface, which the absolute threshold would miss).

Holes are not a rim. Open edges are grouped into loops. If the mesh encloses volume and no body's largest loop spans a quarter of that body, the loops are holes and cracks in a closed shell. The mesh is then classified as the closed body it is, and the tooltip counts the holes. An open surface's rim is its whole outline, so it stays open.

Thin or volumetric. A closed body is **thin** when its thickness is at most 15 % of its mid extent **and** at most 30 m. The 30 m ceiling exists because the plate test alone is scale-free: an open pit with a sublevel-shrinkage block (characteristic thickness 118 m) passes it. A DD plane stands in for an orebody only while the mining width is small compared with the distance at which results are wanted.

Measured	Shape	Suggested Type
Open	Surface	DD (width to be entered)
Closed, thin	Thin solid w m	DD, with the measured thickness as width
Closed, no volume	Sheets	DD
Closed, volumetric	Solid	FF
Thin and volumetric bodies in one mesh	Mixed	FF — split the mesh

3.6 DD surfaces

A DD row's own triangles become **triangular displacement-discontinuity elements**. There is no reference plane, lattice or limit on how the surface folds.

1. **Duplicates.** Coincident triangles — the second sheet of a two-sided export — are removed.
2. **Weld.** The surface is welded into one mesh. A triangle that would be the third on an edge is kept as a loose element.
3. **Refinement.** If the longest edge exceeds $1.5 \times$ the DD element size, the surface is refined: edges are split, then several quality passes collapse short edges and smooth the triangles, projected back onto the source. The outline is kept as a polyline and its corners are pinned. The remesh runs one connected piece at a time and is single-threaded, so identical inputs give identical elements. A finer surface is used exactly as it is; above 200 000 triangles that is noted.
4. **Slivers.** A triangle whose altitude is below 1 % of the element size is dropped and counted. Such a triangle's collocation point sits against a neighbour's edge and makes the system singular; the slit it leaves is negligible.

3.7 Thin closed solids: the mid-surface

A thin closed solid is reduced to its **mid-surface**, and each element takes the roof-to-floor gap measured under it as its width.

The roof and floor are found locally, with no preferred direction:

1. Every face looks **inward** from its centroid and three further points.
2. If it meets an **opposing** face — within 30° of anti-parallel — within three widths, at a consistent distance, it is a **sheet** face. A side wall looks along the plate, meets faces at varying distances, and is not a sheet.
3. Sheet faces join into patches across gentle creases. The patches are then sorted into roof and floor through their pairings.

4. Each roof vertex moves onto the mid-plane of the roof faces around it. At a crease this lands on the exact corner of the two mid-planes.

So a thin solid folds as freely as a surface: a fold through 90° keeps both limbs, each on its own mid-plane.

Further rules handle awkward geometry:

- **Inside-out regions.** A vertical extrusion of an overhanging plane puts the offset copy through the original, so roof and floor swap sides there. These faces find their partner outward instead, and the mid-surface runs between them at the small width they actually have.
- **Rim-wall pairs.** The side walls of a narrow ledge or raise can face each other within three widths and pass as sheets. A patch that meets a sheet across a sharp crease, with a crossing distance more than 1.5 times the sheet's, is demoted to a wall — but only when the patch it pairs with qualifies too.
- **The rim.** The mid-surface's rim is simplified at the scale of the thickness (twice the width, at most half the element size). Where the rim crosses a steep limb, the roof and floor rims sit a thickness or more apart in plan, and the mid-surface rim would otherwise step between them. An open surface's own outline is never touched.
- **Width fallback.** Where the roof-to-floor gap cannot be measured under an element, the row's Width is used and counted.

3.8 Glued and stepped blocks

A stope-shape export often comes as separate closed blocks — one per panel or sub-block — with neighbours sharing walls. GEM4D treats such a row as the **union** of its blocks:

- **Shared walls removed.** A wall two blocks share is in the file twice, often under two different triangulations. Each copy finds its twin facing it at zero distance; both are removed before the roof and floor are sought.
- **Coplanar slivers joined.** A face rejected by the crossing test but coplanar (within 5°) with an adjacent sheet face joins that sheet. This catches the sliver along a cut whose centroid ray leaves through the cut.
- **One block at a time.** Each block's sheet is refined on its own, so blocks touching at a corner cannot be run together.
- **Overlaps.** Where two blocks were drawn over the same ground, the smaller sheet's elements there are dropped and counted.
- **Gaps closed.** Sub-blocks at slightly different levels make roof and floor step at every joint, so each sub-block's mid-surface is a separate piece. Gaps narrower than **1 m** between pieces are closed: each rim vertex moves halfway towards the nearest rim of another piece, so the sheets meet. A move that would fold an element by more than 20° is halved once, then refused, and counted. A wider gap — a real pillar — stays. The 1 m acts on the sheets, which lie at mid-height.

A part of a reef where blocks interpenetrate at different orientations cannot give a clean mid-surface and comes out in fragments. The remedy is to fix the blocks in CAD, or to make one envelope with **Remesh for BEM analysis**.

§7.4 gives the measured effect of open and closed gaps on the results.

3.9 Several surfaces and several steps

- **Pillars.** A hole in a surface is a pillar: nothing is mined there.
- **Prepared once.** Each row is prepared once, and the same elements go into every step it belongs to. An element mined at step 2 is the identical triangle in step 5's model, which is what lets the released-energy bookkeeping recognise it.

- **Ground already covered.** Rows are ordered by their first step. A later row's element is skipped in a step an earlier row shares when the earlier surface passes **through** it — a cast along the element's normal, within a quarter of the element size. An element that merely sits beside the earlier surface's edge is new reef and is kept.

3.10 Refusals

The build is refused, naming the mesh and leaving the existing model untouched, for:

- a DD row with the Mindlin kernel (the combined FF + DD formulation has no free surface);
- a DD row without a width;
- a surface with no usable triangles, or with no fittable plane (all points collinear);
- a "thin solid" with no roof or no floor sheet — usually an open sheet flagged as closed;
- a surface that an earlier surface covers in every one of its steps.

4. Solver paths and memory

This chapter explains how GEM4D chooses the numerical method and manages memory during **Run model**. The user-facing settings are in *User Guide* §5.4.

4.1 Direct and iterative solves

The FF and DD influence equations form a dense system with three unknowns per element.

- **Up to 6 700 elements: direct LU factorisation.** Exact, with no iterations. The matrix is assembled as one array of $(3N)^2$ numbers — 3.2 GB at 6 700 elements. The factorisation uses Intel MKL; without it, the same solve runs roughly 100 times slower, and GEM4D warns at the first solve of a session. The BEM linear algebra: line in GEM4D_diag.log names the provider in use.
- **Above 6 700 elements: iterative GMRES.** The matrix is never formed whole; each iteration applies it element by element (EBE) through the H-matrix of §4.2. Iteration stops when the relative residual falls below the GMRES tolerance set in Setup.

The choice is made per solve, and per step on a sequence.

4.2 The H-matrix

The iterative solver does not store the full matrix. It builds a **hierarchical matrix** (H-matrix) from an octree of the elements:

- **Near-field blocks** — pairs of nearby element groups — are kept dense.
- **Far-field blocks** — well-separated groups — are compressed by adaptive cross approximation (ACA) to a relative tolerance of 1e-4 (or the GMRES tolerance, if that is coarser).

Each GMRES iteration multiplies by this matrix. A block held in RAM costs little per iteration; a block **matrix-free** — recalculated every iteration — costs much more. On a 3 888-element benchmark, one matrix-vector product took 8.7 ms with everything cached, 36 ms with the near field matrix-free, and 248 ms with the far field matrix-free.

4.3 The cache budget

H-matrix cache (GB) in Setup sets how much RAM these blocks may use.

- **Auto** probes the RAM free when each step starts, after releasing the previous step's cache, and keeps a 6 GB reserve, with a floor of 2 GB. The figure shown in the Setup dialog is only a snapshot from when it was opened.
- **A positive value** pins the budget.

- **0** removes the cap.

The near field is charged against the budget first, at its exact measured size. The far field is compressed and cached in what remains. Far-field blocks that do not fit are left matrix-free. The post-build status line reports the total built and the number of blocks left matrix-free (· N MF).

4.4 The near-field threshold

Near-field threshold sets the element count above which the near field is matrix-free from the start. **Auto** derives it from the computer's **total** RAM: about 41 000 elements at 16 GB, 72 000 at 32 GB, 110 000 at 64 GB and 160 000 at 128 GB.

Because the threshold reads total RAM while the budget reads free RAM, the two can disagree on a busy machine. GEM4D resolves this by measurement: when the threshold is on Auto, a step whose dense near field would take more than **75 %** of the budget runs its near field matrix-free instead, so the far field stays cached. Losing the far cache costs far more than losing the near one (§4.2), and the break-even sits close to 75 %. The switch is skipped when the threshold is set by hand, and when the budget is uncapped (0).

4.5 What is logged

Each solved step writes one line to GEM4D_diag.log: the element count, LU or EBE, the near-field mode with its measured size (and whether the 75 % rule switched it), the cache budget (auto, override or uncapped) and the free RAM at the time. The solver's own phase messages, including the H-matrix size actually built, are logged too. Comparing these lines between a fast run and a slow one answers "why did the same model get slower?" — almost always, less free RAM.

4.6 Energy evaluations

The **Stored on surfaces** and **Released on surfaces** maps evaluate the stress at every boundary element **directly**, without the H-matrix. Direct evaluation is exact, needs no block RAM, and measured faster at every size tried (Chapter 8). The field maps (Grid, Octree, datagrid) also evaluate directly, streaming the points.

5. Energy formulations

5.1 Stored strain-energy density

At any point, the elastic strain-energy density follows from the principal stresses and the Setup E and v:

$$u = [\sigma_1^2 + \sigma_2^2 + \sigma_3^2 - 2\nu(\sigma_1\sigma_2 + \sigma_2\sigma_3 + \sigma_3\sigma_1)] / 2E$$

With σ in MPa and E in GPa, u comes out directly in kJ/m³. The expression is quadratic in stress, so the sign convention does not affect it.

The skin map evaluates u at the wall probe of §2.4; the SED checkbox on the Map dialogs evaluates the same u at the user's field points. Both use the same function and the same E and v.

5.2 Released energy on FF boundaries

When mining step k creates new excavation surface Σ_k , the energy released is

$$\Delta W_k = \frac{1}{2} \int \text{over } \Sigma_k \text{ of } \mathbf{t}^{(k-1)} \cdot \Delta \mathbf{u}^{(k)} dS$$

where $\mathbf{t}^{(k-1)}$ is the traction that acted across that surface in the previous equilibrium state (the in-situ field for step 1), and $\Delta \mathbf{u}^{(k)} = \mathbf{u}^{(k)} - \mathbf{u}^{(k-1)}$ is how far the wall moved between the two states. The traction on the new face falls from $\mathbf{t}^{(k-1)}$ to zero while the face moves by Δu , hence the $\frac{1}{2}$. Surfaces

mined at earlier steps are already free of traction and contribute nothing, so the integral runs over the new surface only.

Discretised, each element contributes $\mathbf{w}_e = \frac{1}{2} \mathbf{t}^{(k-1)} \cdot \Delta \mathbf{u}_e$ in kJ/m², and $\Delta \mathbf{W} = \sum \mathbf{w}_e \mathbf{A}_e$, reported in MJ. With stress in MPa, displacement in mm and area in m², **1 MPa × 1 mm = 1 kJ/m²**, so no conversion is needed.

The traction is evaluated at the wall probe of §2.4. On the wall itself the two-sided average would not vanish on a face that is already free.

As-mined bookkeeping. In As-mined mode a step's shape is not all new surface:

- an element whose probe lies inside the previous step's void is **excluded** (ExcludedInterior): that ground was already mined, and no traction acted there;
- an element whose previous traction was already at least half relieved is **kept** but counted (NearZeroTraction): it is old wall re-covered by the new shape, and contributes almost nothing.

The excavated volume increment ΔV_k comes from the closed-body volumes of consecutive steps.

5.3 Released energy on DD elements

On DD elements the release is read directly from the solved discontinuities:

$$\Delta \mathbf{W}_k = -\frac{1}{2} \sum \text{over } \Pi_k \text{ of } \mathbf{t}^{(k-1)} \cdot \Delta \mathbf{D}^{(k)} \cdot \mathbf{A}_e$$

over the elements Π_k mined at step k , with $\mathbf{t}^{(k-1)}$ the traction across the reef before the step and $\Delta \mathbf{D}$ the change in the solved discontinuity.

The leading minus is not a fudge. The solver's normal discontinuity D_n is opening-positive, so a mined seam under compression solves negative — it closes. The minus turns that closure into a positive release. For the same reason the Closure_mm column is reported positive for a seam that closes.

Elements mined at earlier steps are free of traction and contribute nothing. They are recognised by geometric matching — the same centre, area and normal in both steps — and counted as CarriedOver. This is why each DD row is prepared once and the identical elements enter every step (§3.9).

The touchdown count compares each element's closure with its own mining width; the mined volume is $\sum \text{area} \times \text{width}$.

5.4 Path independence

In linear elasticity the total energy released depends only on the final geometry. Summed over any staging, $\sum \Delta \mathbf{W}_k$ equals the single-step $\mathbf{W}$ of the final excavation.

- On FF models GEM4D reproduces this to about **0.4 %** on the validation models.
- On DD elements it is **exact to round-off** (0.000 % on the validation sequence), because the discontinuity is a solved unknown read straight off the element, with no probe or field evaluation in the increment.

5.5 Area normalisations

- $\Delta \mathbf{W} / \Delta \mathbf{V}$ is well defined in both authoring modes.
- $\Delta \mathbf{W} / \mathbf{A}$ uses the new FF surface area. It is reported in Additive mode only, because in As-mined mode part of a step's mesh may re-cover existing wall.
- **ERR on DD** is $\Delta \mathbf{W}$ over the area mined, measured in the reef plane (ERR_MJm2) and in plan (ERR_plan_MJm2). The plan area is the reef-plane area times cos(dip), so the plan figure grows as $1/\cos(\text{dip})$.

6. Excess shear stress and Kirsch tangential stress

6.1 Plane-specific ESS

On a plane with unit normal $\mathbf{n}$, from the total, compression-positive stress tensor $\mathbf{S}$:

- normal stress $\sigma_N = \mathbf{n}^T \mathbf{S} \mathbf{n}$
- shear stress $\tau = \sqrt{(|\mathbf{S} \mathbf{n}|^2 - \sigma_N^2)}$
- **ESS** = $\tau - c - \tan \varphi \cdot \sigma_N$

The slip direction is the direction of the shear traction on the plane; GEM4D reports it as a downward plunge (BEM_SlipPl_deg) and azimuth (BEM_SlipAz_deg).

The tensile clamp. The friction term $c + \tan \varphi \cdot \sigma_N$ is a resistance, physical only while the plane is clamped ($\sigma_N > 0$). With the clamp on, it becomes $c + \tan \varphi \cdot \max(\sigma_N, 0)$. Because σ_N enters ESS with a minus sign, a negative σ_N in the unclamped form *raises* ESS.

6.2 Coulomb ESS

The largest ESS over all plane orientations at a point is a function of σ_1 and σ_3 only (Ryder, 1988, Eq. 13):

$$\text{ESS}_{\text{Coulomb}} = \sec \varphi \cdot (\sigma_1 - \sigma_3)/2 - \tan \varphi \cdot (\sigma_1 + \sigma_3)/2 - c$$

It is reached on the Mohr-Coulomb **critical plane**, whose normal lies in the σ_1 – σ_3 plane at $45^\circ + \varphi/2$ from σ_1 (the plane itself lies at $45^\circ - \varphi/2$ to σ_1). There are two such planes, mirror images about σ_1 and equally likely to slip; GEM4D draws both. On the critical plane σ_N is always compressive, so the tensile clamp does not apply. By construction, Coulomb ESS is never less than the plane-specific ESS at the same point.

6.3 Kirsch tangential stress

For a row with plunge and trend (Dip, DipDir) read as a drive's axis:

1. The axis $\mathbf{a}$ is built from the plunge and trend.
2. Two unit vectors $\mathbf{e}_1, \mathbf{e}_2$ perpendicular to $\mathbf{a}$ span the drive's cross-section.
3. The stress tensor is projected onto the section: the 2×2 in-plane tensor has components $\mathbf{e}_i^T \mathbf{S} \mathbf{e}_j$.
4. Its principal values are **SigMax** $\geq$ **SigMin**.
5. The Kirsch maximum tangential stress on the wall of a circular opening is **Tang** = $3 \cdot \text{SigMax} - \text{SigMin}$.

SigMax and SigMin always lie between the global σ_1 and σ_3 . Appendix A works one row through by hand, including the construction of the section axes.

7. Accuracy and validation

This chapter collects the measurements behind the guidance in the *User Guide*. They were made on test geometries with known answers; they show how the method behaves, not the accuracy of any particular site model.

7.1 Thin bodies and element size

Measured on disc slabs and narrow stopes, with elements from twice a body's thickness down to a third of it, released energy, skin energy and wall stresses **converge as the mesh is refined** — but they have converged only once the elements are **no larger than the thickness**. Coarser than that,

readings wander. On a 0.2 m-thick disc 20 m across, released energy read –5 %, +0.1 % and +2.7 % against the penny-crack value at elements 2.5, 1.7 and 1.25 times the thickness.

On the finest meshes tested, thin FF slabs read slightly **above** the zero-thickness (penny-crack) value, and by less as the slab gets thinner. On a 20 m-diameter disc: about +3 % at 1 m thick (11 880 elements), and –0.1 / +0.6 / +1.3 % at 0.5 m thick on 5 040 / 20 160 / 45 360 elements.

The penny value is a reference, not a floor. The exact released energy of a thin void can lie slightly *below* it: 0.7 % below for a spheroidal void 1 m thick at its centre and 20 m across. A reading a little under the penny value is not by itself an error.

7.2 Triangle shape

The integration decides how close a point is to an element from the element's **area**. A long, thin triangle is therefore treated as farther away than its tips really are. A 4 : 1 triangle is integrated noticeably less accurately at points within about two element sizes of it. Where stresses are read near a wall, keep triangles no more elongated than about 3 : 1.

Extreme slivers (longest edge ÷ height above 100) do more harm: they ill-condition the whole system. On one CAD stope export triangulated with long needle facets throughout, the slivers made the whole solution unreliable.

7.3 Coarsening a thin body

Remesh and smooth at an edge length larger than a body's thickness cannot keep the body's rim. On a 2.3 m-thick stope, a 4.5 m target removed half the volume and a 2.3 m target removed an eighth. **Remesh for BEM analysis**, which samples a thin body at about a third of its thickness, kept the volume within 1.5 % on the same stope.

7.4 Gaps and steps between DD pieces

On a 100 × 50 m test stope 2 m wide (DD, Kelvin, 5 m elements):

Join between two halves of the reef	Closure and released energy, against the continuous stope
Open 0.3 m gap in the plane of the reef	5 % low — the gap acts as a rib
Open 0.3 m step	4 % high — the two offset sheet edges reinforce each other
Gap or step closed (§3.8)	Within 0.3 %

This is why GEM4D closes gaps narrower than 1 m between the pieces of a mid-surface.

7.5 Released energy: the pressurised sphere

A spherical cavity under hydrostatic in-situ stress p releases $W = \pi p^2 a^3 / (2G)$ on excavation (Jaeger, Cook & Zimmerman). GEM4D's faceted spheres converge to it from above: **+3.36 %** at 320 elements, **+2.48 %** at 1 280 and **+1.44 %** at 5 120. Appendix B works one element of this case by hand.

7.6 DD elements: the penny crack and curved reefs

Element size. Against Sneddon's penny-shaped crack, on a circular reef triangulated exactly to its outline:

Element edge	Elements	Error in crack volume ($\approx$ released energy)
span / 12	216	+8.6 %
span / 20	600	+5.4 %

The error falls monotonically with refinement and is always positive: constant-strength elements are slightly too compliant at the mined edge. Appendix C works this case by hand.

Curved and folded reefs. A reef folded onto a cylinder of radius R converges to the flat answer as $(a/R)^2$, with a measured order of 1.99. A thin curved slab modelled as FF and as DD agrees to 2.5 %.

FF and DD on the same stope. A 62×30 m planar stope 2.3 m wide, modelled as a thin FF solid and as a DD surface, gives the same released energy to within a few percent.

Path independence. See §5.4: about 0.4 % on FF models, exact to round-off on DD elements.

8. Result changes from BEM library updates

GEM4D's stress solver is the BasRock BEM library. When a library update changes results, this appendix records what moved and by how much, so you can decide which earlier models to re-run.

8.1 BasRock.BEM 1.2.1.0 (October 2026)

GEM4D now uses version 1.2.1.0 of its BEM library, a follow-up to 1.2.0.0 (§8.2). **Results from 1.2.0.0 and earlier are superseded on every fictitious-force (FF) model;** DD-only results are unchanged.

What changed	What moved
Far-field integration tightened — the library integrated an element with a single point from too short a distance, which put the opposite face of a thin body meshed finer than its thickness into that band	Every FF result: by about 1 % or less on ordinary, bulky stopes; by several % on thin ones — on a $62 \times 30 \times 2.3$ m planar stope meshed at a third of its thickness, released energy +4 % and the hangingwall-centre σ_1 -7 %; skin energy on the broad faces of thin slabs by 10–30 % (a few kJ/m ³). Thin slabs no longer read low (<i>User Guide</i> Chapter 7). It costs some run time: iterative solves about +15 % at 40 000 elements and a few % at 70 000 (where the RAM budget, not the integration, sets the pace — Chapter 4), and skin evaluations +7–14 % (the Energy maps are still faster than before, having moved to direct evaluation — next row)
Stored and Released energy maps corrected element by element — the compressed (H-matrix) field evaluation these two buttons used could put a large error on single elements of a large thin model while the totals stayed within 0.5 %	Per-element Stored-energy and Released-energy maps — on a 45 360-element thin disc one skin element read 829 MPa where the exact value is 0.04 MPa, and 5 % of the elements were more than 1 MPa off. GEM4D now evaluates both energy maps directly (the exact method), which was also the faster one on every model tested
Iterative-solver compression tightened	Models above 6 700 elements move slightly toward the exact answer — skin σ_1 by up to about 1 MPa on large thin bodies, about 0.1 MPa on stopes. The iterative solver also needs more RAM: measured +3 to +28 % in total, most on stopes close together (four stopes with 5 m pillars +28 %, a compact block +8 %), and its near-field share up to ~3× on thin stopes. With plenty of free RAM nothing changes; where RAM is short, more of the matrix is recomputed every iteration and the switch to the matrix-free near field (Chapter 4) comes at smaller models. The thin and compact test models solved as fast or faster (fewer iterations)

Re-run models whose energy maps, thin-stope results or iterative solves came from 1.2.0.0 or earlier.

8.2 BasRock.BEM 1.2.0.0 (October 2026)

GEM4D version 1.2.0.0 of its BEM library was released after an independent audit found and fixed errors in the library's numerical integration and matrix compression. **BEM results computed with earlier versions are superseded — re-run the models** wherever they used one of these paths:

Model or output	What moved
Any model with flat-faced excavations (box stopes, drives)	Wall and skin values (Stored energy, Kirsch, near-wall stresses) by about 5–9 %; stresses more than ~0.25 m from the walls only within discretisation error
Mindlin (half-space) models	Near-wall stresses and skin energy — on a test box 400 m deep the mean skin energy moved 13 % and stresses near the walls by up to ~7 MPa; shallow excavations also change near the ground surface
Models above 6 700 elements (iterative solver)	Element strengths by up to ~7 %; DD reef closures were 56–78 % too low
Stored energy and Released energy on models of 1 024 elements or more	Released energy was 18–26 % too low — worse on thin stopes, where it could turn negative
A Mindlin model whose roof lies exactly on the ground surface	Released energy is now reported (it came back as NaN); the roof elements, which have no rock above them, are left out

Two improvements come with the changes: models with several separated excavations solve faster on the iterative path (four staggered stopes: 337 s → 1.6 s), and direct solves of about 5 500–6 700 elements, which failed with an out-of-memory error, now complete.

Appendix A — Worked example: checking BEM_Tang_MPa by hand

This appendix walks one fictitious minode row through the whole tunnel-section Kirsch calculation, so you can reproduce BEM_SigMax_MPa, BEM_SigMin_MPa and BEM_Tang_MPa with a calculator or a spreadsheet and satisfy yourself that the numbers are what they claim to be. Nothing here is a separate implementation — it is the same arithmetic GEM4D runs, written out.

Before you start: tick **both** *Kirsch tangential* and *Full stress tensor (Sxx...Sxz)* in the *Map to datagrid* options (*User Guide* Chapter 14). Without the tensor columns the row shows you the answer but not the six numbers it was computed from, and there is nothing to check.

A.1 The example row

A minode in a gently declining development drive heading south-east, after *Map to datagrid*:

Column	Value	Where it comes from
Easting / Northing / RL	4250.0 / 7180.0 / -620.0	Your loaded minode CSV — the point the tensor was evaluated at. Not used in the arithmetic below.
Dip	2.5°	Your loaded minode CSV — the tunnel axis plunge (positive = drive falling; a rising incline is negative).
DipDir	130.0°	Your loaded minode CSV — the tunnel axis trend, degrees clockwise from North.
BEM_Sxx_MPa	32.40	Solved total stress (in-situ + induced) at that point.
BEM_Syy_MPa	21.70	Solved total stress (in-situ + induced) at that point.
BEM_Szz_MPa	28.10	Solved total stress (in-situ + induced) at that point.

Column	Value	Where it comes from
BEM_Sxy_MPa	-4.60	Solved total stress (in-situ + induced) at that point.
BEM_Syz_MPa	3.20	Solved total stress (in-situ + induced) at that point.
BEM_Sxz_MPa	7.90	Solved total stress (in-situ + induced) at that point.
BEM_SigMax_MPa	33.47	← reproduced in §A.5
BEM_SigMin_MPa	16.58	← reproduced in §A.5
BEM_Tang_MPa	83.83	← reproduced in §A.6

Axes are mine coordinates: **x = East, y = North, z = up (RL)**, compression positive (*User Guide* §3.5). The stress tensor **S** is symmetric:

$$S = \begin{bmatrix} S_{xx} & S_{xy} & S_{xz} \\ S_{xy} & S_{yy} & S_{yz} \\ S_{xz} & S_{yz} & S_{zz} \end{bmatrix} = \begin{bmatrix} 32.4 & -4.6 & 7.9 \\ -4.6 & 21.7 & 3.2 \\ 7.9 & 3.2 & 28.1 \end{bmatrix} \text{ MPa}$$

One note on decimals before you start. GEM4D works in full double precision and rounds only when writing a cell. Intermediates below are printed to 4–6 decimals, so re-typing a printed intermediate into the next step can land you one unit off in the last decimal shown. That drift never reaches the two decimals the grid displays: the checks that matter (§A.5, §A.6) come out exact.

A.2 Step 1 — the tunnel axis

The CSV gives the drive's plunge *p* and trend *t* as a **line**, so the first move is to build the unit vector **a** pointing along the tunnel. It heads toward the trend azimuth and tilts down by the plunge:

$$\mathbf{a} = (\cos p \cdot \sin t, \cos p \cdot \cos t, -\sin p)$$

Note this is *not* the pole-of-a-plane formula (that one puts *cos* in the *z* slot). Kirsch is solved in the **cross-section of the drive** — the plane perpendicular to **a** — so the tunnel axis itself is the plane normal.

With *p* = 2.5° and *t* = 130° (*cos p* = 0.999048, *sin p* = 0.043619, *sin t* = 0.766044, *cos t* = -0.642788):

$$\begin{aligned} a_x &= 0.999048 \times 0.766044 = 0.765315 && \text{(East)} \\ a_y &= 0.999048 \times -0.642788 = -0.642176 && \text{(North)} \\ a_z &= -0.043619 && \text{(Up – negative, the drive is falling)} \end{aligned}$$

Sanity checks: $|\mathbf{a}| = 1$; the horizontal part of **a** bears $\text{atan2}(a_x, a_y) = \text{atan2}(0.765315, -0.642176) = 130^\circ$ — the trend, as it should (that argument order gives an azimuth clockwise from North); and $\text{atan2}(-a_z, \sqrt{a_x^2 + a_y^2}) = 2.5^\circ$ — the plunge. The **sign** of **a** does not matter to anything that follows (flipping it swaps the two in-plane axes but leaves the answer unchanged), which is also why a rising drive recorded as *Dip* = -2.5° toward 310° gives the identical result.

A.3 Step 2 — the two cross-section axes

The Kirsch relation is two-dimensional, so the tensor has to be projected onto the cross-section. That needs two perpendicular unit vectors lying *in* that section. GEM4D builds them with cross products, and for any drive flatter than about 64° plunge they work out to the two directions you would sketch on a section anyway — **horizontally across the drive** (wall to wall) and **roof to floor**:

$$\begin{aligned} \mathbf{u} \text{ (across the drive)} &= (\cos t, -\sin t, 0) \\ \mathbf{v} \text{ (roof to floor)} &= (-\sin p \cdot \sin t, -\sin p \cdot \cos t, -\cos p) \end{aligned}$$

For this row:

$$\begin{aligned} \mathbf{u} &= (-0.642788, -0.766044, 0) && \text{horizontal, bearing } 220^\circ \text{ (trend } \pm 90^\circ \text{ – square across the drive)} \\ \mathbf{v} &= (-0.033414, 0.028038, -0.999048) && \text{plunging } 87.5^\circ \text{ toward } 310^\circ \text{ (straight down the section: vertical, tipped back by the } 2.5^\circ \text{ grade)} \end{aligned}$$

$\mathbf{u}$, $\mathbf{v}$ and $\mathbf{a}$ are mutually perpendicular unit vectors — check any pair with a dot product and you get 0.

A.4 Step 3 — project the stress tensor onto the cross-section

Squeeze the 3×3 tensor down to the 2×2 stress state seen in the drive's cross-section. Multiply $\mathbf{S}$ into each axis first:

$$\begin{aligned} \mathbf{S} \cdot \mathbf{u} : \\ \begin{aligned} x: & 32.4 \times -0.642788 + -4.6 \times -0.766044 + 7.9 \times 0 = -17.3025 \\ y: & -4.6 \times -0.642788 + 21.7 \times -0.766044 + 3.2 \times 0 = -13.6663 \\ z: & 7.9 \times -0.642788 + 3.2 \times -0.766044 + 28.1 \times 0 = -7.5294 \end{aligned} \\ \\ \mathbf{S} \cdot \mathbf{v} : \\ \begin{aligned} x: & 32.4 \times -0.033414 + -4.6 \times 0.028038 + 7.9 \times -0.999048 = -9.1041 \\ y: & -4.6 \times -0.033414 + 21.7 \times 0.028038 + 3.2 \times -0.999048 = -2.4348 \\ z: & 7.9 \times -0.033414 + 3.2 \times 0.028038 + 28.1 \times -0.999048 = -28.2475 \end{aligned} \end{aligned}$$

then dot each result back onto an axis:

$$\begin{aligned} \sigma_{uu} &= \mathbf{u} \cdot (\mathbf{S} \cdot \mathbf{u}) = -0.642788 \times -17.3025 + -0.766044 \times -13.6663 + 0 \times -7.5294 \\ &= 21.5909 \text{ MPa} \\ \sigma_{vv} &= \mathbf{v} \cdot (\mathbf{S} \cdot \mathbf{v}) = -0.033414 \times -9.1041 + 0.028038 \times -2.4348 + -0.999048 \times -28.2475 \\ &= 28.4566 \text{ MPa} \\ \sigma_{uv} &= \mathbf{u} \cdot (\mathbf{S} \cdot \mathbf{v}) = -0.642788 \times -9.1041 + -0.766044 \times -2.4348 + 0 \times -28.2475 \\ &= 7.7172 \text{ MPa} \end{aligned}$$

Read physically: **21.59 MPa horizontally across the drive, 28.46 MPa roof-to-floor, with 7.72 MPa of shear between them.**

A.5 Step 4 — the in-plane stress extremes

σ_{uu} , σ_{vv} and σ_{uv} are a plain 2-D stress state, so SigMax and SigMin come straight off **Mohr's circle** — centre = the mean, radius = r :

$$\begin{aligned} \text{mean} &= (\sigma_{uu} + \sigma_{vv}) / 2 = (21.5909 + 28.4566) / 2 = 25.0237 \\ \text{half} &= (\sigma_{uu} - \sigma_{vv}) / 2 = (21.5909 - 28.4566) / 2 = -3.4328 \\ r &= \sqrt{\text{half}^2 + \sigma_{uv}^2} = \sqrt{(-3.4328)^2 + (7.7172)^2} = \sqrt{71.3392} = 8.4463 \\ \\ \text{SigMax} &= \text{mean} + r = 25.0237 + 8.4463 = 33.4700 \rightarrow \text{BEM_SigMax_MPa} = 33.47 \quad \checkmark \\ \text{SigMin} &= \text{mean} - r = 25.0237 - 8.4463 = 16.5775 \rightarrow \text{BEM_SigMin_MPa} = 16.58 \quad \checkmark \end{aligned}$$

Both grid columns are now reproduced. Note SigMax $\geq$ SigMin by construction, and both lie *in the cross-section* — they are **not** the same thing as BEM_S1_MPa / BEM_S3_MPa, the full 3-D principals, which take no notice of the drive direction. They are always bracketed by them: for this tensor the global principals are $\sigma_1 = 38.63$, $\sigma_2 = 27.10$, $\sigma_3 = 16.47$, and indeed $38.63 \geq 33.47 \geq 16.58 \geq 16.47$. Only a drive exactly perpendicular to σ_1 gets SigMax all the way up to σ_1 — which is precisely why a drive square to σ_1 reports a higher BEM_Tang_MPa than one along it.

A.6 Step 5 — the Kirsch tangential stress

$$\text{BEM_Tang_MPa} = 3 \times \text{SigMax} - \text{SigMin} = 3 \times 33.4700 - 16.5775 = 83.8324 \rightarrow 83.83 \quad \checkmark$$

This last step is deliberately checkable straight off the grid, with no intermediate working at all: **$3 \times 33.47 - 16.58 = 83.83$** . If those three columns ever fail that arithmetic on a row, something is genuinely wrong — please report it.

A.7 The same thing in a spreadsheet

To check a whole column rather than one row, the projections in §A.4 collapse to three formulas (with the tunnel dip in A2, trend in B2, Sxx...Sxz in C2:H2):

$$\begin{aligned} u_x &= \text{COS}(\text{RADIANS}(B2)) & v_x &= -\text{SIN}(\text{RADIANS}(A2)) * \text{SIN}(\text{RADIANS}(B2)) \\ u_y &= -\text{SIN}(\text{RADIANS}(B2)) & v_y &= -\text{SIN}(\text{RADIANS}(A2)) * \text{COS}(\text{RADIANS}(B2)) \\ u_z &= 0 & v_z &= -\text{COS}(\text{RADIANS}(A2)) \end{aligned}$$

$$\begin{aligned}
\sigma_{uu} &= S_{xx} \cdot u_x^2 + S_{yy} \cdot u_y^2 + S_{zz} \cdot u_z^2 + 2 \cdot (S_{xy} \cdot u_x \cdot u_y + S_{yz} \cdot u_y \cdot u_z + S_{xz} \cdot u_x \cdot u_z) \\
\sigma_{vv} &= S_{xx} \cdot v_x^2 + S_{yy} \cdot v_y^2 + S_{zz} \cdot v_z^2 + 2 \cdot (S_{xy} \cdot v_x \cdot v_y + S_{yz} \cdot v_y \cdot v_z + S_{xz} \cdot v_x \cdot v_z) \\
\sigma_{uv} &= S_{xx} \cdot u_x \cdot v_x + S_{yy} \cdot u_y \cdot v_y + S_{zz} \cdot u_z \cdot v_z + S_{xy} \cdot (u_x \cdot v_y + u_y \cdot v_x) + S_{yz} \cdot (u_y \cdot v_z + u_z \cdot v_y) + \\
&\quad S_{xz} \cdot (u_x \cdot v_z + u_z \cdot v_x)
\end{aligned}$$

then Mohr's circle exactly as in §A.5.

A.8 Notes on the fine print

- **The choice of in-plane axes does not affect the answer.** Across-drive and roof-to-floor are just a convenient pair. Any two perpendicular axes lying in the cross-section give different σ_{uu} , σ_{vv} and σ_{uv} , but the *same* Mohr's circle — the centre and radius are invariant under rotation within the plane, so SigMax, SigMin and hence BEM_Tang_MPa are unchanged. This is why the recipe above is valid at every plunge even though GEM4D internally switches its reference axis for drives steeper than about 64° plunge (shafts and steep declines, where the across-drive vector would degenerate and a different pair is used instead).
- **NaN rows.** If Dip/DipDir or any tensor component is missing or non-numeric, all three columns are written as NaN together — you will never see a stress extreme sitting next to a NaN tangential.
- **Interior rows.** With *Backfill interior* on, an Interior row's tensor columns hold the **in-situ** tensor, and the Kirsch columns are computed from that same in-situ tensor — so the row still checks out internally. See *User Guide* Chapter 13.
- **What the number means** is a separate question from whether it is computed correctly. BEM_Tang_MPa is only meaningful if the row's Dip/DipDir genuinely describes the tunnel direction at that point (*User Guide* Chapter 14 — minode data), and it inherits every assumption in *User Guide* Chapter 20 — linear elastic, dry, constant-stress elements near the wall — plus the Kirsch idealisation itself: a circular opening, long compared with its diameter, evaluated at the boundary.

Appendix B — Worked example: checking a released-energy row by hand

This appendix works **one element** of the released-energy validation case end to end — traction, wall movement, $\frac{1}{2} \cdot t \cdot \Delta u$, and the area weight — so you can reproduce a w_{kJm2} row with a calculator and satisfy yourself the map means what *User Guide* Chapter 17 says it does. As in Appendix A, nothing here is a separate implementation; it is the same arithmetic GEM4D runs, written out.

The case is the one the implementation is validated against: a **spherical cavity in an infinite medium under hydrostatic compression**, which has a closed-form answer to check the total against.

B.1 The example

Quantity	Value	Where it comes from
Cavity radius a	5.0 m	The geometry
Hydrostatic in-situ stress p	30.0 MPa	Setup → In-Situ Stress ($\sigma_v = \sigma_H = \sigma_h = 30$)
Young's modulus E	70 GPa	Setup → Rock Properties
Poisson's ratio v	0.25	"
Shear modulus G = $E / (2(1+v))$	28 GPa = 28 000 MPa	Derived

Quantity	Value	Where it comes from
The element	a facet whose centroid sits on the +x axis of the cavity	Any facet works; this one has the tidiest normal
Element area A_e	0.2443 m ²	From the mesh (≈ 312.7 m ² of faceted surface over 1 280 elements)

Axes are mine coordinates, **compression positive**, and element normals point **into the excavation** (*User Guide* §3.5). Rock occupies $|x| > a$, so at the point ***x* = (+5, 0, 0)** the normal into the void is

$$\hat{n} = (-1, 0, 0)$$

B.2 Step 1 — the traction that acted on the face

This is step 1 of the sequence, so the previous state is the **in-situ** field. With compression-positive hydrostatic stress the tensor is $\sigma = p \cdot I$, and the traction is the ordinary $\mathbf{t} = \sigma \cdot \hat{n}$:

$$\mathbf{t} = 30 \cdot \begin{vmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{vmatrix} \cdot \begin{pmatrix} -1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} -30 \\ 0 \\ 0 \end{pmatrix} \text{ MPa}$$

So ***t* = (−30, 0, 0) MPa** — 30 MPa pointing along $-x$, i.e. *toward* the cavity from this face. (Sanity check on the convention: the material that used to fill the cavity was pushing the wall **outward**, in $+x$. The compression-positive tensor gives the equal-and-opposite bookkeeping vector, which is what pairs with the convergence below to give a positive product. *User Guide* Chapter 17.)

B.3 Step 2 — how far the wall moved

For an unsupported spherical cavity under remote pressure p , the Lamé thick-sphere solution with the outer radius taken to infinity gives a uniform **inward** wall displacement

$$\delta = p \cdot a / (4G) = 30 \times 5 / (4 \times 28\,000) = 150 / 112\,000 = 0.00133929 \text{ m} = 1.33929 \text{ mm}$$

The wall converges **into the void**, i.e. along ***+n***, so at our element

$$\Delta u = \delta \cdot \hat{n} = 1.33929 \times (-1, 0, 0) = (-1.33929, 0, 0) \text{ mm}$$

($u^{(0)} = 0$ because GEM4D reports *induced* displacement, measured from the pre-mining state — so for step 1 the increment Δu is simply $u^{(1)}$.)

B.4 Step 3 — the release density

$$\begin{aligned} w_e &= \frac{1}{2} (\mathbf{t} \cdot \Delta \mathbf{u}) \\ &= \frac{1}{2} [(-30)(-1.33929) + (0)(0) + (0)(0)] \\ &= \frac{1}{2} \times 40.1786 \\ &= 20.089 \text{ kJ/m}^2 \end{aligned}$$

Both vectors point the same way, so the product is **positive** — energy is released, not absorbed. This is the number in the w_kJm2 column, and the unit falls out of the mixed units directly: **MPa × mm = kJ/m²** (*User Guide* Chapter 17).

B.5 Step 4 — the element's contribution

$$w_e \cdot A_e = 20.089 \times 0.2443 = 4.908 \text{ kJ}$$

Every facet of a hydrostatically loaded sphere carries the same density (the geometry is symmetric), so summing over the whole surface is one multiplication:

$$\begin{aligned} W &= w_e \times A = 20.089 \text{ kJ/m}^2 \times 4\pi(5)^2 \text{ m}^2 \\ &= 20.089 \times 314.159 \\ &= 6\,311 \text{ kJ} = 6.311 \text{ MJ} \end{aligned}$$

B.6 The check

The closed form for a pressurised spherical cavity (Jaeger, Cook & Zimmerman) is

$$\begin{aligned} W &= \pi p^2 a^3 / (2G) \\ &= \pi \times 30^2 \times 5^3 / (2 \times 28\,000) \\ &= \pi \times 900 \times 125 / 56\,000 \\ &= 6.3112 \text{ MJ} \end{aligned}$$

with **p** in MPa, **a** in m, **G** in MPa → the answer in MJ. The hand-worked total above reproduces it to the digits shown, which is the point: the per-element density, the ½, the sign and the unit chain are all doing what *User Guide* Chapter 17 claims.

What GEM4D actually reports is slightly higher, because a faceted mesh is not a sphere. On the validation series the computed total converges from above — **+3.36 %** at 320 elements, **+2.48 %** at 1 280, **+1.44 %** at 5 120 — so on a reasonably refined model expect *W* within a percent or two of the closed form, always on the high side. If your own check lands *negative*, or wildly off, that is a real signal: check the mesh — normals, watertightness, slivers and zero-volume bodies (*User Guide* Chapter 7) — and the in-situ stress entered in Setup (*User Guide* Chapter 5).

B.7 Checking a whole column in a spreadsheet

For an arbitrary model you will not know δ analytically, but you can still audit the **relationship**: export the release map (*w_kJm2*) and the corresponding per-step displacement columns, and for each row confirm

$$w_kJm2 = 0.5 \times (t_x \cdot \Delta u_x + t_y \cdot \Delta u_y + t_z \cdot \Delta u_z) \quad [\text{MPa} \times \text{mm}]$$

For **step 1** the traction is the in-situ tensor projected on the element normal, which you can build from the Setup values; for later steps it is the *previous* step's total stress at the same point, which is why the step-1 row is the one that audits cleanly with nothing but the Setup dialog in front of you.

Appendix C — Worked example: checking a tabular panel by hand

Appendix B checks one excavation-wall element of *User Guide* Chapter 17. This appendix does the same for one **reef panel** of *User Guide* Chapter 18: in-situ traction → resolved into the panel's own axes → solved discontinuity → release density → area weight → whole-reef ERR against a closed form. Same arithmetic GEM4D runs, written out.

C.1 The case

The Sneddon **penny-shaped crack** — a circular reef panel in an infinite elastic body under a uniform field normal to it. It is the natural benchmark because a crack of zero thickness is exactly what fictitious-force elements cannot represent — they need a closed volume — and it has an exact answer.

Quantity	Value
Reef radius <i>a</i>	100 m (a 200 m span)
Field normal to the reef, <i>p</i>	60 MPa
Young's modulus <i>E</i>	70 GPa
Poisson's ratio <i>v</i>	0.25
Shear modulus $G = E/2(1+v)$	28 GPa
Panel size <i>g</i>	5 m (= span / 40)

The reef is horizontal, so its local axes are $\hat{s}$ = East, $\hat{t}$ = North, $\hat{n}$ = up. Take the panel at the centre of the reef.

C.2 Step 1 — the traction that acted on the reef before mining

In-situ is vertical compression p with no shear, so the stress tensor at the panel is $\sigma_{zz} = 60$ MPa, everything else zero. The traction on a plane with normal $\hat{n}$ is $\mathbf{t} = \boldsymbol{\sigma} \cdot \hat{n}$:

$$\mathbf{t} = \boldsymbol{\sigma} \cdot \hat{n} = (0, 0, 60) \text{ MPa} \quad (\hat{n} = (0, 0, 1))$$

Resolved into the panel's own axes (this is the step that matters on a dipping reef — here it is trivial because the axes are already aligned):

$$\begin{aligned} t_s &= \mathbf{t} \cdot \hat{s} = 0 \text{ MPa} \\ t_t &= \mathbf{t} \cdot \hat{t} = 0 \text{ MPa} \\ t_n &= \mathbf{t} \cdot \hat{n} = 60 \text{ MPa} \quad \leftarrow \text{the reef was carrying 60 MPa of normal load} \end{aligned}$$

For step 1 this is simply the in-situ field. For a later step it is the **previous solved state's** total stress at the panel centre — the reef sees whatever the earlier mining left behind, which is exactly why an abutment panel releases more than a panel in the middle of a mined-out area.

C.3 Step 2 — how far the reef closed

This is a **solved unknown**, not a field evaluation — GEM4D reads it straight off the panel. The Closure_{mm} column for the centre panel reads

$$\text{closure} = 207.14 \text{ mm} = 0.20714 \text{ m}$$

Sneddon's closed form for the centre of a pressurised penny crack gives the total opening

$$\begin{aligned} \delta_{\max} &= 4(1-\nu) \cdot p \cdot a / (\pi \cdot G) \\ &= 4 \times 0.75 \times 60 \times 100 / (\pi \times 28\,000) \quad [\text{MPa}, \text{m}, \text{MPa}] \\ &= 18\,000 / 87\,965 \\ &= 0.20463 \text{ m} = 204.63 \text{ mm} \end{aligned}$$

so the solved value is **+1.2 %** on the closed form — the panel-size discretisation, exactly as *User Guide* Chapter 18's table predicts.

Sign. Internally the solver stores this as $D_n = -0.20714 \text{ m}$: its normal discontinuity is *opening-positive*, and the seam closed. GEM4D reports the closure **positive**. Both conventions are in play in the formula below, which is why it carries a leading minus.

C.4 Step 3 — the release density

$$\begin{aligned} w_e &= -\frac{1}{2} (t_s \cdot D_s + t_t \cdot D_t + t_n \cdot D_n) \\ &= -\frac{1}{2} [0 + 0 + (60)(-0.20714)] \quad [\text{MPa} \times \text{m}] \\ &= -\frac{1}{2} \times (-12.428) \\ &= 6.214 \text{ MJ/m}^2 = 6\,214 \text{ kJ/m}^2 \end{aligned}$$

Positive — energy released. The two minuses cancel: the reef was in compression (t_n positive) and it closed (D_n negative). **A negative value here means the sign or the frame was mishandled**, never a physical result for a mined seam.

Unit fact: MPa $\times$ m = MJ/m², so with closure in metres the density lands in MJ/m² directly; the w_{kJm^2} column is this $\times$ 1000.

C.5 Step 4 — the panel's contribution, and the reef total

$$\begin{aligned} A_e &= g^2 = 5 \times 5 = 25 \text{ m}^2 \\ \Delta W_e &= w_e \cdot A_e = 6.214 \times 25 = 155.4 \text{ MJ} \quad \leftarrow \text{this one panel} \end{aligned}$$

Summing that over all 1 264 panels of the disc gives what GEM4D reports:

	Value
ΔW (whole reef)	133.56 GJ
Reef-plane area	31 600 m ² (1 264 × 25 m ²)
ERR = ΔW / A_{reef}	4.226 MJ/m²

C.6 The check

Sneddon's released energy for the whole crack is

$$\begin{aligned}
 W &= 4(1-\nu) \cdot p^2 \cdot a^3 / (3G) & (= p^2 a^3 / G \text{ at } \nu = 0.25) \\
 &= (60)^2 \times (100)^3 / 28\,000 & [\text{MPa}^2, \text{m}^3, \text{MPa}] \\
 &= 3\,600 \times 1\,000\,000 / 28\,000 \\
 &= 128\,571 \text{ MJ} = 128.57 \text{ GJ}
 \end{aligned}$$

and the mean ERR over the true disc area $\pi \cdot a^2 = 31\,416 \text{ m}^2$ is

$$\text{ERR} = 128\,571 / 31\,416 = 4.093 \text{ MJ/m}^2$$

	Closed form	GEM4D	Error
ΔW	128.57 GJ	133.56 GJ	+3.9 %
ERR	4.093 MJ/m ²	4.226 MJ/m ²	+3.3 %

Both high by the amount *User Guide* Chapter 18's panel-size table predicts for $g = \text{span}/40$ — the bias is positive because a rasterised outline is slightly larger than the true circle and constant-strength panels are slightly too compliant at the edge. Halving g roughly halves the error: the same disc at $\text{span}/20$ (316 panels, the **Suggest** default) comes out **+6.9 %**, and at $\text{span}/10$ (80 panels) **+15.1 %**.

C.7 On a dipping reef

Nothing above changes except §C.2. For a reef dipping θ the traction resolves into three non-zero components, the two ride terms contribute real work, and the reef-plane area is the sum of the panel areas while the plan area is that $\times \cos \theta$. GEM4D reports **both** — ERR_MJm2 (reef plane, the headline) and ERR_plan_MJm2. On this flat example the two columns are identical; on a 30° reef the plan figure is 15 % higher, and on a 70° reef it is nearly three times higher, which is why *User Guide* Chapter 18 uses the reef-plane figure and says so.