

Physics Guide

Wakefields, material losses and waveguide scattering

Wake potentials, Green functions, coupling impedance, and the extraction of transverse monopole, dipole and quadrupole contributions in EMUZAG Studio; bunch factors, RMS spreads, material energy losses and waveguide S-parameters.

Scope of this guide

The three chapters connect physical definitions to the current Studio workflow. The User Guide remains the reference for general application controls. This is a derivation and workflow guide, not a report of newly performed simulations.

Release 1.1.2 / 26 October 2026

Exact Solver and Studio revisions: see the package manifest

Contents

1 Wakefields and Impedances **1**

1.1 Geometry representation and numerical convergence . . . 1
1.2 The beam–structure problem . . . 1
1.3 Longitudinal and transverse wake potentials . . . 2
1.4 What changes below the speed of light? . . . 3
1.5 Direct and indirect integration . . . 4
1.6 The Panofsky–Wenzel relation . . . 6
1.7 Wake potential and Green function . . . 6
1.8 Wake and coupling impedance . . . 7
1.9 Loss factor, kick factor and RMS spread . . . 9
1.10 Taylor expansion and the meaning of transverse components . . . 12
1.11 Full-structure extraction in Studio . . . 13
1.12 Synchronous opposite-charge pair: one-run dipole extraction . . . 14
1.13 Derivative with respect to observation position . . . 16
1.14 One reflection plane: the one-run PEC dipole method . . . 16
1.15 Two reflection planes: quarter-domain calculations . . . 18
1.16 Checks before publishing a component or impedance . . . 19
1.17 References and software scope . . . 19

2 Material Losses and Power Accounting **21**

2.1 Beam energy loss is not necessarily material heat . . . 21
2.2 Material interfaces and mesh refinement . . . 21
2.3 Physical mechanisms and their representations . . . 22
2.4 What the power-loss monitor accumulates . . . 24
2.5 CSV normalization: power, cumulative loss and joules . . . 24
2.6 Studio workflow and interpretation checks . . . 25

3 Definition and Calculation of S-Parameters **27**

3.1 Modal amplitudes and the S-matrix . . . 27
3.2 Supported setup in Studio . . . 28
3.3 Reading and checking the result . . . 29

CHAPTER 1

Wakefields and Impedances

1.1 Geometry representation and numerical convergence

STL approximates a boundary with triangles. The curved route retains supported surfaces and their trimmed boundaries; a hybrid scene combines explicit per-body choices. These describe the geometry supplied to the common Cartesian mesher. They do not remove spatial discretization error from the field calculation.

For a reproducible comparison, hold dimensions, material laws, priorities, domain and excitation fixed. Separate STL-approximation refinement from grid refinement. The default face-area rule reconstructs finite-volume material areas from edge lengths; the optional surface rule integrates the represented, priority-resolved boundaries. Compare the quantity of interest, such as wake, absorbed power or a scattering coefficient, under the chosen refinement. A finer preview, smaller source file or successful meshing is not a convergence result.

A zero-thickness PEC or resistive screen has a separate two-sided topology and is not a thin conductive volume. Both STL and native screens must resolve holes, gaps and their contacts in that topology. A rejected unresolved coincidence or subcell opening is not permission to fill the hole or change the physical domain. The User Guide describes explicit representation choices and recovery. Geometry admission does not extend the applicability of a wall or sheet material law.

1.2 The beam–structure problem

We prescribe a rigid bunch moving on a straight line through a structure at constant speed $v = \beta c$, with $0 < \beta \leq 1$. Maxwell's equations, the material laws and the boundary conditions determine its electromagnetic field. A witness samples that field without changing the prescribed source or the field solution. This separates the wake calculation from a self-consistent beam-dynamics calculation.

Use the same Cartesian axes as Studio: x is longitudinal, while y, z are transverse. The source and witness positions are independent:

$$\boldsymbol{r}_q = (y_q, z_q), \qquad \boldsymbol{r}_w = (y_w, z_w), \qquad s = vt - x. \tag{1.1}$$

The bunch centre crosses $x = 0$ at $t = 0$. A witness at lag s crosses position x at $t = (x + s)/v$. Positive s means behind the bunch centre; it is not the global longitudinal coordinate of a point in the geometry. Source and witness have the same constant speed in this guide.

For a transversely idealized line source,

$$\rho(x, \boldsymbol{r}_\perp, t) = Q\lambda(vt - x)\delta^{(2)}(\boldsymbol{r}_\perp - \boldsymbol{r}_q), \qquad \boldsymbol{J} = v\rho\boldsymbol{e}_x, \qquad \int_{-\infty}^{\infty} \lambda(s)\,\mathrm{d}s = 1. \tag{1.2}$$

Q is the signed drive charge; λ is a normalized line profile, in m^{-1}. A finite transverse profile may replace the transverse delta. The Gaussian used in the ordinary Studio bunch workflow is

$$\lambda(s) = \frac{e^{-s^2/(2\sigma_s^2)}}{\sqrt{2\pi}\sigma_s}, \qquad I(s) = Qv\lambda(s). \tag{1.3}$$

σ_s is an RMS length, not the full bunch length. A grid with five cells per RMS length means $\sigma_s/\Delta x = 5$.

EMUZAG also accepts several synchronous line sources, each with its own signed charge Q_j and fixed transverse position $\boldsymbol{r}_j$. They share this Gaussian profile, RMS length and velocity, with no relative longitudinal delay:

$$\rho = \lambda(vt - x) \sum_j Q_j \delta^{(2)}(\boldsymbol{r}_\perp - \boldsymbol{r}_j), \qquad \boldsymbol{J} = v\rho \boldsymbol{e}_x. \tag{1.4}$$

Source and observation coordinates are separate explicit lists of absolute (y, z) positions, including when a single source and witness coincide. An observation adds no charge and does not move when a source is edited. Section 1.12 describes an opposite-charge pair.

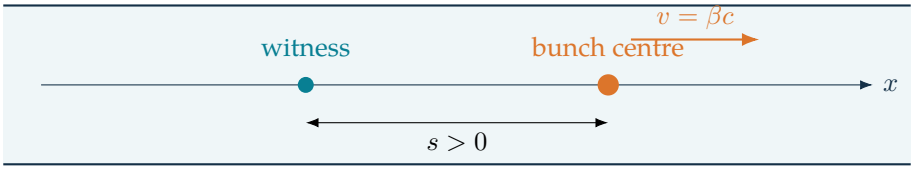


Figure 1.1: Longitudinal separation at a common time. Transverse source and witness positions need not coincide. The drawing is schematic, not a field simulation.

The geometry includes the vacuum aperture and all relevant materials, including lossy walls or volume media. Radiation boundary conditions remove outgoing waves; PML is a numerical absorber, not a physical material to include in a beam-coupling impedance budget. Mirror reduction is legitimate only when both geometry and material arrangement have the stated reflection symmetry.

1.3 Longitudinal and transverse wake potentials

For any component, put a bar over the field to denote evaluation on the witness:

$$\overline{\boldsymbol{E}}(x, s) = \boldsymbol{E}\big(x, \boldsymbol{r}_w, (x + s)/v; \boldsymbol{r}_q\big), \qquad \overline{\boldsymbol{B}}(x, s) = \boldsymbol{B}\big(x, \boldsymbol{r}_w, (x + s)/v; \boldsymbol{r}_q\big). \tag{1.5}$$

The transverse Lorentz force per unit witness charge is

$$\boldsymbol{F}_\perp = \boldsymbol{E}_\perp + v\boldsymbol{e}_x \times \boldsymbol{B}_\perp, \qquad F_y = E_y - vB_z, \qquad F_z = E_z + vB_y. \tag{1.6}$$

Consequently, a transverse electric-field plot alone is not a transverse wake. Electric and magnetic forces must be combined on the same witness trajectory.

For a fixed integration interval $[x_1, x_2]$, define

$$W_\parallel^{\mathrm{dir}}(\boldsymbol{r}_q, \boldsymbol{r}_w, s) = -\frac{1}{Q} \int_{x_1}^{x_2} \overline{E}_x(x, s)\, \mathrm{d}x, \tag{1.7}$$

$$\boldsymbol{W}_\perp^{\mathrm{dir}}(\boldsymbol{r}_q, \boldsymbol{r}_w, s) = \frac{1}{Q} \int_{x_1}^{x_2} \overline{\boldsymbol{F}_\perp}(x, s)\, \mathrm{d}x. \tag{1.8}$$

The longitudinal convention here is positive for deceleration of a like-sign witness; the transverse convention is the signed physical kick. These are bunch wake potentials: the finite longitudinal profile is already included.

One sign conversion when reading Studio

Studio's X curve uses the accelerating convention:

$$W_X^{\mathrm{Studio}} = -W_{\parallel}, \qquad W_Y^{\mathrm{Studio}} = W_y, \qquad W_Z^{\mathrm{Studio}} = W_z.$$

Keep this sign conversion when checking Panofsky–Wenzel or comparing longitudinal impedance formulas. No sign conversion is needed for Y or Z.

Current Solver output stores the physical witness voltage for every source list:

$$V_\alpha^{\Sigma}(\boldsymbol{r}_w, s) = \sum_j Q_j W_\alpha^{\mathrm{Studio}}(\boldsymbol{r}_j, \boldsymbol{r}_w, s). \tag{1.9}$$

There is no independently authored normalization charge. Linearity makes this the response of one combined field calculation. The source sum may vanish; it is never used as a denominator. For an equal opposite pair, a separate artifact stores the source-moment coefficient described in Section 1.12. Other source lists retain physical voltage. The single-source energy and bunch-factor formulas below keep their single-bunch meaning; a neutral source pair is not a zero-energy physical bunch.

For witness charge q, the first-order changes are

$$\Delta\mathcal{E} = -qQW_{\parallel}, \qquad \Delta\boldsymbol{p}_\perp = \frac{qQ}{v}\boldsymbol{W}_\perp, \qquad \boldsymbol{\theta} \simeq \frac{qQ}{pv}\boldsymbol{W}_\perp, \tag{1.10}$$

where p is the incoming momentum magnitude. Multiplying a V/pC curve by the drive charge in pC gives a voltage, not an angle. In particular, replacing pv by pc is only justified at $\beta = 1$.

For the same normalized Gaussian witness distribution, define

$$k_{\parallel} = \int \lambda(s) W_{\parallel}(s)\,\mathrm{d}s, \qquad \boldsymbol{k}_\perp = \int \lambda(s)\boldsymbol{W}_\perp(s)\,\mathrm{d}s. \tag{1.11}$$

The bunch energy loss in this convention is $Q^2 k_{\parallel}$ and the signed transverse voltage is $Q\boldsymbol{k}_\perp$. Studio's loss/kick voltage metrics already include the drive charge; do not multiply those metric rows by charge again. Port-dependent space-charge exchange is not automatically irreversible dissipation.

1.4 What changes below the speed of light?

The field equations do not change. The source transport speed, the magnetic part of the Lorentz force and the incident/reference field do change. The field of a rigid bunch in a uniform, perfectly conducting pipe makes the distinction particularly clear. With $\zeta = x - vt$ and $\gamma^{-2} = 1 - \beta^2$, its scalar potential satisfies

$$\left(\nabla_\perp^2 + \gamma^{-2}\partial_\zeta^2\right) V^{\mathrm{p}} = -\frac{\rho}{\varepsilon_0}, \qquad V^{\mathrm{p}} = 0 \quad \text{on a PEC wall.} \tag{1.12}$$

The corresponding fields and force obey

$$\boldsymbol{E}_\perp^{\mathrm{p}} = -\boldsymbol{\nabla}_\perp V^{\mathrm{p}}, \qquad E_x^{\mathrm{p}} = -\gamma^{-2}\partial_\zeta V^{\mathrm{p}}, \tag{1.13}$$

$$\boldsymbol{B}_\perp^{\mathrm{p}} = \frac{v}{c^2}\boldsymbol{e}_x \times \boldsymbol{E}_\perp^{\mathrm{p}}, \qquad \boldsymbol{F}_\perp^{\mathrm{p}} = \gamma^{-2}\boldsymbol{E}_\perp^{\mathrm{p}}. \tag{1.14}$$

These relations describe the nonradiating uniform-pipe reference, not arbitrary fields in a cavity or a lossy wall.

At $\beta = 1$, this ideal pipe has zero longitudinal force and exact transverse electric–magnetic cancellation. At $\beta < 1$, both longitudinal space charge and residual transverse force generally remain. Along a fixed-lag witness they are independent of x: integrating them over an infinitely long pipe generally diverges. A finite integration length or an explicitly subtracted reference is therefore part of the physical definition, not merely a numerical convenience.

For equal pipes, a globally pipe-subtracted wake is possible. For unequal pipes, one must identify the entrance and exit references separately. A comparison at different β must retain the same geometry, source position, witness position and integration planes, and must state which space-charge contribution is included.

Causality does not mean setting every negative lag to zero

At $\beta = 1$, the point-charge wake of the usual straight passive structure has no response ahead of the source. A Gaussian bunch wake can still be nonzero at $s < 0$, ahead of its centre. At $\beta < 1$, fields can overtake the source and a point-source response itself need not vanish at negative s; electromagnetic causality is not violated. Do not truncate that response by imposing an ultrarelativistic one-sided wake assumption.

In the current Release, finite velocity is supported for stationary FDTD explicit and stationary dispersion-free explicit or implicit ATI. Moving dispersion-free remains $\beta = 1$ only. The optional beam self-wake correction compensates a numerical reference-transport error; it is not a subtraction of physical pipe space charge. Release dispersion-free/ATI silently ignores that preference. New and imported projects default this option off. It changes the evolving fields as well as the wake and can introduce artificial excitation outside a shielded beam channel. Existing explicitly enabled settings are preserved. Open-space incident-field treatment remains intrinsic and independent of this option. These software choices do not redefine Eqs. (1.7)–(1.8).

1.5 Direct and indirect integration

Direct integration samples the complete fields at the appropriate witness times and accumulates Eqs. (1.7)–(1.8). Radiation may continue to act on the witness in the uniform pipes. Stopping the direct integral near the cavity can omit this contribution even when the field evolution is accurate.

Indirect integration reconstructs the remaining uniform-pipe integrals from field data at a cross section or a suitable end snapshot, using the field equations. It avoids following the entire tail in a long three-dimensional domain. It is neither smoothing nor a Fourier deconvolution. Classical indirect methods are discussed in [3]; the two-reference convention below follows the owner's unequal-pipe formulation [4].

1.5.1 Two local pipe references

Let Ω_1, Ω_2 be the uniform entrance and exit pipe sections. In each tail subtract the rigid-bunch field of that pipe only:

$$(\boldsymbol{E}_j^{\mathrm{sc}}, \boldsymbol{B}_j^{\mathrm{sc}}) = (\boldsymbol{E}, \boldsymbol{B}) - (\boldsymbol{E}_j^{\mathrm{p}}, \boldsymbol{B}_j^{\mathrm{p}}), \qquad j = 1, 2. \tag{1.15}$$

The finite-reference assembled wake is

$$QW_{\parallel}^{[x_1,x_2]} = -\int_{-\infty}^{x_1} \overline{E}_{x,1}^{\mathrm{sc}}\,\mathrm{d}x - \int_{x_1}^{x_2} \overline{E}_x\,\mathrm{d}x - \int_{x_2}^{\infty} \overline{E}_{x,2}^{\mathrm{sc}}\,\mathrm{d}x, \tag{1.16}$$

$$Q\boldsymbol{W}_{\perp}^{[x_1,x_2]} = \int_{-\infty}^{x_1} \overline{\boldsymbol{F}}_{\perp,1}^{\mathrm{sc}}\,\mathrm{d}x + \int_{x_1}^{x_2} \overline{\boldsymbol{F}_{\perp}}\,\mathrm{d}x + \int_{x_2}^{\infty} \overline{\boldsymbol{F}}_{\perp,2}^{\mathrm{sc}}\,\mathrm{d}x. \tag{1.17}$$

The centre contains the complete field, including space charge. Only the two tails are reference-subtracted. No reference pipe is artificially continued through the cavity or joined to the other reference inside it.

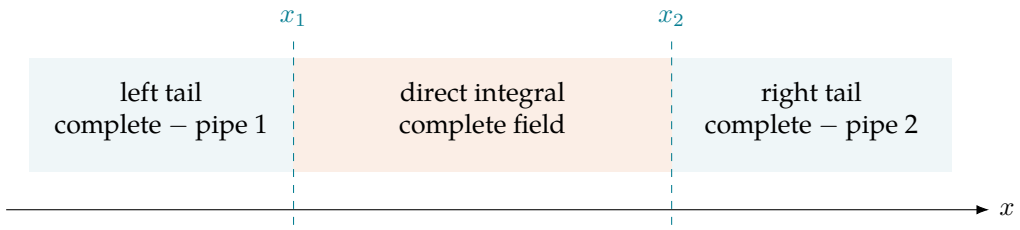


Figure 1.2: The two-port definition. Subtraction is local to each tail; the complete central field must not also be replaced by a scattered field.

The transverse correction reconstructs the electric-plus-magnetic tail, including both longitudinal-electric and longitudinal-magnetic information. It is not, in general, obtained by differentiating just a longitudinal correction. In the finite-velocity formulation the operator $\nabla_\perp^2 + \gamma^{-2}\partial_s^2$ couples longitudinal lags; at $\beta = 1$ it reduces to a transverse Poisson operator.

If the pipes are identical, the assembled result differs from the globally pipe-subtracted wake by the complete uniform-pipe contribution over $L = x_2 - x_1$. In particular,

$$W_\parallel^{[x_1,x_2]} = W_\parallel^{\mathrm{sc}} - \frac{L}{Q} E_x^{\mathrm{p}}, \qquad \boldsymbol{W}_\perp^{[x_1,x_2]} = \boldsymbol{W}_\perp^{\mathrm{sc}} + \frac{L}{Q} \boldsymbol{F}_\perp^{\mathrm{p}}. \tag{1.18}$$

Thus a uniform pipe at finite β can have a nonzero assembled wake even though its scattered wake is zero. This is not a failure of indirect integration. Extending only the exit plane by δL adds $-\delta L E_{x,2}^{\mathrm{p}}/Q$ longitudinally and $\delta L \boldsymbol{F}_{\perp,2}^{\mathrm{p}}/Q$ transversely. Port invariance at finite β requires removing or accounting for that known change of definition.

1.5.2 What to inspect in Studio

The saved Direct result and the assembled Indirect result coexist. Their difference is the reconstructed tail contribution, not automatically an error. Do not average them. Compare complete quantities with identical references; refine the grid and extend the relevant recording when checking convergence.

For finite β, use the actual terminal coordinates in the Solver log message

```
direct_wake_calculated_between=z1=...,z2=...
```

The historical log names z_1, z_2 denote the longitudinal planes called x_1, x_2 here. They are not necessarily the inner TF/SF interfaces or the ends of the visible cavity. Scattered field on the outer side of a TF/SF interface can still lie inside the complete-field direct integration interval.

Uniform reference sections must satisfy the assumptions of the selected indirect method. A physical closed reference pipe may lie inside a larger transverse-PML domain; PML alone does not supply a closed-pipe reference. Direct integration is available without indirect integration. The current open-space initial-field workflow uses Direct with pipe indirect integration disabled. None of this is a blanket requirement that every finite-beta calculation have two open ends or transversely closed outer domain boundaries.

1.6 The Panofsky–Wenzel relation

The original theorem is due to Panofsky and Wenzel [2]. For constant v, Faraday's law gives a local identity on the witness trajectory:

$$\partial_s \overline{\boldsymbol{F}_\perp} = \frac{\mathrm{d}\overline{\boldsymbol{E}}_\perp}{\mathrm{d}x} - \boldsymbol{\nabla}_\perp \overline{E}_x. \tag{1.19}$$

For example,

$$\begin{aligned} \partial_s(E_y - vB_z) &= v^{-1}\partial_t E_y - \partial_t B_z \\ &= v^{-1}\partial_t E_y + \partial_x E_y - \partial_y E_x. \end{aligned} \tag{1.20}$$

The first two terms in the last line form the total x derivative along the witness. No ultrarelativistic approximation was made.

Integrating over a fixed finite interval yields

$$\partial_s \boldsymbol{W}_\perp^{\mathrm{dir}} = \boldsymbol{\nabla}_\perp W_\parallel^{\mathrm{dir}} + \frac{\overline{\boldsymbol{E}}_\perp(x_2, s) - \overline{\boldsymbol{E}}_\perp(x_1, s)}{Q}. \tag{1.21}$$

The derivative is with respect to the *witness* coordinates, with source position fixed. The extra term contains transverse electric fields, not the transverse Lorentz force.

Adding both local-reference tails cancels the complete device fields at the ports and leaves

$$\boxed{\partial_s \boldsymbol{W}_\perp^{[x_1,x_2]} = \boldsymbol{\nabla}_\perp W_\parallel^{[x_1,x_2]} + \frac{\boldsymbol{E}_{\perp,2}^{\mathrm{p}}(\boldsymbol{r}_w, -s) - \boldsymbol{E}_{\perp,1}^{\mathrm{p}}(\boldsymbol{r}_w, -s)}{Q}.} \tag{1.22}$$

This is a relation between the longitudinal and transverse wake potentials themselves. No additional "Panofsky wake potential" is introduced.

When the familiar endpoint-free formula is valid

$\partial_s \boldsymbol{W}_\perp = \boldsymbol{\nabla}_\perp W_\parallel$ requires the relevant boundary term to vanish. Identical pipe reference fields provide one such case. Unequal pipes may require the boundary term even at $\beta = 1$: their transverse electric fields can differ although the pure-pipe Lorentz force vanishes. With Studio's X sign, the gradient term is $-\boldsymbol{\nabla}_\perp W_X^{\mathrm{Studio}}$.

To reconstruct a transverse component from this relation, integrate its right-hand side from a lag where the integration constant is known, usually a sufficiently remote vanishing-response limit. Do not impose a zero initial value inside a finite-beta space-charge signal. Keep identical source coordinates, port convention and effective witness nodes when taking transverse differences. Modern Studio reads the directly evaluated Y/Z wake artifacts; a longitudinal finite difference and Panofsky–Wenzel are an independent consistency check, not its ordinary way of manufacturing the displayed Y/Z curve.

1.7 Wake potential and Green function

For a linear, time-invariant structure with prescribed motion, define $g_\alpha(\boldsymbol{r}_q, \boldsymbol{r}_w, s; \beta)$ as the charge-normalized response to a longitudinal delta source, for $\alpha = \parallel, y, z$. Geometry, material response, velocity, transverse source model and reference planes are held fixed. The bunch wake is the convolution

$$W_\alpha(\boldsymbol{r}_q, \boldsymbol{r}_w, s; \beta) = \int_{-\infty}^{\infty} g_\alpha(\boldsymbol{r}_q, \boldsymbol{r}_w, s - s'; \beta)\lambda(s')\,\mathrm{d}s'. \tag{1.23}$$

Both g_α and W_α have units V/C. Some literature calls g the wake function and W the wake potential; specify this distinction rather than using "wake" to mean an electric-field snapshot. Introductory definitions and their frequency-domain interpretation are discussed in [1].

The integration is deliberately written as two-sided. A one-sided point-source kernel can be used when its physical causal-in-s property is established; it is not a general finite-beta assumption. A wake calculated with one bunch length is not the Green function simply because the bunch is short.

A longitudinal delta is not automatically a resolved 3D point charge

Dividing out the Gaussian spectrum removes longitudinal smoothing only. It does not undo the transverse beam profile or the grid representation of the source. At finite β, a coincident point source and witness can have a singular space-charge self-field. Specify a finite transverse distribution or another well-defined reference/regularization before interpreting such a limit.

Equation (1.23) requires the same linear system and reference definition throughout. Nonlinear material response, source evolution, or mixing wakes from different reference lengths invalidates a simple deconvolution.

1.8 Wake and coupling impedance

An impedance is a frequency-domain description of the same linear response. Its sign, transverse factor of i, Fourier exponential and velocity convention must be stated together. Two curves with different conventions cannot be compared by their real and imaginary parts alone.

1.8.1 Temporal frequency at the beam speed

Use a spatial Fourier transform and a normalized bunch spectrum

$$\widehat{g}_\alpha(k) = \int_{-\infty}^{\infty} g_\alpha(s) e^{-\mathrm{i}ks}\,\mathrm{d}s, \qquad \Lambda(k) = \int_{-\infty}^{\infty} \lambda(s) e^{-\mathrm{i}ks}\,\mathrm{d}s. \tag{1.24}$$

For temporal angular frequency $\omega = 2\pi f$, the delay is s/v, so $k = \omega/v$. With the decelerating longitudinal convention and the transverse $+\mathrm{i}$ convention,

$$Z_\parallel^{(v)}(\omega) = \frac{1}{v}\widehat{g}_\parallel(\omega/v), \tag{1.25}$$

$$\boldsymbol{Z}_\perp^{(v)}(\omega) = \frac{\mathrm{i}}{v}\widehat{\boldsymbol{g}}_\perp(\omega/v). \tag{1.26}$$

These physical trajectory impedances have units Ω. The transverse quantity is not yet divided by an offset. A dipole or quadrupole *coefficient* impedance has units Ω/m after the corresponding spatial derivative.

The inverse longitudinal relation is

$$g_\parallel(s) = \frac{1}{2\pi}\int_{-\infty}^{\infty} Z_\parallel^{(v)}(\omega) e^{\mathrm{i}\omega s/v}\,\mathrm{d}\omega. \tag{1.27}$$

For a transverse component the prefactor is $1/(2\pi\mathrm{i})$. These formulas use the full two-sided transform. Reality implies $Z_\parallel(-\omega) = Z_\parallel(\omega)^*$, whereas the chosen transverse factor gives $Z_y(-\omega) = -Z_y(\omega)^*$.

Where the endpoint-free Panofsky–Wenzel relation applies and Fourier end terms vanish, it implies, for nonzero ω,

$$\boldsymbol{Z}_\perp^{(v)}(\omega) = \frac{v}{\omega}\boldsymbol{\nabla}_\perp Z_\parallel^{(v)}(\omega). \tag{1.28}$$

This is still a derivative with source position fixed. If Eq. (1.22) has a nonzero electric-field boundary term, its transform must also be retained; the endpoint-free impedance relation is then insufficient.

1.8.2 Bunch response versus point-charge impedance

Fourier transforming the convolution gives

$$\widehat{W}_\alpha(k) = \widehat{g}_\alpha(k)\Lambda(k). \tag{1.29}$$

Define the bunch-weighted spectral response $Z^{(v)}_{\alpha,B}$ by replacing g with W in Eqs. (1.25)–(1.26). Then

$$Z^{(v)}_{\alpha,B}(\omega) = Z^{(v)}_\alpha(\omega)\Lambda(\omega/v), \qquad Z^{(v)}_\alpha = \frac{Z^{(v)}_{\alpha,B}}{\Lambda(\omega/v)}. \tag{1.30}$$

The first quantity is often called the bunch or beam impedance in a simulation plot. More precisely, it is the structure response *still filtered by the chosen bunch*. It is not a different material or geometry impedance.

For the temporal transform of current and decelerating voltage, $\widetilde{I} = Q\Lambda$ and $\widetilde{V}_{\text{loss}} = QZ^{(v)}_{\parallel}\Lambda$. Thus $\widetilde{V}_{\text{loss}}/\widetilde{I} = Z^{(v)}_{\parallel}$, wherever the denominator is usable. Dividing voltage by total charge rather than by the current spectrum retains the bunch factor. The transverse convention inserts the corresponding i in the voltage/current ratio.

For a Gaussian,

$$\Lambda(\omega/v) = \exp\left[-\frac{1}{2}\left(\frac{\omega\sigma_s}{v}\right)^2\right]. \tag{1.31}$$

Dividing by a very small Λ amplifies noise and discretization error. Deconvolution estimates only the excited, resolved part of the point-charge response. It cannot recover missing high-frequency information or remove a finite observation window exactly. Changing the bunch length changes W and Z_B; a trustworthy deconvolved Z should agree in the common reliable band, provided β and all physical/reference definitions are unchanged.

1.8.3 Studio's spatial c-referenced convention

Studio uses s/c to label its spatial spectrum at every beta, not s/v. Define $\omega_c = ck$ for its displayed frequency and write

$$Z_{\parallel,c}(\omega_c) = \frac{1}{c}\widehat{g}_\parallel(k), \qquad Z_{y,c}(\omega_c) = \frac{\mathrm{i}}{c}\widehat{g}_y(k). \tag{1.32}$$

In software the longitudinal minus sign is applied to the accelerating X wake before transforming. Gaussian deconvolution divides by $\exp[-(k\sigma_s)^2/2]$. Studio applies this transform only to a single-source wake, dividing physical voltage by that source's signed charge first. A composite voltage has no unique ordinary per-charge impedance; multiple-source voltages and observation derivatives are excluded from this analysis. Studio's separate dipole-coefficient artifacts support curve display and CSV export; they are not accepted by its ordinary impedance analysis. Transforming an exported coefficient requires preserving its extra inverse-length unit and the same bunch deconvolution and frequency convention.

Do not mix frequency conventions at finite beta

At $\beta = 1$, Z_c and the temporal impedance $Z^{(v)}$ are identical. At $\beta < 1$, the same spatial Fourier sample corresponds to

$$f_v = \beta f_c, \qquad Z^{(v)}_\alpha(2\pi f_v) = \frac{Z_{\alpha,c}(2\pi f_c)}{\beta}.$$

This applies to bunch-weighted and deconvolved spectra, and to offset-normalized coefficients. Preserve the Studio axes when reporting Z_c. If a beam-dynamics comparison requires the

temporal v convention, convert both the frequency and the complex amplitude and label them explicitly. Do not rescale the wake itself.

For example, at $\beta = 0.8$ a Studio sample at $f_c = 10$ GHz corresponds to $f_v = 8$ GHz. An amplitude $Z_c = 100 + 20\mathrm{i}\ \Omega$ at that sample corresponds to $Z^{(v)} = 125 + 25\mathrm{i}\ \Omega$. This is an arithmetic illustration of the convention, not a simulated impedance result.

1.8.4 A defensible spectral workflow

In Results, select Analysis: Impedance, the recorded Integration, Component and Trajectory, then Load / Analyze.

1. Keep Gaussian bunch deconvolution off to inspect the bunch-weighted response. Turn it on to estimate the longitudinal-delta-source impedance.
2. Inspect real and imaginary parts as well as magnitude. Compare phases only with identical lag origins and Fourier conventions.
3. Compare Rectangular with Hann tail (last 10%) if a tail remains at the record end. A window changes the transformed signal; it does not supply unrecorded tail data.
4. Choose a frequency band supported by the bunch spectrum, spatial/time discretization and actual wake length. Studio's automatic band uses a Gaussian spectrum threshold of 0.05; an explicit maximum does not remove the reliability checks.
5. For finite beta, state whether the published spectrum uses Studio's c reference or a converted temporal v convention.

For N samples separated by Δs, Studio's sampling scales are

$$f_{N,c} = \frac{c}{2\Delta s}, \qquad \Delta f_{\mathrm{record,c}} \simeq \frac{c}{N\Delta s}. \tag{1.33}$$

Multiply both by β for the temporal v convention. Nyquist is only a sampling bound, not a guarantee of field accuracy. Zero padding decreases the plotted bin spacing but does not improve record-length resolution. A finite-beta signal or a Gaussian wake may require data at negative s; deleting it changes both the spectrum and its phase.

1.9 Loss factor, kick factor and RMS spread

A wake curve describes individual longitudinal slices. A loss or kick factor summarizes its effect over the bunch. For the normalized longitudinal profile of Eq. (1.3), define

$$\langle f \rangle_\lambda = \int_{-\infty}^{\infty} f(s)\lambda(s)\,\mathrm{d}s, \qquad \langle 1 \rangle_\lambda = 1. \tag{1.34}$$

Source and witness transverse positions, beta and integration convention are held fixed during this average. An unweighted average of plotted sample values is not a bunch average.

1.9.1 Mean loss and signed kick

With the decelerating longitudinal sign used in this guide,

$$\boxed{k_\parallel = \langle W_\parallel \rangle_\lambda, \qquad k_y = \langle W_y \rangle_\lambda, \qquad k_z = \langle W_z \rangle_\lambda.} \tag{1.35}$$

$k_\parallel$ is the longitudinal *loss factor*. k_y, k_z are signed *kick factors on the stated witness trajectory*. All three have units V/C, commonly V/pC. A positive loss factor means that the bunch gives up energy in the stated self-wake problem. A kick factor can have either sign, or vanish by symmetry while individual slices still receive kicks.

On the bunch trajectory, with the same transverse source model and consistent physical integration/reference definition, the energy transferred from the whole bunch is

$$U_{\text{beam loss}} = Q^2 k_\parallel. \tag{1.36}$$

Use SI units in this equation. The corresponding mean longitudinal loss voltage is $Qk_\parallel$. The mean kick voltage is Qk_y, and a witness particle of charge q and longitudinal momentum p has

$$\langle \Delta p_y \rangle = \frac{qQ}{v} k_y, \qquad \langle \theta_y \rangle \simeq \frac{qQ}{pv} k_y. \tag{1.37}$$

An average calculated on an additional witness trajectory is its response to the source bunch; it is not the energy lost by the driving bunch. At finite beta, changing the included pipe-space-charge length changes the definition being averaged. Do not identify that change with additional material heating.

The word *kick factor* is also used for an offset-normalized dipole coefficient. Keep that quantity distinct from the physical trajectory factor. Using the expansion in Section 1.10,

$$k_y = k_{M_y} + y_q k_{D_y} + y_w k_{K_y} + \cdots, \quad k_{M_y} = \langle M_y \rangle_\lambda, \quad k_{D_y} = \langle D_y \rangle_\lambda, \quad k_{K_y} = \langle K_y \rangle_\lambda. \tag{1.38}$$

k_{D_y} and k_{K_y} have units V/pC/m. In the one-run PEC dipole protocol, $k_{D_y} \simeq \langle W_y^{\text{PEC}} \rangle_\lambda / (2a)$. Do not divide a general Y curve by the source offset before separating any monopole and quadrupole contributions.

1.9.2 RMS value versus RMS spread

For $\alpha \in \{\parallel, y, z\}$, distinguish two nonnegative quantities:

$$\begin{aligned} k_{\alpha,\text{rms}} &= \sqrt{\langle W_\alpha^2 \rangle_\lambda}, && \text{RMS value about zero,} \\ \sigma_{k_\alpha} &= \sqrt{\langle (W_\alpha - k_\alpha)^2 \rangle_\lambda} = \sqrt{k_{\alpha,\text{rms}}^2 - k_\alpha^2}, && \text{RMS spread about the mean.} \end{aligned} \tag{1.39}$$

The factors and their spreads have the same units. For a constant wake W_0, the mean is W_0, the RMS value is $|W_0|$, and the RMS spread is zero. A zero mean does not imply a zero spread.

Studio reports RMS spread, not RMS about zero

The table entry named Loss voltage RMS spread or Kick voltage Y/Z RMS spread is the standard deviation over the Gaussian bunch. It is not $\sqrt{\langle V^2 \rangle_\lambda}$ and is not a statistical uncertainty of the numerical calculation.

Wake-induced slice spreads are

$$\sigma_{V,\alpha} = |Q| \sigma_{k_\alpha}, \qquad \sigma_{\Delta\mathcal{E}} = |qQ| \sigma_{k_\parallel}, \qquad \sigma_{\theta_y} \simeq \frac{|qQ|}{pv} \sigma_{k_y}. \tag{1.40}$$

These expressions describe the wake contribution alone. They do not include the incident energy spread, its correlation with the wake, intrinsic angular spread, or transverse phase-space correlations. In particular, a kick RMS is not an emittance.

The same averaging and RMS definitions apply to a charge-normalized $D_y(s)$ or $K_y(s)$, giving V/pC/m. Studio displays physical responses in V/m: multiply the pair's saved D_y by one source

charge magnitude q, and differentiate physical observation voltage for the quadrupole/gradient view. Extract the physical curve first, then compute its signed mean and centred RMS spread. For both views all components retain the plotted sign in the mean; only the ordinary longitudinal Loss voltage uses the negative-X convention. The common normalized Gaussian profile supplies the weights even when the source pair's total charge and total current vanish. Both statistics have V/m units; the unrooted second central moment would have units $(\mathrm{V/m})^2$. The RMS of a sum or difference cannot in general be obtained by adding or subtracting the separate RMS values: correlations between the curves matter.

1.9.3 Loss and kick voltages in Studio

For the charge-normalized X companion, $W_X^{\mathrm{Studio}} = -W_{\parallel}$. For any displayed component $\alpha \in \{X, y, z\}$, the slice voltage and its bunch mean are

$$\begin{aligned} V_\alpha(s) &= Q W_\alpha^{\mathrm{Studio}}(s), \\ \overline{V}_\alpha &= \int_{-\infty}^{\infty} V_\alpha(s)\lambda(s)\,\mathrm{d}s = Q\int_{-\infty}^{\infty} W_\alpha^{\mathrm{Studio}}(s)\lambda(s)\,\mathrm{d}s. \end{aligned} \tag{1.41}$$

The normalized Gaussian $\lambda(s)$ is given in Eq. (1.3); $W_{y,z}^{\mathrm{Studio}} = W_{y,z}$. The physical RMS spread is

$$\sigma_{V,\alpha} = \left[\int_{-\infty}^{\infty} \big(V_\alpha(s) - \overline{V}_\alpha\big)^2 \lambda(s)\,\mathrm{d}s\right]^{1/2}. \tag{1.42}$$

Studio reports voltage metrics, not charge-normalized factors:

Studio metric	Definition
Loss voltage	$-\overline{V}_X$
Kick voltage Y/Z	$\overline{V}_y$ or $\overline{V}_z$
Loss/Kick voltage RMS spread	$\sigma_{V,\alpha}$ from Eq. (1.42)
Source charge	Signed bunch charge Q, displayed in pC

Every voltage row is in V and already includes charge: a wake in V/pC multiplied by the charge in pC gives V. To obtain the corresponding mean factor in V/pC, divide the displayed signed mean voltage by the signed charge in pC; for its RMS spread, divide by the absolute charge in pC. No source-offset division is applied to these ordinary curves. These definitions refer to the complete bunch profile and to the charge and RMS bunch length of the selected calculation.

For a composite excitation, the saved ordinary curves already contain V_α^{Σ} in volts. Apply the same Gaussian mean and RMS directly, without multiplying by a charge. The current overlay is the total signed current; it vanishes for a synchronous equal opposite pair. A value at one witness is not the total energy lost by all sources. Separate dipole and observation-derivative coefficient curves have no voltage metrics or current overlay. Their normalization is included in the displayed units; see Section 1.12.

1.9.4 Connection to impedance

For the same drive and witness longitudinal profile, convolution and the Fourier definitions in Section 1.8 give

$$\begin{aligned} k_{\parallel} &= \frac{1}{\pi}\int_0^{\infty} \mathrm{Re}\, Z_{\parallel}^{(v)}(\omega)|\Lambda(\omega/v)|^2\,\mathrm{d}\omega, \\ k_y &= \frac{1}{\pi}\int_0^{\infty} \mathrm{Im}\, Z_y^{(v)}(\omega)|\Lambda(\omega/v)|^2\,\mathrm{d}\omega. \end{aligned} \tag{1.43}$$

The second sign follows specifically from this guide's transverse $+\mathrm{i}$ convention. These are point-source impedances, so the squared bunch spectrum is applied once; a bunch-filtered impedance must not be treated as an already deconvolved one. With Studio's spatial convention, replace v by c and integrate over ω_c consistently. The resulting factor is unchanged. These identities assume convergent integrals and the same physical reference definition. They are not permission to extend a noisy numerical spectrum beyond its reliable band.

1.10 Taylor expansion and the meaning of transverse components

Choose a reference orbit at $y = z = 0$. If the physical axis has another coordinate, all offsets below are measured from that axis. A transverse wake depends independently on the source and witness positions: $y_q = y_{\text{source}}$ and $y_w = y_{\text{witness}}$. The four-coordinate Taylor construction is discussed for the longitudinal wake in Appendix B of [5]; here we apply it directly to W_y. For one component, at fixed s, beta and integration convention, let $\boldsymbol{u} = (y_q, z_q, y_w, z_w)$. Where the response is smooth,

$$W_y(\boldsymbol{u}, s) = W_y(\mathbf{0}, s) + \sum_i A_i(s) u_i + \frac{1}{2} \sum_{i,j} B_{ij}(s) u_i u_j + \cdots . \tag{1.44}$$

Here $A_i = \partial_{u_i} W_y|_{\boldsymbol{u}=\mathbf{0}}$ and $B_{ij} = \partial_{u_i}\partial_{u_j} W_y\big|_{\boldsymbol{u}=\mathbf{0}}$, with $i, j = 1, \ldots, 4$. No transverse harmonicity is assumed in this expansion. To first order this can be written

$$W_y = M_y + D_{yy} y_q + D_{yz} z_q + K_{yy} y_w + K_{yz} z_w + O(\|\boldsymbol{u}\|^2). \tag{1.45}$$

The constant M_y is the transverse monopole contribution. The D coefficients describe source displacement (driving or dipolar terms), and the K coefficients describe witness displacement (detuning or quadrupolar terms).

For the co-planar scan $z_q = z_w = 0$, the three terms are explicitly

$$W_y(y_q, y_w, s) = \underbrace{M_y(s)}_{\text{monopole}} + \underbrace{y_q D_y(s)}_{\text{dipole}} + \underbrace{y_w K_y(s)}_{\text{quadrupole}} + \cdots . \tag{1.46}$$

Their definitions, with all other offsets fixed at zero, are

$$\begin{aligned} M_y(s) &= W_y(0, 0, s), \\ D_y(s) &= \left.\frac{\partial W_y}{\partial y_q}\right|_{y_q = y_w = 0}, \qquad \text{source derivative}, \\ K_y(s) &= \left.\frac{\partial W_y}{\partial y_w}\right|_{y_q = y_w = 0}, \qquad \text{witness derivative}. \end{aligned} \tag{1.47}$$

Here $D_y = D_{yy}$ and $K_y = K_{yy}$. The same construction applies to W_z after exchanging y and z. Cross-plane coefficients need independent scans of the other coordinate; do not assume they vanish in an asymmetric structure.

Quantity	Meaning at the reference orbit	Units
M_y	$W_y(0, 0)$; a kick with both trajectories centred	V/pC
D_y	$\partial_{y_q} W_y(0, 0)$; response to source displacement	V/pC/m
K_y	$\partial_{y_w} W_y(0, 0)$; response to witness displacement	V/pC/m
aD_y	Dipolar contribution at source offset a	V/pC
bK_y	Quadrupolar contribution at witness offset b	V/pC

"Quadrupolar" does not mean the term is quadratic in witness offset in the *transverse kick*. The force is linear in that offset; a corresponding longitudinal potential can be quadratic. With an endpoint-free Panofsky–Wenzel relation, a local longitudinal expansion containing $A(s)y_q y_w + B(s)y_w^2/2$ gives $\partial_s D_y = A$ and $\partial_s K_y = B$. Boundary contributions must be added when Eq. (1.22) requires them.

The bunch trajectory alone does not separate dipole and quadrupole

Moving source and witness together gives $W_y(a, a) = M_y + a(D_y + K_y) + \cdots$. Dividing that single full-domain curve by a generally mixes the two slopes and, if present, M_y/a. Hold the witness fixed when extracting D_y and hold the source fixed when extracting K_y.

In a structure with y-reflection symmetry, $W_y(-y_q, -y_w) = -W_y(y_q, y_w)$ in this co-planar scan. Thus $M_y = 0$ and even total powers vanish. Two reflection planes also remove the cross-plane linear terms at their intersection. They do not imply cylindrical symmetry: D_y and D_z need not be equal.

In the special ultrarelativistic, source-free witness region with a harmonic longitudinal wake and an endpoint-free relation, the detuning matrix is trace-free: $K_y + K_z = 0$ after fixing the integration constant. Do not impose that equality for arbitrary finite-beta complete-field wakes; space charge and reference boundary terms can invalidate the assumptions. Likewise, a rotationally symmetric ultrarelativistic example is not evidence that every structure has zero K_y.

1.11 Full-structure extraction in Studio

Studio displays ordinary X/Y/Z wakes on explicit observation trajectories and separate dipole-coefficient artifacts for an equal opposite-charge pair. Such a pair gives the source difference in one field calculation, as described below. Independent single-source runs remain useful for comparison and for extracting the monopole and witness derivative. The Compare action overlays curves; it does not subtract them. Export those curves with Export CSV for the finite-difference arithmetic in this subsection.

Choose small, resolved positive offsets a and b. Use the same beta, charge, RMS bunch length, geometry, materials, grid, integration planes, recording and Direct/Indirect convention in all runs. Keep distinct case/run identities so a source displacement cannot overwrite the reference result.

Run	Source y_q	Required witness y_w	Absolute observation rows
Centre	0	$0, +b, -b$	$(0,0)$, $(+b,0)$, $(-b,0)$
Plus	$+a$	0	$(0,0)$
Minus	$-a$	0	$(0,0)$

All z coordinates are zero in this table. Every observation is authored explicitly, including $(0,0)$. Source and observation lists are independent; the Plus and Minus runs use the same on-axis witness.

Read Results > Wake potential > Component Y for each specified trajectory. The centred-difference estimates are

$$M_y(s) = W_y(0, 0, s), \tag{1.48}$$

$$D_y(s) \simeq \frac{W_y(+a, 0, s) - W_y(-a, 0, s)}{2a}, \tag{1.49}$$

$$K_y(s) \simeq \frac{W_y(0, +b, s) - W_y(0, -b, s)}{2b}. \tag{1.50}$$

For a smooth response, the centred derivative errors are $O(a^2)$ and $O(b^2)$. The Centre run supplies both M_y and K_y, so three full-domain calculations suffice for this symmetric finite-difference protocol. A two-run one-sided drive difference is possible but generally has a larger $O(a)$ truncation error.

1.11.1 Effective positions and units

Studio/Solver select nearest physical grid nodes for the source and witness trajectories. Read the effective coordinates, not just the authored decimal offsets. If two requested offsets select the same node, no derivative has been resolved. Duplicate nodes within either list are rejected; a source and an observation may coincide because the two lists serve different purposes.

Keep the effective witness identical in the Plus and Minus cases. Use the actual source separation in the denominator. If the two effective source positions are not symmetric about the chosen axis, the difference estimates a derivative at their midpoint and is not the stated centred formula about zero. The same caution applies to witness differences. Refinement or a nonuniform polynomial fit is preferable to silently pretending the nodes are symmetric.

For coefficients in V/pC/m, convert a, b to metres before division. Dividing a V/pC curve by an offset in millimetres initially gives V/pC/mm; multiply that numerical value by 10^3 to express it per metre. No extra charge normalization is applied to a wake curve after conversion to per-charge units.

To verify linearity, compare several resolved offsets and a refined mesh. An offset too large includes higher multipoles; one too small amplifies roundoff, grid error and cancellation. At finite beta, retain the same pipe-reference convention across the scan: the coefficients then include the specified finite-length space-charge contribution. Subtracting a separate physical reference defines a different coefficient and must be documented.

1.12 Synchronous opposite-charge pair: one-run dipole extraction

No reflection symmetry is required for this method. Use the complete structure and put equal opposite charges at two resolved source positions:

$$Q_+ = q, \quad Q_- = -q, \quad q > 0, \qquad \boldsymbol{r}_\pm = \boldsymbol{r}_0 \pm \frac{d}{2}\boldsymbol{e}, \quad |\boldsymbol{e}| = 1. \tag{1.51}$$

Here d is the *full* transverse separation, and $\boldsymbol{e}$ points from the negative source to the positive source. Keep every witness fixed while changing the source separation. Equation (1.9) gives

$$V_\alpha^\Sigma = q[W_\alpha(\boldsymbol{r}_+, \boldsymbol{r}_w, s) - W_\alpha(\boldsymbol{r}_-, \boldsymbol{r}_w, s)]. \tag{1.52}$$

The fields, initial/reference fields and currents are combined within one calculation. The part independent of source displacement cancels before the final wake is written, which avoids subtracting two separately calculated large background wakes. This can improve accuracy when the dipole signal is small.

For a smooth source response about $\boldsymbol{r}_0$,

$$\frac{V_\alpha^\Sigma}{qd} = (\boldsymbol{e} \cdot \nabla_{\boldsymbol{r}_q}) W_\alpha(\boldsymbol{r}_0, \boldsymbol{r}_w, s) + O(d^2). \tag{1.53}$$

All even powers of displacement cancel. Higher odd powers remain at finite separation, so the pair yields a directional source derivative in the small, resolved-separation limit. It is not an exact decomposition into azimuthal modes of an asymmetric structure. At a centred witness, a y-directed

pair gives D_{yy} from component Y and D_{zy} from component Z; a z-directed pair gives the other source-direction column. A displaced witness gives the derivative at that witness, including its higher-order position dependence.

1.12.1 Automatic moment normalization and saved results

Solver uses the *effective grid positions* to form the signed dipole moment

$$\boldsymbol{p} = q(\boldsymbol{r}_+ - \boldsymbol{r}_-), \qquad D_\alpha^{\mathrm{pair}} = \frac{V_\alpha^\Sigma}{|\boldsymbol{p}|}. \tag{1.54}$$

$\boldsymbol{p}$ has units C m. Its direction specifies the source derivative, and D_α^{pair} is saved in V/pC/m. There is no additional reference charge. Never divide by the zero net charge, and never divide a saved coefficient by the separation again. If snapping makes the pair asymmetric about the requested centre, the derivative is centred on the effective midpoint. Use a sufficiently resolved symmetric pair when the derivative is required about a particular reference orbit.

An exactly equal opposite-charge two-source list produces both the ordinary voltage curves in V and separate files ending in `_dipole.bin`, including available Direct/Indirect X/Y/Z components. In Studio, choose Result Analysis, then Analysis → Dipole coefficient, and select the integration, component and observation. Studio multiplies the matching saved coefficient by $q > 0$, the actual magnitude of one source charge, and displays $D_{V,\alpha} = qD_\alpha^{\mathrm{pair}} = V_\alpha^\Sigma/d$ in V/m. The signed mean, centred RMS spread and CSV use the same physical units. The saved coefficient keeps its original units; there is no normalization switch. The option is unavailable, with an explanatory tooltip, if this run and selection have no coefficient artifact; selecting it never generates a missing result. The result metadata `excitation.json` records physical charges, requested/effective positions, moment and normalization. General source lists retain composite voltage curves; an automatic dipole coefficient is defined only for this two-source pair.

Contribution and coefficient are different outputs

For sources at $y = \pm a$, the physical pair voltage divided by $2q$ is the exact odd-in-source contribution at offset a. The separate coefficient file already contains $V_\alpha^\Sigma/(2qa)$, using the effective separation in metres. In the linear regime these quantities are aD_α in V/pC and D_α in V/pC/m, respectively.

1.12.2 Studio setup and convergence

1. Use a transverse domain containing both sources and the actual structure. Do not impose a mirror cut unless its geometry and excitation parity are valid.
2. In Excitation Source, create a dipole pair. Enter its centre, full separation along Y/Z and charge magnitude. Both sources use the common Gaussian profile, RMS length and beta. Enter the physical charges only.
3. In Wake & Outputs, enter absolute observation coordinates. The pair preset may add its centre as an observation; existing observations stay fixed.
4. Check the effective source positions and moment. Build the mesh and run, then inspect both the physical voltage and the separate dipole coefficient.
5. Compare coefficients at several resolved separations and on a finer mesh, keeping beta, bunch length, witnesses and integration convention unchanged. Small separation alone does not establish convergence if grid or roundoff error dominates. The same stationary finite-beta and moving beta-one limits apply as for a single source.

Studio supports viewing and CSV export of coefficient curves. Voltage metrics apply to physical voltage curves. Spectral analysis requires a single-source normalized wake; no arbitrary charge is assigned to a multi-source voltage. A dipole impedance requires a consistently normalized transform of the exported coefficient; the present coefficient viewer does not perform it.

1.13 Derivative with respect to observation position

With one fixed source of charge q, two recorded observations at $\boldsymbol{r}_a = (y_a, z_0)$ and $\boldsymbol{r}_b = (y_b, z_0)$ give

$$K_{\alpha y}^{\mathrm{FD}}(s) = \frac{V_\alpha(\boldsymbol{r}_b, s) - V_\alpha(\boldsymbol{r}_a, s)}{q(y_b - y_a)}, \qquad \alpha \in \{x, y, z\}. \tag{1.55}$$

For small resolved separation this estimates the derivative with respect to witness position. For transverse components Y/Z this is conventionally the detuning or quadrupolar coefficient. For component X it is the transverse gradient of the longitudinal wake. Keep the source fixed. A central pair gives second-order truncation error about its midpoint for a smooth response. A one-sided pair estimates the derivative at an endpoint only to first order in the separation.

In an asymmetric structure, V_α on the source trajectory may be nonzero. Dividing a single off-source curve by its distance from the source would retain that constant kick. The difference above removes it. Simple division is justified only when the reference value vanishes and the response is in its linear regime. This witness derivative is different from the driving coefficient obtained by varying the source position.

In Studio choose Wake gradient (quadrupole) in Result Analysis, a component X, Y or Z, a target Trajectory, a Reference trajectory, and Derivative axis Y or Z. This is the second Analysis option, after Wake potential and before Dipole coefficient. The two sampled positions must differ only along that axis. Studio subtracts reference from target and divides by the signed effective separation stored in the result, converted to metres. Reversing both observations therefore preserves the coefficient. A zero separation is rejected. Authored coordinates are not substituted for the sampled positions.

The Current on plot selector appears only with multiple sources and shows an individual source, their total, or all source currents in amperes, without differentiation or separation scaling. A single source current is shown automatically. A pair initially shows its positive source. Its total is exactly zero, while individual currents $I_i = q_i v \lambda$ remain nonzero. The common Gaussian weight for wake metrics is independent of the selected current curve.

For one or several sources the displayed physical voltage difference is divided only by separation and is in V/m. The signed mean and centred RMS spread also use V/m, as does CSV. No source-charge sum or invented normalization is used. Check convergence with smaller resolved offsets and a refined mesh; exact separation does not remove finite-difference or field-discretization error.

1.14 One reflection plane: the one-run PEC dipole method

Assume the *complete physical structure and its materials* are symmetric about $y = 0$. Simulate only $y \geq 0$. The cut is a computational symmetry plane, not an extra metal sheet inserted into the physical aperture.

For the normal component W_y, a PEC plane enforces odd source-potential parity; a PMC plane enforces even parity. The resulting component parity is

Cut at $y = 0$	PEC	PMC
Electric field	Tangential $E_x, E_z = 0$	Normal $E_y = 0$
Wake on cut	Tangential $W_X, W_z = 0$; normal W_y can remain	Normal $W_y = 0$; tangential components can remain
Excitation sector	Odd under source reflection	Even under source reflection

Swapping PEC and PMC selects a different symmetry sector of the full problem. It is not a generic numerical stabilization choice.

1.14.1 Dipolar contribution: one calculation, then divide by two

Put the source at $y_q = a > 0$, use PEC on $y = 0$, and record the Y wake at a fixed witness $y_w = b$. Linearity and reflection give

$$W_y^{\mathrm{PEC}}(a, b, s) = W_y(a, b, s) - W_y(-a, b, s). \tag{1.56}$$

This identity can be understood as the source and its opposite-sign reflected image. No second source is authored in Studio: the PEC boundary imposes the required parity. Here every term is normalized by the charge of the authored source, not by a fictitious net charge of the odd pair (which would be zero).

The factor of two

The odd-in-source contribution of one displaced source is

$$W_y^{\mathrm{odd}}(a, b, s) = \frac{W_y^{\mathrm{PEC}}(a, b, s)}{2}.$$

In the linear regime this is the dipolar contribution $aD_y(s)$. To obtain the coefficient per unit source offset, divide by a as well:

$$D_y(s) \simeq \frac{W_y^{\mathrm{PEC}}(a, 0, s)}{2a}.$$

One field calculation is sufficient. Dividing by two gives the contribution; dividing by $2a$ gives the coefficient. These are different quantities and units.

The odd-source identity is exact for the symmetric linear problem. Calling the entire odd part a *linear dipole* is an approximation at finite offset; higher odd source orders survive. With witness on the symmetry axis the derivative estimate has $O(a^2)$ error. The centred on-axis witness also avoids unnecessary witness-offset dependence in the extracted coefficient.

1. Keep the real geometry and materials in the half-domain $y \geq 0$ and set the cut face to PEC. Retain the correct physical outer boundaries.
2. Set the source to $y_q = a$, keeping its other coordinate fixed on the intended reference orbit. Do not put the odd-sector source exactly on the PEC cut.
3. In Wake & Outputs, enter the absolute observation $(y, z) = (0, 0)$. Verify the effective witness is on the cut and the source offset is resolved.
4. Run once, select the desired Integration, Component Y and that observation, then export its curve.

5. Divide by two for the dipolar contribution at offset a; divide by $2a$ for D_y. A single authored source with a PEC image does not create the automatic two-source coefficient file; perform this normalization on export.

As an arithmetic example only, if $a = 0.3$ mm and one sample of W_y^{PEC} is 0.6 V/pC, the odd contribution is 0.3 V/pC and the linear dipole coefficient is 1000 V/pC/m. No new simulation is implied by these numbers.

1.14.2 Centred excitation and quadrupolar coefficient

For the even sector use PMC on $y = 0$ and place the source on that plane. Then $M_y = 0$ and the witness-axis wake vanishes by parity, but a witness at $b > 0$ can have a nonzero Y wake. The reduced calculation represents the centred physical source; its normalization on a PMC source plane is already handled by Solver.

$$K_y(s) \simeq \frac{W_y^{\mathrm{PMC}}(0, b, s)}{b}. \tag{1.57}$$

There is *no extra division by two* in this formula. The apparently one-sided sample is a centred derivative by odd witness parity, and has $O(b^2)$ error for a smooth symmetric response. In Studio use the centred source, an explicit observation $(y, z) = (b, 0)$, and the Y curve. An additional explicit $(0, 0)$ observation supplies the longitudinal wake on the source trajectory.

1.14.3 If the only symmetry plane is the other transverse plane

If the structure has only $z = 0$ reflection symmetry, that plane does not force M_y to vanish and cannot by itself isolate the y-driving dipole with PEC. Keep the source and the relevant witnesses at $z = 0$, use PMC there for their even z sector, and perform the full y-scan of Section 1.11 in the reduced domain. Conversely, a PEC z cut gives the one-run construction for the normal Z dipole. Always associate the odd cut with the source-offset direction of the dipole being extracted.

1.15 Two reflection planes: quarter-domain calculations

Assume both $y = 0$ and $z = 0$ are genuine reflection planes of geometry and materials, and retain $y, z \geq 0$. The following small set of calculations provides the uncoupled on-axis coefficients.

Target	y cut	z cut	Source (y_q, z_q)	Witness and reduction
Longitudinal; K_y, K_z	PMC	PMC	$(0, 0)$	$(0, 0)$ for longitudinal; $(b, 0)$: W_y/b; $(0, b_z)$: W_z/b_z
D_y	PEC	PMC	$(a, 0)$	$(0, 0)$: $W_y/(2a)$
D_z	PMC	PEC	$(0, a_z)$	$(0, 0)$: $W_z/(2a_z)$

The transverse monopoles vanish at the intersection by symmetry. The centred PMC/PMC run supplies both quadrupolar coefficients using explicit off-axis witness trajectories. The y-dipole and z-dipole each need one additional odd-sector run. To study only the Y component, the centred run and the PEC/PMC run suffice.

Two symmetry planes do not imply a factor of four for the dipole

For the Y dipole the source is displaced from the PEC y cut but lies *on* the PMC z cut. Solver already applies the on-plane source weight for that PMC plane. The odd y pair still produces the factor two, not four: $W_y^{\mathrm{odd}} = W_y^{\mathrm{PEC/PMC}}/2$ and $D_y \simeq W_y^{\mathrm{PEC/PMC}}/(2a)$ in the linear regime. For a source on both PMC cuts, keep the authored full-source charge and do not multiply or divide its per-charge analysis curve by four.

These normalization statements require the stated on-plane placement. An off-plane source with PMC mirrors represents same-sign reflected sources, not automatically one displaced bunch in the complete device. Arbitrarily moving that source off the second cut changes the excitation and its normalization. The power-loss monitor's separate symmetry multiplier is an extensive-energy scaling control; do not apply it to per-charge analysis curves.

PEC/PEC instead selects a response odd in both transverse source coordinates. It is not the ordinary on-axis Y dipole, Z dipole or centred quadrupole protocol. Do not select it merely because the model is a quarter of the complete geometry.

1.16 Checks before publishing a component or impedance

1. State β, σ_s, source and witness coordinates, the physical geometry/material symmetry, and whether the reported wake includes finite-length pipe space charge. For finite beta, give the actual integration planes.
2. Identify the result as M_y, aD_y, D_y, bK_y, K_y or the complete W_y. Do not label a raw Y curve as a dipole coefficient solely because its source is displaced.
3. Check source and witness offsets independently and use effective grid coordinates. Preserve a common lag grid before subtracting exported curves; never subtract by row number when their s coordinates differ.
4. For an ordinary odd-sector wake divide by two once, then by the offset only if a coefficient is desired. A pair's saved dipole coefficient is already normalized. No additional PMC on-plane factor is needed in the quarter-domain protocols above.
5. Keep the same window, bunch deconvolution, frequency convention and usable frequency band when transforming or comparing coefficients. Linear subtraction and normalization commute with a common transform; selecting magnitudes first does not. For example, subtract complex spectra, not $|Z_+| - |Z_-|$.
6. Verify convergence of the quantity actually reported. Use the relevant Panofsky–Wenzel relation with its boundary term as an independent check, not as permission to remove a physically present space-charge contribution.

All symmetry identities and source/witness derivative definitions above also apply at finite beta, for a linear reflection-symmetric problem and a consistently transformed reference definition. They do not require the pure-pipe force to vanish. The numerical accuracy of a specific calculation must still be established.

1.17 References and software scope

The derivations use the axes and explicit signs defined in this guide; referenced papers may use another longitudinal axis or wake sign. Two-port wake formulas are expressed in the owner's

complete-centre/local-reference-tail convention, without introducing a separate corrected longitudinal wake potential.

[1] M. Dohlus and R. Wanzenberg, *An Introduction to Wake Fields and Impedances*, CERN Yellow Reports: School Proceedings 3 (2017), p. 15. doi:10.23730/CYRSP-2017-003.15. General physical background and definitions.

[2] W. K. H. Panofsky and W. A. Wenzel, *Some Considerations Concerning the Transverse Deflection of Charged Particles in Radio-Frequency Fields*, Review of Scientific Instruments 27, 967 (1956). doi:10.1063/1.1715427.

[3] I. Zagorodnov, *Indirect methods for wake potential integration*, Physical Review Special Topics – Accelerators and Beams 9, 102002 (2006). doi:10.1103/PhysRevSTAB.9.102002; arXiv:physics/0606049.

[4] I. Zagorodnov and D. Bazyl, *Indirect Integration of Longitudinal and Transverse Wake Potentials for Unequal Beam Pipes and Arbitrary Beam Velocity*. Submitted to arXiv and PRAB.

[5] I. Zagorodnov, M. Dohlus and S. Tomin, *Accelerator beam dynamics at the European X-ray Free Electron Laser*, Physical Review Accelerators and Beams 22, 024401 (2019), Appendix B, p. 024401-16. doi:10.1103/PhysRevAccelBeams.22.024401. Four-coordinate source/witness Taylor expansion. This guide writes the Y-component terms explicitly in Studio axes, without assuming harmonicity for general finite-beta complete-field wakes.

[6] *EMUZAG Studio User Guide*, EMUZAG 1.0.2, edition 1.21. The package manifest identifies the exact Solver and Studio source revisions. The current result controls read X/Y/Z artifacts on selected trajectories, evaluate Gaussian-weighted voltage metrics, display material-loss CSVs and read the two-port multimode S-matrix. Saved dipole coefficients use V/pC/m; Studio displays charge-scaled V/m curves, means, RMS spreads and CSV. Ordinary-wake impedance uses the spatial c-referenced convention.

The Studio workflows were checked against the current source and public artifact contract. The pair coefficient is computed by Solver from the synchronous signed sources; it is distinct from the removed paired-run postprocessor and from manual witness-derivative extraction. Where software labels and physical notation differ, the mapping is stated explicitly.

Numerical qualification is recorded separately from this derivation and workflow guide. The arithmetic examples in the text demonstrate units and normalization, not new simulation measurements. A source/evidence ledger and a reproducible LaTeX build accompany the document.

CHAPTER 2

Material Losses and Power Accounting

This chapter separates energy extracted from the beam from heat deposited in particular material mechanisms. Solver advances the material states and accumulates their passive energy losses. Studio reads and presents those recorded quantities; it does not infer heat from a field screenshot or from the longitudinal wake. The software labels and units below follow the current User Guide and result reader [6]; the accompanying source ledger also identifies the Solver accumulation and output routines.

2.1 Beam energy loss is not necessarily material heat

In a consistent physical energy balance, work supplied by the beam can become material heat, electromagnetic or material stored energy, or energy leaving through the boundary:

$$U_{\text{beam loss}} = U_{\text{heat}} + \Delta U_{\text{stored}} + U_{\text{out}}. \tag{2.1}$$

Here U_{out} is the signed net outgoing energy, with any imposed incoming field work accounted for consistently. Stored energy includes dispersive and wall internal states when present. In a truncated simulation, numerical PML absorption accounts for outgoing energy; it is not heat in a physical wall.

Consequently, the wake loss factor $k_{\parallel}$ of Section 1.9 need not equal a material cumulative-loss curve. A PEC cavity can have a nonzero beam loss factor with zero wall heating: energy may remain in its modes or radiate into the pipes. Equality requires compatible source/reference definitions, the same physical domain and symmetry normalization, all relevant dissipative channels, and an appropriate final-time energy balance. At finite beta, complete-field pipe space charge and unequal-pipe references require particular care.

2.2 Material interfaces and mesh refinement

A jump in positive, isotropic, frequency-independent ε_r or μ_r changes the field even when neither material absorbs energy. Solver automatically corrects certified single planar interfaces parallel to the Cartesian grid axes, including interfaces between grid lines. Each electric or magnetic component uses its own staggered averaging region. There is no project switch to enable it.

Supported aligned interfaces with volume conductivity or electric/magnetic Lorentz dispersion use the complete constitutive response: arithmetic averaging along the interface and passive series composition across it. Averaging only the high-frequency constants does not reproduce that dispersive normal response. The same complete, single-interface support certificate is required.

The separate Oblique dielectric tensor option defaults to off. It couples electric components at certified inclined planar boundaries between positive, lossless, nondispersive isotropic dielectrics with vacuum permeability on the local support. It supports stationary FDTD and explicit Dispersion-Free with a stationary or moving mesh. Implicit ATI and combinations with conductivity, dispersion, conductive walls or thin screens are not qualified. Sources remain in vacuum. It can increase memory and calculation time; aligned interfaces keep their ordinary diagonal treatment. The Solver checks stability at the selected step and never silently reduces it. Moving-tail accuracy requires application and mesh-refinement checks.

General curved interfaces, corners and sheet/PEC contacts do not acquire this tensor treatment. Inspect the preparation diagnostics and check convergence of the physical observable as the mesh is refined. An improved local coefficient does not guarantee a sufficiently resolved wave or geometry.

Rebuild calculation meshes made before this correction. Keep their saved jobs and results for comparison; do not resume an old checkpoint against a mesh with changed material coefficients. Reference example plots in the guides retain their original numerical observations.

Release 1.1.2 also uses indexed conducting-wall and thin-screen meshes. Rebuild older affected meshes when Solver reports an unsupported index or format; do not resume an incompatible checkpoint against the rebuilt mesh. Enabled dispersive volumes must satisfy the mesher's physical open-boundary admission checks. Read the object, material and boundary diagnostic before changing a model.

2.3 Physical mechanisms and their representations

2.3.1 Volume conductivity

For an isotropic conducting volume with physical SI conductivity $\sigma \geq 0$,

$$\boldsymbol{J}_\sigma = \sigma \boldsymbol{E}, \qquad p_\sigma = \boldsymbol{J}_\sigma \cdot \boldsymbol{E} = \sigma |\boldsymbol{E}|^2, \qquad P_\sigma(t) = \int_{\mathcal{V}} \sigma |\boldsymbol{E}|^2 \, \mathrm{d}V. \tag{2.2}$$

The local heat density rate is in $\mathrm{W/m^3}$. The field is the actual field in the conducting volume, not the vacuum field extrapolated into a conductor. Resolve the relevant penetration and material interfaces on the volume mesh.

For a nondispersive homogeneous material, the pure conduction substep obeys $\partial_t \boldsymbol{E} = -(\sigma/\varepsilon)\boldsymbol{E}$. A substep of duration δt gives

$$\boldsymbol{E}_{\mathrm{after}} = e^{-\sigma \delta t/\varepsilon} \boldsymbol{E}_{\mathrm{before}}, \qquad \Delta U_\sigma = \frac{1}{2} \int_{\mathcal{V}} \varepsilon (|\boldsymbol{E}_{\mathrm{before}}|^2 - |\boldsymbol{E}_{\mathrm{after}}|^2) \, \mathrm{d}V \geq 0. \tag{2.3}$$

Solver uses the discrete counterpart of this exact relaxation in two material half-steps around the shared field update. The reported volume heat is their electric-energy decrease, evaluated with the assembled edge energy weights. It is not a coarse-time estimate obtained by squaring preview samples.

2.3.2 Conductive walls and normal-depth lines

A conductive-wall object closes the active volume geometrically, like PEC, but adds a physical surface response. It has no finite-volume constitutive weight in the three-dimensional mesh. Its conductivity is not applied again as bulk σ inside the same wall.

With normal $\boldsymbol{n}$ directed from the field volume into the wall, the input power is

$$P_{\mathrm{port}} = \int_{\mathcal{S}} (\boldsymbol{E} \times \boldsymbol{H}) \cdot \boldsymbol{n} \, \mathrm{d}A = \int_{\mathcal{S}} \boldsymbol{E}_t \cdot (\boldsymbol{H}_t \times \boldsymbol{n}) \, \mathrm{d}A. \tag{2.4}$$

For a wall model with internal states, this is not all instantaneous heat:

$$\Delta U_{\mathrm{port}} = \Delta U_{\mathrm{wall,stored}} + D_{\mathrm{wall}} + U_{\mathrm{back}}, \qquad D_{\mathrm{wall}} \geq 0. \tag{2.5}$$

U_{back} is the separately accounted rear-termination work when a depth stack has a non-PEC backing. Solver's wall heat channel uses D_{wall}, not the full input work. A transient wall can return stored energy to the field while its dissipated energy still increases.

Studio physical model	Representation of the wall response
Thick finite-conductivity wall	A passive auxiliary-differential-equation (ADE) approximation to the thick, homogeneous good-conductor surface response; no explicit depth grid.
Homogeneous depth-resolved wall	A passive local 1D normal-depth line, with an automatically prepared grid and a PEC back plane.
Coating or multilayer wall	An ordered set of physical layers represented by a 1D depth line, with the selected PEC, resistive or bulk-conductor backing.

For a good conductor, the skin depth and surface resistance are

$$\delta(\omega) = \sqrt{\frac{2}{\omega\mu\sigma}}, \qquad R_s(\omega) = \sqrt{\frac{\omega\mu}{2\sigma}}. \tag{2.6}$$

In sinusoidal steady state, for peak-amplitude phasors, $\langle P_{\text{port}} \rangle = \frac{1}{2}\int_{\mathcal{S}} \operatorname{Re} Z_s \, |\boldsymbol{H}_t|^2 \, \mathrm{d}A$. This frequency-domain check does not replace transient energy accounting. Depth lines resolve normal penetration; they do not create a resolved three-dimensional coating or resolve lateral currents within the wall. Inspect the prepared Solver plan. The homogeneous Maximum depth cells is not an authored coating thickness; reducing it can also change the inferred depth in the current policy, so check both response accuracy and physical depth.

2.3.3 Resistive sheets and dispersive volumes

A zero-thickness resistive sheet is a two-sided screen, not a thick wall. For sheet resistance $R_\square > 0$ and tangential surface current $\boldsymbol{K}_s$, its heat rate is $\int_{\mathcal{S}} R_\square |\boldsymbol{K}_s|^2 \, \mathrm{d}A$. In the current implementation, its dissipated increments are added to the same CSV channel labelled Conductive wall; there is no separate sheet curve.

A real, frequency-independent ε_r or μ_r alone does not describe dielectric or magnetic damping. For a dispersive volume the electric and magnetic Lorentz states can absorb, store and return energy; positive damping produces irreversible loss. Tabulated complex response is fitted to such passive states. For sinusoidal peak fields and the material convention $e^{+\mathrm{i}\omega t}$, the physical cycle-average heat density can be written

$$\langle p \rangle = \frac{1}{2}\left[(\sigma + \omega\varepsilon_0\varepsilon_r'')|\boldsymbol{E}|^2 + \omega\mu_0\mu_r''|\boldsymbol{H}|^2\right], \tag{2.7}$$

where ε_r'', μ_r'' denote nonnegative loss parts ($-\operatorname{Im}$ of the relative response). Do not count the same conductivity both in σ and in an imported electric loss part. This is a physical harmonic identity, not an instantaneous formula for the CSV monitor. Solver accumulates irreversible damping at actual material-state updates; reversible polarization or magnetization energy is not heat.

Scope of the current Total curve

The complete eleven-column result sums *wall/sheet dissipation, bulk conductivity, electric dispersion and magnetic dispersion*. Historical seven-column results cover only wall/sheet and conductivity losses; their absent dispersion channels are unknown, not measured zeros. Both formats exclude rear-termination work, reversible stored energy, radiation and PML absorption. A moving-window record excludes subsequent relaxation after a material state leaves the window; discarding that state is not counted as heat. Use a reader compatible with the recorded schema; an older Studio may reject the new header.

For a multilayer wall, heat in the explicitly modelled depth layers is included in wall dissipation. The rear-termination contribution remains in the separate back-port report, including a resistive or conducting backing; do not silently add it to the CSV wall curve. There is currently no per-object

or per-layer breakdown in this artifact. Material Inspector identifies material placement; it is not a spatial heat-deposition map.

The separate `results/power_loss_materials.csv` records material IDs and names with their mechanism contributions. Studio can select all materials or one saved material. This is a material breakdown, not a spatial, per-object or per-layer temperature map. Aggregate-only historical files do not contain enough information to reconstruct it from the current project.

2.4 What the power-loss monitor accumulates

Let $D_a[n]$ be the physical dissipated-energy increment in each recorded channel in joules on the simulated domain for completed field step n. For a recording interval $\mathcal{I}_j$ containing m_j steps of duration Δt,

$$\overline{P}_{a,j} = \frac{S}{m_j \Delta t} \sum_{n \in \mathcal{I}_j} D_a[n], \qquad U_a(t_j) = S \sum_{n \le n_j} D_a[n], \quad a \in \{\text{wall}, \sigma, \text{electric}, \text{magnetic}\}. \tag{2.8}$$

S is the authored output symmetry factor; t_j is the end of the interval. Every field-step increment is retained in the sum, including a shortened final interval. Increasing Sample interval steps changes the plotted time resolution, not the field time step and not which steps contribute heat. It smooths peak power, but preserves final cumulative energy up to summation roundoff.

The monitor uses energy reports already produced by the shared field/material updates. It does not perform an independent field scan. Cartesian MPI owners accumulate their physical contributions, and the completed traces are reduced for one output writer. Studio need not reconstruct a whole three-dimensional field history to display this result. Field-preview recording is independent of material-loss recording.

2.5 CSV normalization: power, cumulative loss and joules

The complete artifact `results/power_loss.csv` contains time in seconds and ten curves: four mechanisms and their total, each as interval-average normalized power and cumulative normalized loss for one source. For its physical charge Q in coulombs, define

$$\mathcal{P}_a(t) = \frac{\overline{P}_a(t)}{Q^2}, \qquad k_a(t) = \frac{U_a(t)}{Q^2}. \tag{2.9}$$

Their SI units are $\mathrm{W/C^2}$ and $\mathrm{J/C^2} = \mathrm{V/C}$, respectively. For multiple synchronous sources, the CSV instead contains physical powers in W and cumulative energies in J, without charge normalization. The fields contain the full signed source combination, so quadratic dissipation includes interference terms; it is not the sum of separately calculated source losses.

For one source the CSV displays $\mathrm{W/pC^2}$ and V/pC. These curves are normalized by charge *squared*, although cumulative energy is expressed as a loss factor. The symmetry factor in Eq. (2.8) is already included.

Recorded channel	Meaning
Conductive wall power/cumulative loss	Wall-model and resistive-sheet dissipation, summed over active surface owners.
Volume conductivity power/cumulative loss	Heat from physical bulk σ relaxation, summed over active volume owners.
Electric/magnetic dispersion power and cumulative loss	Irreversible damping of the respective volume material states.
Total power/cumulative loss	Pointwise sum of the four recorded channels (two in historical files); not an independent balance measurement.

Let $Q_{\mathrm{pC}} = Q/(10^{-12}\ \mathrm{C})$ be the numerical charge in pC. If p_a and κ_a are the *numerical values* read from the CSV in $\mathrm{W/pC^2}$ and V/pC, the conversions are

$$\boxed{P_a[\mathrm{W}] = p_a Q_{\mathrm{pC}}^2, \qquad U_a[\mathrm{J}] = 10^{-12}\kappa_a Q_{\mathrm{pC}}^2.} \tag{2.10}$$

For example, $2\ \mathrm{W/pC^2}$ at 3 pC is 18 W, whereas a cumulative 2 V/pC at 3 pC is 18×10^{-12} J. Neither normalized curve should be multiplied by charge only once to obtain watts or joules. The corresponding loss-equivalent voltage would instead be $U_a/Q = Qk_a$.

When checking the CSV itself, integrate its *interval-average* power over each represented interval, not by interpreting the end timestamps as instantaneous power samples:

$$\kappa_a(t_j) - \kappa_a(t_{j-1}) = 10^{12} p_{a,j}(t_j - t_{j-1}), \tag{2.11}$$

with times in seconds and cumulative value zero at the start of the complete record. The 10^{12} converts $\mathrm{W/pC^2}$ integrated in seconds to V/pC.

2.5.1 Repetition rate is a separate question

For independent identical bunch passages separated long enough for the relevant field memory to decay, the average heat load is

$$\langle P_{\mathrm{heat}} \rangle = f_{\mathrm{rep}}\, U_{\mathrm{heat,one\ bunch}}. \tag{2.12}$$

The transient power peak in the CSV is not this repetition-averaged power. If long-lived modes survive between bunches, a single-bunch multiplication does not include coherent multibunch buildup. A thermal calculation also needs spatial heat deposition and cooling/material thermal data; the global CSV curves do not calculate temperature.

2.6 Studio workflow and interpretation checks

1. In Materials, select the physical representation, enter SI conductivity or the appropriate dispersion/layer data, and assign the material to the intended solid or screen. Inspect priorities and the resolved material mesh. A named but unused material contributes no loss.
2. In the always-visible Power-loss monitoring controls, enable recording. Set Sample interval steps and an explicitly justified Symmetry factor. Nonzero material conductivity alone does not enable this monitor.
3. Use $S = 1$ for a full structure. For a centred physical source on a genuine PMC symmetry cut, a half-domain uses $S = 2$, or a PMC/PMC quarter uses $S = 4$. Studio's domain fitting does not automatically justify or set this factor.

4. Run with the selected material models. On completion, Studio selects the power-loss artifact when available. To revisit it, select the run and `results/power_loss_materials.csv`, or the aggregate `results/power_loss.csv`, then Load / Analyze.
5. Use Quantity to select power or cumulative loss on one vertical axis. For one source their units are $\mathrm{W/pC^2}$ and V/pC; multi-source results use W and J. The initial plot shows total power. Enable Show loss mechanisms to inspect the contributing channels and use Material for all materials or one saved material when available. Export CSV retains the recorded channels and quantities independently of these display selections.
6. Check nonnegative interval power, nondecreasing cumulative loss, and the sum of the recorded mechanisms equals total. Extend the physical recording time until any required late dissipation is captured. A low instantaneous field or power does not by itself prove the cumulative asymptote has been reached.
7. Compare the wake loss factor with material losses only using the energy balance in Eq. (2.1). Include unrecorded mechanisms, residual stored energy, outgoing flux and any reference-field contributions before attributing a mismatch to the solver.

Wake parity normalization is not a heat multiplier

The PEC dipole rule $W_y^{\mathrm{odd}} = W_y^{\mathrm{PEC}}/2$ in Chapter 1 is linear in the field response. Heating is quadratic in fields. A mirrored PEC odd-sector calculation represents an antisymmetric source excitation, not automatically the heat of one off-axis bunch in the full structure. Do not reuse the wake division by two, or choose $S = 2$ or 4 merely from the number of cuts, to infer that single-bunch heat. Even and odd excitation sectors must be combined consistently for the physical source.

The frozen release examples provide two useful material configurations without changing any of their saved parameters: Example 04, *Lossy Bellows*, uses a thick-wall ADE conductor with $\sigma = 5.8 \times 10^7$ S/m and records every step with $S = 4$. Example 05, *Lossy Pillbox*, combines the same wall conductivity with a volume layer having $\varepsilon_r = 4$, $\mu_r = 1$ and $\sigma = 2$ S/m; it records every four steps with $S = 4$. They illustrate wall-only and combined wall/volume channels. No new numerical results for these examples are claimed in this guide.

CHAPTER 3

Definition and Calculation of S-Parameters

S-parameters describe how an electromagnetic wave incident through a waveguide port is reflected, transmitted or converted into other modes. Unlike a wakefield calculation, this problem has no drive bunch, witness trajectory or beam beta. This chapter covers the two-port multimode workflow already available in Studio [6].

3.1 Modal amplitudes and the S-matrix

A channel is a particular mode at a particular port. Let $a_{p,m}(f)$ be its incident complex amplitude, directed into the device, and $b_{p,m}(f)$ its outgoing amplitude. The port index p and mode label m identify physical channels, not time samples. At a fixed frequency,

$$\boldsymbol{b}(f) = \boldsymbol{S}(f)\boldsymbol{a}(f), \qquad S_{(p,m),(q,n)}(f) = \left.\frac{b_{p,m}(f)}{a_{q,n}(f)}\right|_{\text{all other incident amplitudes}=0}. \tag{3.1}$$

The first index is the output and the second is the excitation. Returning to the same port is reflection, including conversion to another mode at that port. Leaving through the other port is transmission.

With one declared mode at each of the two ports, this becomes

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}. \tag{3.2}$$

Excitation from port 1 gives reflection S_{11} and transmission S_{21}; excitation from port 2 gives S_{22} and S_{12}.

For a propagating mode in a lossless port, power normalization means $P_{\text{in}} = |a|^2$ and $P_{\text{out}} = |b|^2$: amplitudes have units $\sqrt{\text{W}}$ and S-parameters are dimensionless. This normalization refers to the time-averaged Poynting flux of the travelling mode,

$$P = \frac{1}{2}\,\text{Re}\int_A (\boldsymbol{E} \times \boldsymbol{H}^*) \cdot \boldsymbol{n}_{\text{travel}}\,\text{d}A, \tag{3.3}$$

where $\boldsymbol{E}, \boldsymbol{H}$ are peak-amplitude phasors and $\boldsymbol{n}_{\text{travel}}$ points along that wave's travel direction. Thus $|S_{21}|^2$ is a transmitted power fraction, not $|S_{21}|$ itself; $\arg S_{21}$ is its phase.

Studio supports numerical TE and TM port modes. Relative to longitudinal X, TE means $E_x = 0$ and TM means $H_x = 0$. These physical modal names are independent of the selected time-stepping method. Below cutoff a mode is evanescent rather than a propagating power channel; the present output is not an evanescent S-matrix.

3.2 Supported setup in Studio

Use Excitation Source to select Waveguide ports – S-parameters. Create a separate calculation if the original wakefield setup must be retained: conversion replaces the bunch, initial-field and wake-monitor sections rather than preserving a second hidden setup.

Setting	Current supported scope
Domain and method	Stationary; Open/Open X boundaries, PEC/PEC Y and Z boundaries. Use a supported stationary FDTD/Explicit or Dispersion-free Explicit/Implicit configuration and its compatible mesh.
Ports	Exactly two X-normal ports, facing inward from opposite ends. Their uniform corridors must contain vacuum and PEC only. Material-loaded or conductive-wall sections belong in the device, outside those port corridors.
Modes	At least one per port, with an ID, TE or TM family, and a positive family ordinal starting at 1. This ordinal selects a numerical mode, not a pair of analytic rectangular mode indices. Edit complete mode lists in YAML Source.
Excitation	A Gaussian-modulated pulse, with positive centre frequency, RMS bandwidth and amplitude in $\sqrt{\mathrm{W}}$. All frequency inputs are in Hz, not GHz.
Output band	A positive frequency interval within the pulse centre frequency plus/minus three RMS bandwidths. Only resolved propagating frequencies with sufficient incident signal can be interpreted.
Duration	An explicit positive Scaled simulation time, equal to cT in the case length unit; automatic duration is not supported here.

The *modal plane* is where the field is projected onto modes. The *reference plane* is where the reported amplitudes are defined. Studio's automatic placement uses the X domain, mesh spacing and PML depth. After changing them, use Reset automatic port placement and review the result. Each reference plane must be inward from its modal plane, still within the same uniform guide, and the two reference planes must remain ordered.

Translating amplitudes from the modal plane to the reference plane is *de-embedding*. In a lossless uniform port it changes propagation phase, not power. It does not move the geometry or remove a discontinuity or lossy section. Compare phases only at the same reference planes and in the same modal basis.

A declared mode is both an input and an output

With N total declared modes, one case requests N independent field evolutions and N^2 complex S-parameter traces at resolved frequencies. There is no single active-incident-mode selector. Reducing the mode lists reduces the matrix workload, but can omit real mode conversion.

3.3 Reading and checking the result

Build and inspect the mesh, then run the scattering calculation. In Results, open the scattering artifact:

```
results/scattering_parameters.csv
```

Studio shows each excitation-to-output channel's linear magnitude $|S|$ and phase in degrees against frequency in GHz. The CSV also retains real and imaginary parts, incident amplitude, reference-plane coordinates and a validity flag. Invalid rows are not shown as trustworthy coefficients. The optional incident-relative threshold, default 10^{-8}, rejects frequencies with insufficient incident amplitude relative to its spectral maximum; it is not an accuracy guarantee.

For one incident mode at port 1 and only one propagating mode at each port,

$$R = |S_{11}|^2, \qquad T = |S_{21}|^2, \qquad A = 1 - R - T. \tag{3.4}$$

If these are the only outgoing power channels, a lossless steady response has $R + T = 1$; a passive absorbing device has $A \geq 0$. For multiple modes, the corresponding balance includes the outgoing power in *every* propagating channel. Equivalently, for a single excited channel with incident power P_{in},

$$A = 1 - \frac{\|\boldsymbol{b}\|_2^2}{P_{\text{in}}}. \tag{3.5}$$

Here $\|\boldsymbol{b}\|_2^2$ is the total outgoing modal power, not a time average over numerical samples.

Studio labels this remainder Dissipated power fraction. It forms the curve only where the valid channel set is complete relative to the channels seen for that excitation in the result, and does not clamp negative values. This checks the available output inventory; it cannot establish that the declared mode list contains every physical propagating mode.

A power-balance remainder is not a material heat monitor

An omitted outgoing mode, an insufficient decay time, residual boundary reflection or a numerical error can also produce a nonzero remainder. Interpret it as absorption only after checking modal completeness and convergence. It does not separate wall and volume heat, or locate losses inside individual objects. The bunch-normalized loss curves of Chapter 2 are a different measurement; the scattering workflow does not retain those bunch-monitor settings.

Before interpreting a peak or deficit, check that the relevant modes propagate across the band, that incident content is adequate, and that port corridors remain uniform. Refine the mesh and extend the duration to check the spectrum: the physical frequency resolution is of order $1/T$, and selecting more output frequencies cannot recover a transient that has not been recorded. Changing reference planes within a uniform lossless corridor should preserve magnitudes while changing phases consistently.

S-parameters and beam coupling impedance are different responses of the same structure to different excitations. The former map incident port waves to outgoing port waves; the latter describe the force on a witness due to a moving charge. Studio does not convert an S-matrix into a wake or a beam impedance.