Liminus Validation Benchmarks

Six standard electromagnetic problems covering analytic validation and published-device reproduction. Every number on this page comes from one benchmark run on a single NVIDIA A100, with the setup, the reference and the error stated side by side.

Hardware **NVIDIA A100 80 GB PCIe** Compiler **NVIDIA HPC SDK 25.11**

Precision **FP32 field storage, FP64 arithmetic**

Fully reproducible: raw logs, reference data and post-processing scripts available on request

Results at a glance

Accuracy is always relative to the reference named in the row, never to another simulation code. Solve time is the time-stepping loop alone; throughput is Yee cells times time steps per second of solve time. Timed rows are the median of three repeats, which agreed within 2%.

Benchmark	Reference	Accuracy	Cells	Steps	Solve	Gcells/s
1. Dielectric sphere scattering	Mie series	RCS within 0.36 dB rms at 500 nm; scattering spectrum within 2%	16.0 M	62,600	117 s	8.6
2. Thin-film slab and Fresnel interface	Transfer matrix	R and T within 0.002; energy conserved to 1e-5	38 k	25,000	20 s	—
3. Metasurface unit cell	Khorasaninejad et al., Science 2016	Peak conversion 97.7% vs 95.4% published; 3.7% rms over the band	0.8 M	40,000	2 × 45 s	0.7
4. Gold nanoparticle (dispersive metal)	Mie with Rakić gold	Absorption peak within 1.4%, resonance within 4 nm	12.3 M	120,200	497 s	3.0
5. Microsphere resonator	Exact Mie poles	Five modes located within 1.2 nm; Q within 6%	16.0 M	83,922	155 s	8.7
6. Waveguide and directional coupler	Exact slab modes	Effective index within 9×10^{-4}; coupling length within 0.2%	5.3 M	12,600	12.5 s	5.4

Rows 1, 4 and 5 use the 252^3 and 231^3 graded meshes generated by Tidy3D for its own published versions of these examples, so the cell counts and step counts are directly comparable with that solver's job reports. Row 2 has no throughput figure: at 38,400 cells the run is limited by kernel-launch latency rather than by the GPU, so cells per second would not describe the solver.

The benchmarks in detail

01 Dielectric sphere scattering

Reference: Mie series (exact) · Bohren & Huffman

A 1 µm diameter sphere of permittivity 4 in air, illuminated by a plane-wave pulse centred at 500 nm through a six-face total-field/scattered-field box. The bistatic radar cross-section is projected to the far field from a near-field box and compared with the Mie series in both principal planes. Mesh: the auto-generated 252^3 grid of Tidy3D's published example (12.5 nm cells around the sphere), 12-layer PML, run to the same 10^{-5} energy shutoff.

RCS error, E-plane	RCS error, H-plane	Forward / back σ	Solve time
0.36 dB rms	**0.40** dB rms	**1.011 / 1.026** × Mie	**117** s

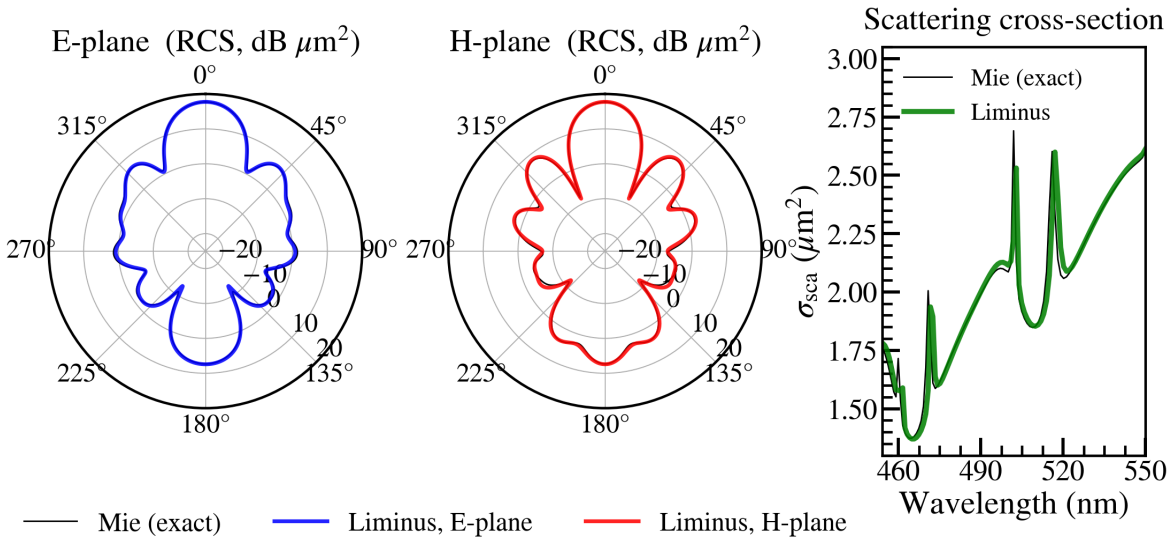

Figure 1. Bistatic RCS of the 1 µm dielectric sphere at 500 nm in the E- and H-planes (Liminus, thick; Mie series, thin), and the scattering cross-section across 454–550 nm.

Across 454–550 nm the scattering cross-section tracks the Mie curve to 2%, except within a few nanometres of three narrow Mie resonances, where the staircased sphere surface shifts the resonance by about one nanometre (benchmark 5 quantifies that shift directly). A uniform 12.5 nm mesh with three times the cells gives the same error level, so the graded mesh produces comparable error for this benchmark.

For the same problem on the same mesh, Tidy3D's job log reports a run time of 133 s on its cloud hardware; Liminus completes it in 118 s on one A100 (117 s of time stepping), and the step counts agree within 7%.

> **Sub-half-decibel scattering patterns on a vendor-generated mesh, at the same wall-clock time as a cloud solver.** FP64 field storage gives identical results in 151 s.

02 Thin-film slab and Fresnel interface

Reference: transfer matrix (exact) · Born & Wolf

A 1 µm film of refractive index 1.5 in air, and separately a single air-to-glass interface, illuminated at normal incidence over 400–700 nm with 5 nm cells on the periodic production path. Complex reflection and transmission coefficients, including phase, are compared with the transfer-matrix solution in 300 wavelength bins.

R, T error (film)	Phase of r, t	Energy closure	Interface R, T error
0.002 absolute	**0.2 / 0.5** deg rms	**6×10^{-6}** max |1−R−T|	**$1 / 4 \times 10^{-4}$**

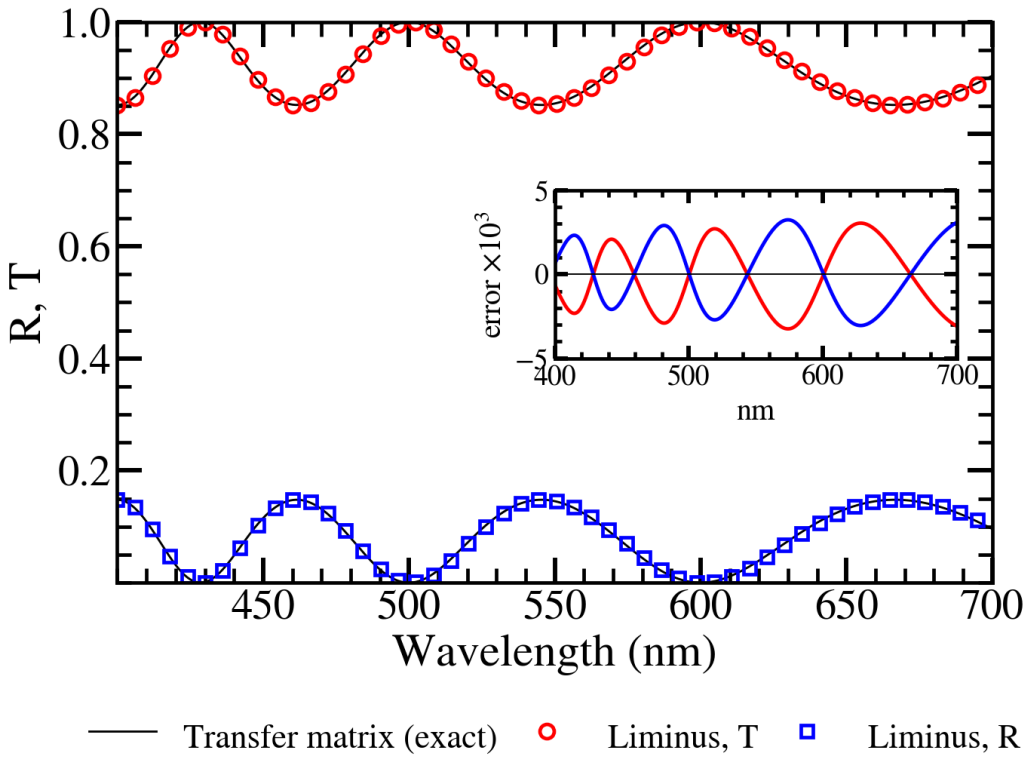

Figure 2. Reflectance and transmittance of the 1 µm n = 1.5 film: Liminus markers over the transfer-matrix curves; inset, the difference ×10³.

The Fabry-Pérot fringes land on the analytic curve to the third decimal across the whole band, including at the transmission maxima where R falls to zero. Energy is conserved to parts per million, while the agreement with the analytic R and T solution provides an additional check on boundary performance. The residual half-degree in the transmitted phase is the known numerical dispersion of the Yee scheme at 53 cells per wavelength inside the film.

> **Phase-accurate thin-film response with energy closure at the 10^{-5} level.** At 38,000 cells this case is limited by kernel launch latency, so its time is not a throughput figure.

03 Metasurface unit cell

Reference: Khorasaninejad et al., Science 352, 1190 (2016)

The 405 nm design of the Khorasaninejad metalens: a 150 × 40 × 600 nm titanium-dioxide nanofin on glass with a 200 nm pitch, dispersive TiO_2, periodic boundaries. Two linearly polarised runs give the Jones-matrix elements along and across the fin; the circular polarisation conversion efficiency is compared bin by bin with the curve digitised from the paper's Figure 1F.

Peak conversion	Published peak	Half-wave phase	Band agreement
97.7% at 402 nm	**95.4%** at 412 nm	**π — 2°** at 404 nm	**3.7%** rms, 410–700 nm

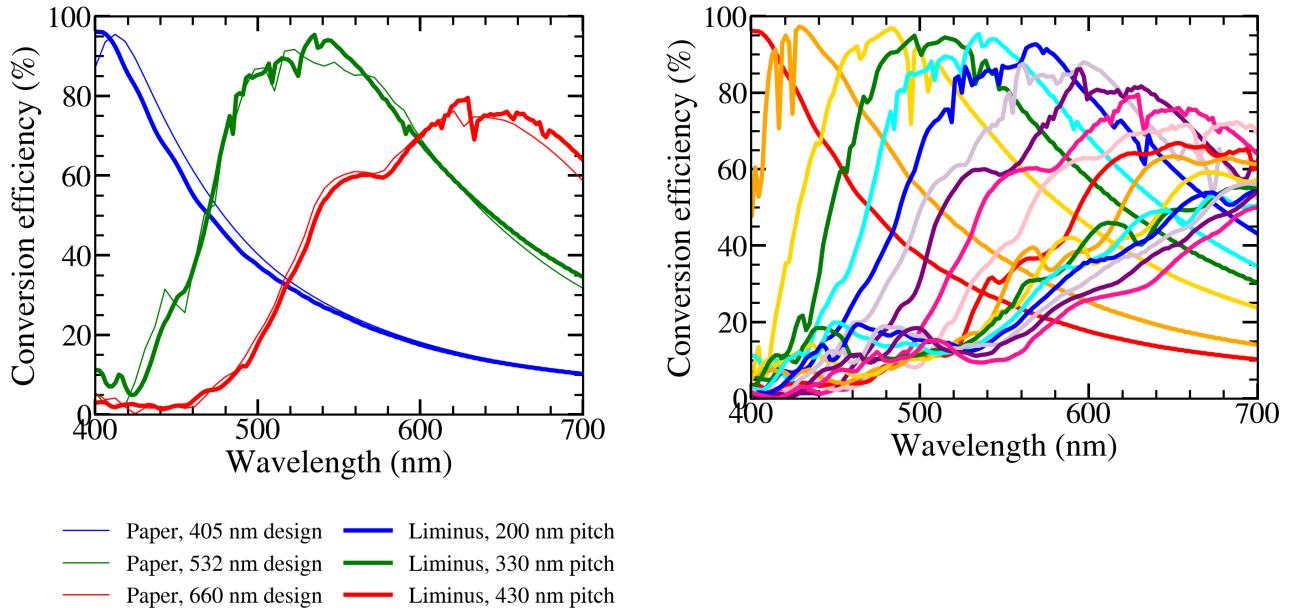

Figure 3. Cross-polarised conversion efficiency of the TiO_2 fin unit cells from the design sweep: (left) fins at 200, 330 and 430 nm pitch (thick) against the digitised curves of the paper's 405, 532 and 660 nm designs (thin), (right) the full sweep from 200 to 685 nm pitch, each curve one unit cell.

The fin reaches its half-wave retardation within two degrees of the design wavelength, and the conversion spectrum follows the published one within a few percent across the visible band, with the largest difference (5%) at 559 nm. Energy closure R + T stays within 0.1% of unity and the cross-polarised leakage of the source is below 0.4%, which are the internal checks that the two-run Jones-matrix method relies on.

A published metasurface design reproduced within a few percent across the reported band. Each polarisation run takes 45 s.

04 Gold nanoparticle, dispersive metal

Reference: Mie series with Rakić et al. gold, Appl. Opt. 37, 5271 (1998)

A 10 nm gold sphere described by the seven-pole Rakić Lorentz-Drude model, excited by a 442–662 nm pulse, on the 231^3 mesh of Tidy3D's published example, with 0.165 nm cells over the particle. Absorption and scattering cross-sections come from two flux boxes, forward RCS from the far-field projection; all three are compared with Mie theory evaluated with the identical permittivity model.

σ_{abs} peak	σ_{sca} peak	Resonance (Mie 489.5 nm)	Red HWHM (Mie 51.5 nm)
0.986 × Mie	**0.970** × Mie	**493.1** nm	**53.9** nm

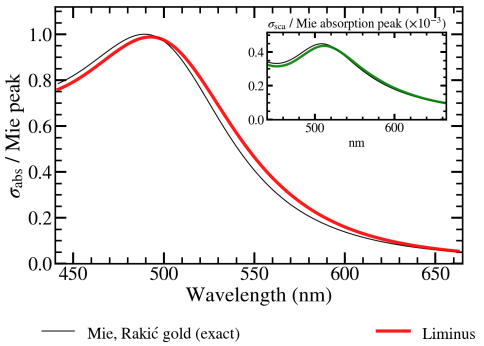

Figure 4. Absorption cross-section of the 10 nm gold sphere relative to the Mie peak (Liminus, thick; Mie with Rakić gold, thin); inset, scattering, three orders of magnitude smaller.

The plasmon resonance height, position and width are reproduced with a dispersive material advanced by an exponential-integrator pole-residue update, through 120,200 time steps of 0.31 attoseconds. The 4 nm red shift and the 4% broader line are consistent with staircasing of the sphere surface; the convergence study confirms that both decrease approximately in proportion to cell size. Rerun on meshes built with the same grading rules at 30, 60, 90 and 120 cells across the particle, the shift falls from 6.7 to 1.6 nm and the linewidth excess from 10% to 2% (Table 4.1). The peak height ratio, by contrast, drifts from 1.06 to 0.96 across the series and is not yet converged at the finest mesh. Scattering, the smaller of the two cross-sections by three orders of magnitude, is still within 4% rms over the band.

Table 4.1. Mesh convergence on graded meshes built with the Tidy3D rules; resonance from a parabolic fit to the peak (Mie 489.3 nm).

Cells across particle	Cell size (nm)	Resonance shift (nm)	Linewidth ratio	Peak height ratio	Band rms error
30	0.333	+6.7	1.104	1.059	29%
60	0.167	+3.3	1.048	0.989	10%
90	0.111	+2.2	1.042	0.971	5.2%
120	0.083	+1.6	1.022	0.960	3.5%

Multi-pole metals with the same numerical fidelity as dielectrics, at 3.0 Gcells/s. The seven-pole update costs about one third of the plain-dielectric throughput.

05 Microsphere resonator

Reference: exact Mie poles · Lam, Leung & Young, JOSA B 9, 1585 (1992)

The permittivity-4 sphere of benchmark 1 driven by a point dipole inside it, recorded by three field probes for a fixed 2 ps window. Harmonic inversion of the ring-down yields each mode's resonance wavelength and quality factor, which are compared with the exact poles of the Mie coefficients: five whispering-gallery modes between 460 and 553 nm with Q from 190 to 1,350.

Mode	λ exact (nm)	λ Liminus (nm)	Shift (nm)	Q exact	Q Liminus	Ratio
TE, order 10	460.22	460.91	+0.69	1346	1306	0.970
TM, order 9	470.81	472.03	+1.22	399	376	0.944
TE, order 9	502.08	502.57	+0.48	663	689	1.039
TM, order 8	515.90	517.04	+1.14	192	186	0.968
TE, order 8	552.89	553.28	+0.40	332	338	1.019

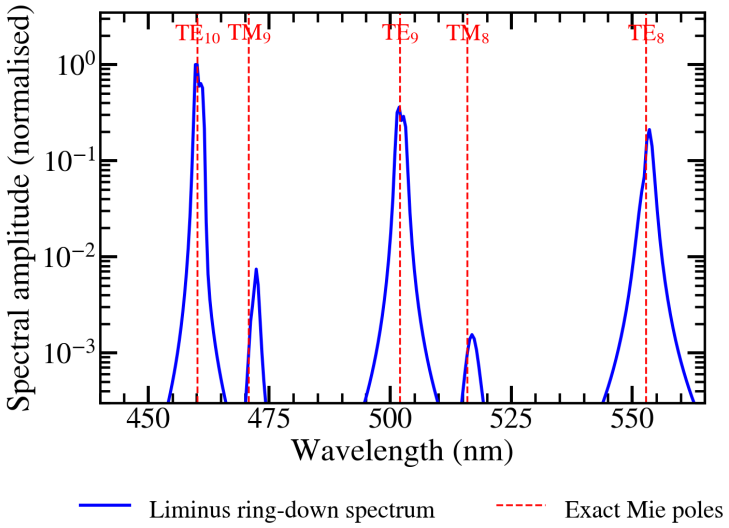

Figure 5. Ring-down spectrum of the dielectric microsphere with the five exact Mie pole positions marked; the doublet beside each peak is the cubic-grid splitting of the degenerate mode.

All five modes are found, each red-shifted by 0.4 to 1.2 nm (0.17% in frequency) by the staircased surface on a 12.5 nm mesh, with quality factors within 6% of the exact values. On a cubic grid each degenerate spherical mode splits into a multiplet spread over up to 1.9 nm; the table reports the strength-weighted multiplet. The lower-Q modes converge within about 1 ps of ring-down; the Q ≈ 1,300 mode needs the full 2 ps window.

Resonance positions to 0.2% and Q factors to a few percent from a single broadband ring-down. 83,922 steps in 155 s.

06 Waveguide and directional coupler

Reference: exact slab modes · Snyder & Love, Optical Waveguide Theory

A 0.5 µm wide slab core of index 1.99 in index-1.444 cladding, single-moded at 1.55 µm, excited by a mode-profile source and monitored over 60 µm of propagation at 1.41–1.72 µm. A second run adds an identical guide 0.4 µm away. The effective index comes from the phase accumulated along the guide; the coupling length comes both from the supermode phases and from the power exchange between the two guides.

n_{eff} error, 20 nm cells	n_{eff} error, 10 nm cells	Coupling length	Power conservation
$+9.1 \times 10^{-4}$	**$+2.2 \times 10^{-4}$**	**17.50** µm (exact 17.53)	**4×10^{-4}** along 53 µm

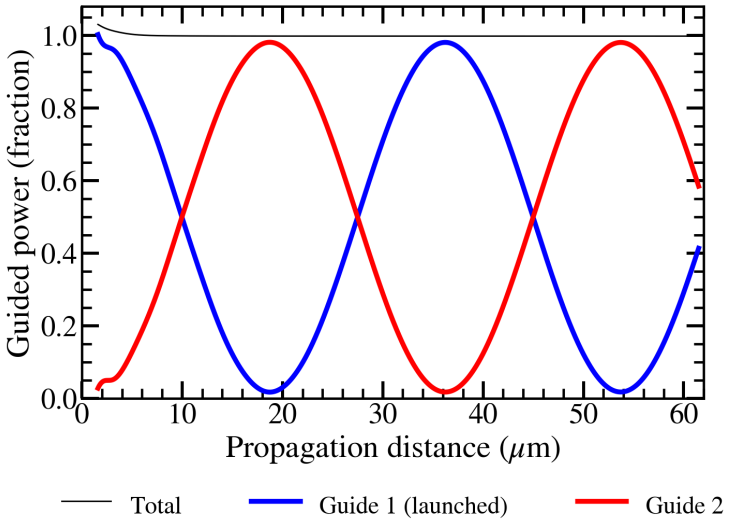

Figure 6. Guided power in the launched guide and its neighbour along the 0.4 µm-gap coupler at 1.55 µm, normalised to the total: the power crosses over every 17.5 µm, the exact coupling length, while the total stays flat.

The effective-index error is consistent with the Yee scheme's numerical dispersion: it is the same in the single guide and in both coupler supermodes, remains uniform across the band, and falls by a factor of 4.2 when the cell size halves. Because the coupling length depends on the difference between the two supermode indices, that error cancels and L_c lands within 0.2% of the exact value by two independent methods. 98% of the power transfers between the guides, the expected figure for a pair this strongly coupled. Guided power is flat to 0.04% over the monitored length, confirming lossless propagation and clean absorption at the end of the guide.

Second-order convergent waveguide modes and coupling lengths to 0.2%, in 12.5 s per run. The 10 nm convergence run takes 77 s.

The same job on different GPUs

Benchmark 1 (16 million cells, 62,600 steps) run unchanged across the GPUs we had access to. Solve times in seconds; shorter is better. The Tidy3D figure is the full run time of 133 s reported in its job log for the same mesh and run, shown for scale; the Liminus total including start-up and output is 118 s on the A100.

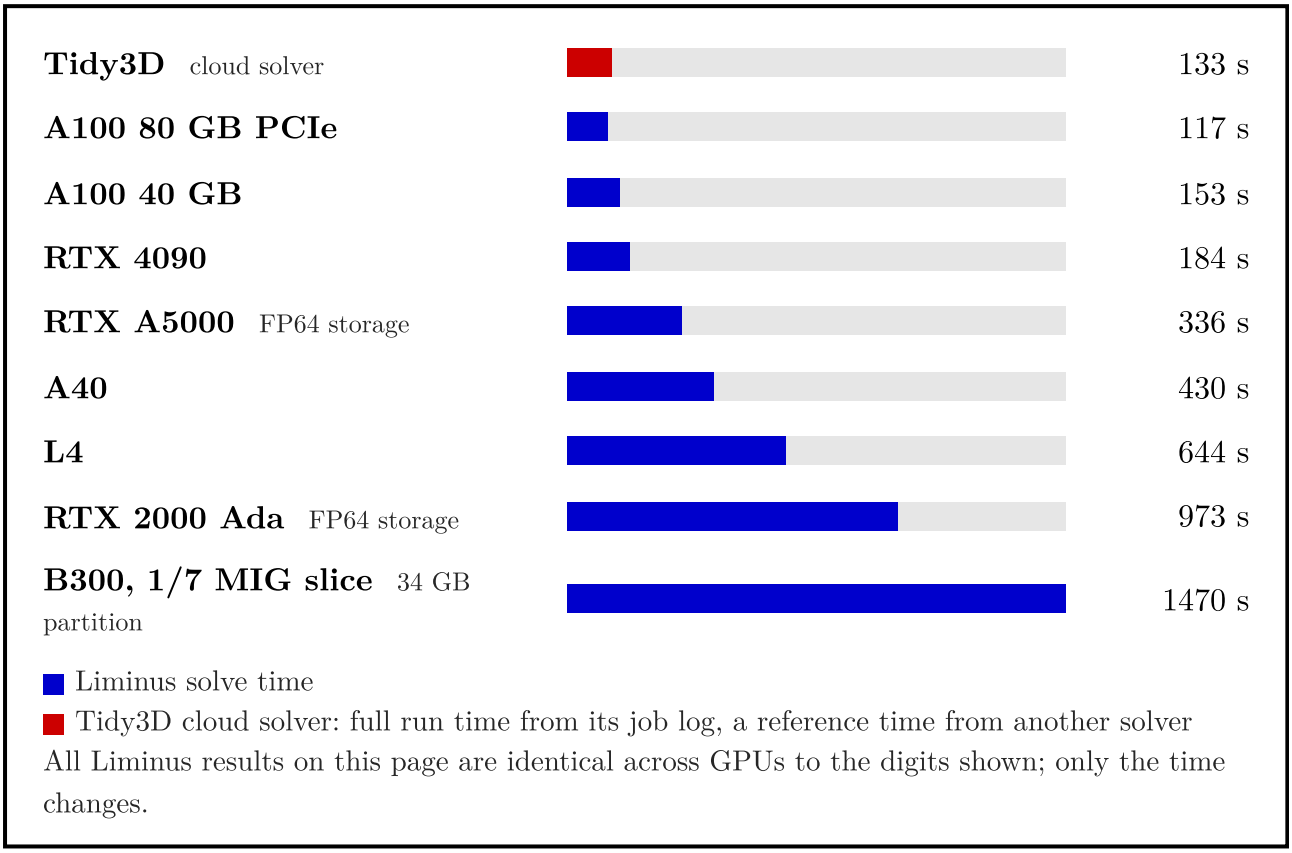

On cards with weak FP64 units (A5000, RTX 2000 Ada) the FP64-storage build is faster than FP32 storage because the arithmetic is FP64 either way and the conversions cost more than the bandwidth they save. A full FP32-arithmetic build is the obvious next step for that class of hardware.

How we measured

Solve time Wall-clock time of the time-stepping loop, from the first step to the last, excluding grid setup, device transfers and output. Total process time is typically 1–8% longer.

Throughput Yee cells × time steps ÷ solve time, in billions of cell-updates per second (Gcells/s). It includes the cost of all absorbing boundaries and frequency-domain monitors active in that run.

Repeats Timed rows are the median of three identical runs. Results are bitwise identical between repeats; the spread in solve time was below 2%.

Precision Fields are stored in FP32 and all arithmetic is carried out in FP64. Benchmark 1 was repeated with FP64 storage: the results agree to seven significant figures.

Meshes Benchmarks 1, 4 and 5 run on the graded meshes that Tidy3D generates for its own published versions of these problems, with the same time step, pulse, run length and energy shutoff. The remaining rows use uniform cells of 5–20 nm chosen for 39–53 cells per wavelength in the densest material.

Boundaries Convolutional PML on all open faces; periodic boundaries for the unit-cell rows. Every benchmark reports an energy-closure or power-conservation figure alongside its comparison with the reference; together they check the boundaries indirectly.

Geometry Objects are rasterised cell by cell (staircased), with no sub-pixel averaging. Where the reference is exact, the residual error is dominated by that staircase or by the Yee scheme's numerical dispersion. Convergence with cell size is demonstrated for the gold nanoparticle (benchmark 4, four resolutions) and the waveguide (benchmark 6, two resolutions); the other rows were run at a single resolution.

References Mie series and resonance poles evaluated from the Riccati-Bessel functions (Bohren & Huffman 1983; pole formalism after Lam, Leung & Young 1992); transfer-matrix (Born & Wolf) and slab-mode (Snyder & Love) solutions from closed-form expressions solved to machine precision; gold permittivity from Rakić et al. 1998; the metasurface curve digitised from Figure 1F of Khorasaninejad et al. 2016.

READ THESE NUMBERS WITH

- Timings are for one A100 80 GB PCIe in a shared cloud host. The small periodic cases (under a million cells) are limited by kernel-launch latency rather than by the GPU, and their times are sensitive to host CPU load.
- The Tidy3D run time quoted for benchmark 1 is taken from that service's own job log for its published example; the hardware behind it is not disclosed, so it is a like-for-like job comparison, not a like-for-like hardware comparison.
- Accuracy figures are specific to the stated mesh. Finer meshes reduce the staircase and dispersion errors; benchmark 4 demonstrates first-order convergence of the staircased resonance shift and benchmark 6 second-order convergence of the waveguide effective-index error, and the other rows were run at a single resolution.

Reproducing the results

Every row is produced by one scripted suite that builds the solver variant, runs the repeats, samples GPU memory, runs the post-processing against the reference and writes a metrics table. The suite specifications, the reference data with its provenance, the post-processing scripts and the raw logs of every run on this page are archived and available on request, so any row can be rerun on your own hardware and compared line by line.
