TERA verification

How TERA Has Been Checked, and How to Read HFTA

2026-11-26. A short guide to three long reports: the solver verification report, the Burke 1994 comparison and the HFTA study. The numbers come from those reports, which give the method and limits behind each one.

1 In one minute

- TERA's solver has been checked four ways: against exact answers, against another research code, against measurements made by helicopter, and in the configuration the product actually runs. The reports were critiqued over six rounds by an AI model (Claude Opus 5.5), given the reports and asked to find what was wrong. That is not an independent expert review. Every point it raised was checked against the code or a calculation before being accepted or rejected, and all six critiques are published with the reports.
- At the settings the product uses, the solver itself is good to roughly 0.2 to 0.9 dB near the peak of the pattern. The larger uncertainties are elsewhere: the ground constants you choose, and, below about 3°, how far the terrain profile extends.
- TERA is not presented as definitively correct. It is presented as checked, with the limits of those checks written down.
- HFTA was a large step forward when it appeared, and above about 10° it agrees with TERA to within about 1 dB. Its weak spot is the lowest angles, 1 to 5°, where it usually reads optimistic and shows about half the gain that extra height really buys. That comes from the method available to a program of its time (its calculation engine dates from 2003), not from mistakes in the code.

2 What "verified" means here

No program can be proven right on real terrain, because real terrain has no exact answer to check against. What can be done is to check the program everywhere an answer is known, check it against independent work where it is not, and say plainly what is left over. Each check below was built so that it could fail. Each exact check carries a control, a version known to be wrong, to show the test would catch a fault.

3 The four kinds of check

1. Exact answers. For a few shapes of ground the radiated field can be worked out exactly with pencil-and-paper physics: flat ground of any kind, a curved bump on flat ground, a round cylinder of lossy material, and a "made to order" field built to satisfy the equations over open, lossy terrain. TERA matches these to a fraction of a percent at fine settings and, at the settings the product uses, to about 0.3 dB on the bump and 0.27 dB (horizontal) and 0.54 dB (vertical) over flat ground at 3 to 20°.

2. Another research code. In 1994 Gerald Burke of Lawrence Livermore National Laboratory, a co-author of NEC, published patterns over a Gaussian hill and over Cedar Valley, Utah, computed with the same kind of solver TERA uses. His code was never archived, so TERA was compared against his printed figures. The two agree in pattern shape on every panel, and in absolute signal level to a median of 0.18 dB on 31 of 32 panels. Chasing the remaining differences found two defects in TERA, both since fixed. Fixing them improved its accuracy against the exact flat-ground answer by roughly a factor of ten.

3. Measurements. The Cedar Valley study (J. K. Breakall, WA3FET, and colleagues) measured elevation patterns from a helicopter. These are the only published measurements of this kind we have found. TERA matches the measured shapes to 1.2 to 4.2 dB RMS, the same as Burke's own model. Breakall advises the TERA project; the measurements predate TERA and are used as published.

4. The product as shipped. Checks were rerun at the mesh density, profile length and code path the web service, TERASweep and the command line actually use, on real profiles with real antennas. A test now fails if those three ever disagree. This is where the reviews paid off most: they found two faults that affected what users saw, not in the solver itself but in how terrain was fed to it. One modeled the ground behind the antenna as flat, which on hillside sites read 45 to 90° up to 12 dB low. The other thinned the terrain input on the low bands. Both are fixed in the current release.

4 How far to trust a TERA pattern

Rough sizes of each source of error at production settings, smallest first:

Source	Typical size	Notes
Elevation data	about 0.1 dB	against 1 m lidar
The solver	0.2 to 0.9 dB near the peak	horizontal median 0.22 dB
Ground approximation	0.2 dB on average ground, up to 0.7 dB on the City and Dry Ground presets	very thin custom grounds are flagged on the card
2D terrain	0.5 dB or less above about 6°	estimated at one site; not estimated below 6°
End of the profile, below 3°	usually within 1 dB, several dB on some bearings	the plot carries a note below 3°
Your ground constants	not bounded	usually estimated, rarely measured

Known limits, each tracked as an open issue:

- Like HFTA, TERA works along one bearing at a time and treats the terrain as extending unchanged sideways (2D).
- A vertical dipole above 45° can read up to 1.5 dB low.
- Over seawater and perfect ground the edges of the modeled strip add a fine ripple of up to about 1 dB. The card notes it.
- Ground loss in the immediate neighborhood of the antenna is not modeled. This matters mainly for vertical antennas, whose figures will read optimistic.

5 Reading HFTA results

Many stations chose their antenna heights with HFTA, written by Dean Straw, N6BV, and distributed with the ARRL Antenna Book. TERA was compared with HFTA using HFTA's own calculation engine, the YTWCore.dll released with it in 2003, given exactly the inputs the HFTA program gives it. The comparison covers three named sites and a random draw of 100 US stations on 40, 20, 15 and 10 m, with horizontal Yagis from 30 to 100 ft.

Why the two differ

HFTA's calculation engine was released in 2003. It traces rays: it follows the direct path, the reflections off each patch of ground, and the bending of signal over ridges (diffraction), and adds them up. Ray tracing is fast, which suited the home PCs of the time, and it works well where the paths are clear-cut. It struggles at very low angles, where signal skims the ground. There it can show gain that isn't there, or miss signal that is.

This is a limit of the method, not of HFTA's programming. TERA has its own ray tracer, and on one of the worst bearings in the study it overstates low-angle gain even more than HFTA does. That is why TERA's default engine is not a ray tracer. It solves for the currents the antenna induces in the ground along the whole profile (the method of moments). This captures reflection, diffraction and ground loss together, with nothing to add up path by path. It takes far

more computing than ray tracing, but on today's hardware a bearing takes seconds.

N6BV was candid about all this. HFTA's manual estimates its accuracy at ±3 dB, warns of artificial spikes and holes at particular heights, and asks readers who have validation data to get in touch.

Where HFTA holds up

- Above about 10° the two agree to within about 1 dB at every site studied.
- Over flat ground HFTA's worst error near the pattern peak is 2.95 dB, inside its author's own ±3 dB (TERA's is 0.90 dB). It gets the gain from added height right to within 0.13 dB.
- From a summit with ground falling away on all sides, HFTA is a fair guide, reading a dB or two optimistic at low angles.
- Its shortcuts on input (thinning and smoothing the terrain) cost under 0.3 dB.

Where to be careful

There is no exact answer on real terrain, so these are differences between HFTA and TERA rather than errors measured against the truth. TERA's own error was measured in the same 1 to 3° window over flat ground: it accounts for none of the 40 m difference and about 0.4 dB of the 20 and 15 m differences.

- Low angles read optimistic. At 1 to 3° HFTA reads higher than TERA at most stations, by a typical 5.4 dB on 40 m, 2.8 dB on 20 m, 1.4 dB on 15 m and 0.6 dB on 10 m. The lower the band, the larger the gap.
- Height looks less valuable than it is. HFTA shows about half the low-angle gain from added height that TERA does: roughly 0.7 dB per 10 ft against 1.2 to 1.5 dB. This held at 92 to 97 of 100 stations, the most consistent finding in the study. Following HFTA's best height lost more than 2 dB on at least one bearing at three-quarters or more of the stations.
- Sharp low-angle spikes are suspect. A lobe more than 10 dB above TERA turned up somewhere at nearly every station. N6BV's own advice applies: rerun at 1 ft above and below, and distrust a feature that comes and goes.
- "No signal" into rising terrain is not literal. At about half the stations HFTA shows no field at all at some low angles where diffraction over the ridge still carries a weak signal. The direction may be poor, but extra height still helps there.
- HFTA stops at 4.41 km. At about 4 stations in 10, terrain beyond that distance changes the 1 to 3° answer by more than 2 dB on some bearing. At one site a ridge 7.5 to 9.5 km out cost up to 10 dB that HFTA could not see. TERA uses 10 km by default.
- Ground absorption is left out of HFTA's main reflection calculation. It keeps the phase of each bounce but not how much signal the ground absorbs, so its lobes read up to 1 dB high even over flat average ground.

What this means for an existing HFTA analysis

If you picked a height with HFTA and your interest is high-angle work, the answer is likely to stand. If your interest is DX at the lowest angles, especially on 40 m, rerun the site in TERA: expect less low-angle gain than

HFTA showed, more benefit from added height, and a different picture on bearings that look into rising ground or toward distant ridges.

6 Where to read more

- Solver verification against exact answers, with the full error budget:
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-solver-verification.pdf
- Comparison with Burke 1994 and the Cedar Valley measurements:
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-burke-comparison.pdf
- HFTA compared with TERA at three sites and 100 stations:
  https://paas.k6tu.net/wp-content/uploads/tera/reports/hfta-tera-comparison.pdf
- What HFTA's author published about its accuracy: the "Caveats" section of Operating Instructions for HFTA, distributed by the ARRL with the Antenna Book:
  http://www.arrl.org/files/file/Product%20Notes/Antenna%20Book/hfta.pdf
- The six AI critiques of the two reports, in order, each with the prompt it was given. Each report revision says how it answered the round before. The critiques cite issue numbers, commits and scripts from the TERA repository, which is not public. Rounds 1 to 6:
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-critique-round1.pdf
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-critique-round2.pdf
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-critique-round3.pdf
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-critique-round4.pdf
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-critique-round5.pdf
  https://paas.k6tu.net/wp-content/uploads/tera/reports/tera-critique-round6.pdf