

Holographic PNF Filtering Based On Known Volumetric AUT Bounds

Scott T. McBride, Pieter N. Betjes

NSI-MI Technologies

Suwanee, GA USA

smcbride@nsi-mi.com, pbetjes@nsi-mi.com

Abstract— There has been much discussion in the last few decades regarding redundancy in conventional near-field sampling, and that redundancy is most pronounced in the planar geometry. There has also been much discussion regarding modal filtering of near-field data to attenuate the effects of stray signals. Both discussions revolve around the limited local spatial bandwidth that can be produced at each probe location when the antenna under test’s (AUT’s) radiating sources are all contained within a known geometric boundary.

This paper discusses a novel filtering technique that exploits the inherent sampling redundancy in conventional planar near-field acquisitions. The filtering is based solely on the known location and shape in the scanner’s coordinate system of a closed 3D boundary around the radiators of interest. The paper describes the algorithm and presents results from both measured and synthesized input. The new filter is also compared to other available filters in terms of speed, attenuation of stray signals, and preservation of AUT signals.

Keywords: modal filtering, spatial filtering, holographic filtering, stray signals, planar near field

I. INTRODUCTION

The Nyquist sampling theorem says that a planar near-field (PNF) acquisition must sample the field at a rate slightly greater than the local spatial bandwidth[1] (or an increment $< 1/\text{bandwidth}$). A conventional PNF acquisition uses constant increments in X and Y, and those two fixed increments must be suitable for the largest possible spatial bandwidth. For a conventional (constant-increment) PNF acquisition with the probe close to the antenna under test (AUT), that constant increment needs to be close to $\lambda/2$. Because the AUT field’s spatial bandwidth decreases dramatically beyond its aperture[1], a conventional PNF acquisition contains far more data than are required (in the absence of measurement errors) to characterize AUT parameters. The filtering technique in this paper takes advantage of that oversampling to resolve received power that can be radiating from the AUT from power that cannot.

Section II describes this new filter’s algorithm. Key metrics of any filtering technique ought to include its preservation of AUT signal, and then its suppression of non-AUT signal. Section III investigates the algorithm’s ability to address these two competing goals, as well its speed. Synthesized data are used in Section III so that the correct AUT pattern is known. Section IV then applies the filter to gain-standard data measured in the presence of numerous stray signals. The true pattern of the AUT

in Section IV is not known, but Section IV provides a high-level confidence check of the algorithm.

Modal-filtering techniques exist [2]-[6] that have the same goals as this new holographic PNF filter. Most of those techniques involve a full PNF transform, interpolation to a different (and very dense) angular grid, coordinate-origin translation, and then a far-field-to-far-field transformation with modal filtering. By contrast, the holographic PNF filter has no interpolation and primarily uses several 2D fast Fourier transforms (FFTs) to produce filtered PNF data that can then be input to the standard PNF transform. Section V compares the new filter’s performance to that of existing filtering algorithms in terms of AUT-pattern preservation, stray-signal rejection, and processing complexity.

The basis of this algorithm is illustrated in Figure 1, which shows the following:

- A 3D box bounding all radiating sources of interest as a black rectangle in the Z-Y plane, though any shape(s) could be used
- The central Y scan of PNF data at the right edge ($Z=0$) synthesized from those radiating sources, contrived to produce a challenging pattern with no stray signals
- The central Y scan of several X-Y holograms at different Z values, forming a “Z-Y hologram” (used only for illustration)
- All power from the AUT received at the probe plane will have come from the AUT volume, so all holographic flow lines (other than those due to scan-plane truncation) pass from $Z=0$ through the AUT volume.

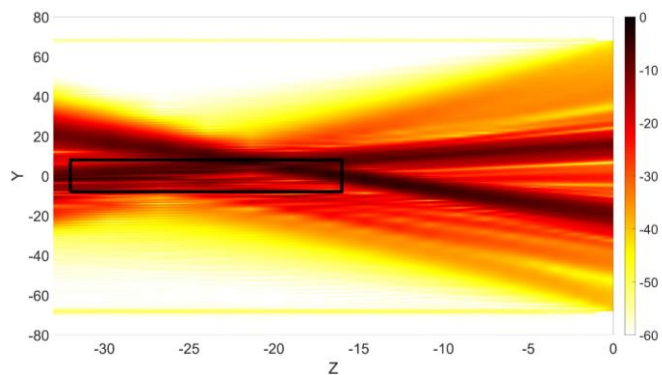


Figure 1 Z-Y hologram for a volume of radiators.

II. ALGORITHM

This filter uses a combination of microwave holography[6] and spatial filtering to separate stray signals from AUT signals. The only assumption is that all radiators of interest are contained within a known volumetric boundary. The boundary chosen here is a rectangular box, though more complicated (and/or multiple) shapes should be equally valid.

A. Overview

For a box with zero thickness in Z, this filter's output would be identical to the following (and to earlier methods[6]):

- Use holography to estimate the complex fields vs. X and Y at the AUT's known Z location.
- Taper the estimated fields outside the known AUT boundary to zero smoothly over a prescribed distance to minimize added truncation errors.
- Use holography to estimate the complex fields without stray signals vs. X and Y at the probe's Z location.

The perceived drawback of the above approach is its limitation to a zero-thickness volume. For an AUT (including support structures or other features that might contribute to the field) with any thickness, one should expect the out-of-plane sources to be defocused in the one hologram. That defocusing could cause it to extend outside the specified X-Y bounds, and the filter would attenuate that portion of true AUT contribution.

The approach taken with the new filter is to back-project to several Z locations, or slices, within the specified AUT volume. Because we expect some out-of-focus content in every Z hologram, (and because we expect that a larger Z size should retain more AUT signal than a shallow box) we cannot zero out the field outside the X-Y box at each Z. Instead, knowing that the field is the complex sum of the AUT signal and the stray signals, we choose instead to accumulate those signals that *cannot* be originating in (or passing through) the AUT volume. To do this, we zero out everything *inside* the X-Y box at each Z location, tapering that interior truncation to unity smoothly over a specified (perhaps zero) distance outside the AUT box. The taper profile used herein is shown on a subset of the x-y plane in Figure 2.




Figure 2 Taper for zeroing field inside AUT box.

After emptying the signal out of the AUT box in each of the several Z slices, we should have an estimate of stray signals. The

measured data are the complex sum of these stray signals and the AUT signals we wish to evaluate. The final step is then to propagate the stray-signal distribution back to the measurement plane (nulling the spectrum outside the sine-space unit circle) and subtract it from the original complex measured data.

To summarize the method, the loop through the AUT volume starts with all signals, and then removes any content that could possibly come from the AUT in each Z slice. That final stray-signal estimate (propagated back to the probe plane) then gets subtracted from the measured data, leaving any contribution that might have come from the AUT.

Figure 3 illustrates the holographic PNF filter stepping through the AUT volume, initially assuming that all measured signals are stray, subsequently removing any signals that might come from the AUT, and therefore accumulating only the stray signals to be removed. Each plot shows a Z-Y hologram just before the AUT-signal removal in each of the five specified Z slices. The scan plane is on the right edge of each plot, and the black rectangle represents the modeled and specified AUT volume. The modeled PNF data (described in Section III) represent a complicated pattern with no stray signals. Had there been a sixth plot in Figure 3 (after the final filtering slice), it would have been blank (as desired with no stray signals present).

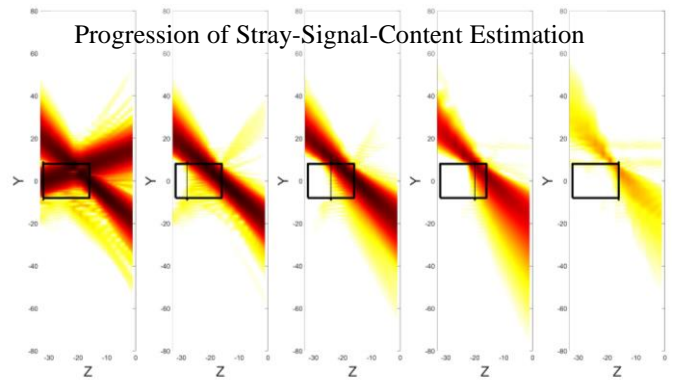


Figure 3 Successive removal of AUT signal (no stray).

It should not be assumed that all the stray signals will have been removed. A broad-beamed scattering source directly behind the AUT volume (and somehow not occluded), for example, could have most of its contributions to the scan plane pass through the AUT volume, and thus be largely unattenuated.

B. Details

1) Pad with Zeroes

The first step is to zero-pad the measured complex probe voltage $v_p(x,y)$ in both dimensions, placing the non-zero data near the center of the larger stray-signal matrix v_s as the initial (and later updated) estimate of stray-signal content. This zero padding might be necessary if stray signals come from locations well outside the X-Y probing extent. If rays from the scan plane toward such a stray signal run off the edge of the zero pad, then they will alias to the opposite edge and perhaps continue toward the AUT volume. It will also be beneficial to choose the dimensions of the padded matrix to be consistent with the FFT algorithm in use, usually powers of two.

2) Define the Phase-Propagation Matrices P_{Z1} and $P_{\Delta Z}$

We first use (1) through (4) to compute the matrix K_z , and then (5) and (6) to define two constant propagation matrices:

$$K_x = K_0 \left(-\frac{N_x}{2} : \frac{N_x}{2} - 1 \right) * \frac{\lambda}{N_x \Delta x} \quad (1)$$

$$K_y = K_0 \left(-\frac{N_y}{2} : \frac{N_y}{2} - 1 \right) * \frac{\lambda}{N_y \Delta y} \quad (2)$$

$$K_{r2} = \frac{\text{fftshift}(K_x^2 + K_y^2)}{K_0} \quad (3)$$

$$K_z = \begin{cases} K_0 \sqrt{1 - K_{r2}}, & K_{r2} < 1 \\ 0, & K_{r2} \geq 1 \end{cases} \quad (4)$$

$$P_{Z1} = e^{-jK_z Z1} \quad (5)$$

$$P_{\Delta Z} = e^{-jK_z \Delta Z} \quad (6)$$

where

$[K_x, K_y, K_z]$ represents a propagation vector away from the AUT for each element of the FFT output matrix,

N_x and N_y are the dimensions of the zero-padded matrix, with even numbers assumed in (1) and (2),

Δx and Δy are the sample spacings in X and Y,

K_0 is the wave number $2\pi/\lambda$,

λ is the wavelength in the same units as Δx and Δy ,

$Z1$ is the negative Z location (in the same units as λ) where the AUT box is furthest from the probe, and

ΔZ is the user-specified positive spacing (in the same units as λ) among slices through the AUT box.

The standard `fftshift()` function, which moves the broadside aspect ($K_z = K_0$) to the first spectral indices, is applied to the K_z matrix one time (per frequency) so that there is no need to shift the spectral FFT output at each Z position.

3) Define Tapering profile T

At each Z slice through the AUT volume, we want to zero out the spatial-domain content within the bounds of that volume. Since there might be significant defocused AUT signal outside the X-Y bounds in that slice, it seems beneficial to taper that zeroing in a way that minimizes truncation effects on the spectrum. Figure 2 illustrates the taper profile used herein, with a general formulation shown in (7):

$$T = \begin{cases} 0, & |x| \leq x_{max} \text{ and } |y| \leq y_{max} \\ 1, & |x| > x_{max} + x_t \text{ and } |y| > y_{max} + y_t \\ 0 \text{ to } 1, & \text{in between} \end{cases} \quad (7)$$

where

$\pm x_{max}$ and $\pm y_{max}$ are the bounds of the AUT volume, assuming a box centered at $x=y=0$,

x_t and y_t are the X and Y distances over which to apply the taper, and

the in-between region changes continuously from 0 to 1 as shown in Figure 2.

4) Loop Z Through AUT Box

The first step is to propagate the zero-padded measurement v_s to the back of the AUT volume with (7) and (8):

$$S = \text{fft}^{-1}(v_s) \quad (7)$$

$$v_s = \text{fft}(S \times P_{Z1}) \quad (8)$$

Then we multiply the blanking taper T by this initial v_s at $z=Z1$ in (9):

$$v_s = v_s \times T \quad (9)$$

Then, for each slice through the AUT volume, we repeat (7) to find the modified stray-signal spectrum, substitute (10) for (8) to estimate the remaining stray signal at the new slice, and then apply (9) again to blank out any signal within the AUT volume.

$$v_s = \text{fft}(S \times P_{\Delta Z}) \quad (10)$$

5) Propagate Back to Scan Plane

Now we have removed all potential AUT signal from spectrum S at the front of the AUT box. We need its spatial-domain v_s at the probe plane, and we get that from (11):

$$v_s = \text{fft} \left(\frac{S}{P_{\Delta Z}^n P_{Z1}} \right), \quad (11)$$

where n is the number of ΔZ propagations performed.

6) Subtract Accumulated Stray Signal From Measured

The final $v_s(x,y)$ found in (11) now represents the complex probe-plane voltage that cannot have come to the probe directly from the AUT volume. The measured probe voltage $v_p(x,y)$ is the complex sum of all signals that reached the probe. Therefore we use (12) to remove $v_s(x,y)$, leaving only the signals that might have come from the AUT in the filter output $v_f(x,y)$.

$$v_f(x,y) = v_p(x,y) - v_s(x,y) \quad (12)$$

III. SYNTHESIZED DATA

In this section we model the AUT primarily as a pair of simple rectangular dipole arrays. One array, with a Hanning aperture distribution, is at the back of the AUT box steered upward at 25° . The other array, with a Blackman-Harris aperture distribution, is at the front of the AUT box steered downward at 50° . Several other uniformly weighted dipoles are scattered randomly within the AUT volume (as part of the AUT). No interaction is modeled among any dipole pairs. The “Truth” data were determined by propagating each AUT dipole to a dipole probe at a very large distance. In this section, equivalent stray-signal level (ESSL) is reported relative to this Truth data, and will thus include PNF truncation effects.

Two PNF datasets were produced and processed with this synthesized AUT. The first has no stray signals injected, and is included to ensure that the new filter does not perturb the contributions from within the AUT volume. The second has two simple stray signals injected, with attenuated reflections of the $+25^\circ$ and -50° beams through simulated horizontal plates so that they appear at elevations of -25° and $+50^\circ$, respectively. These stray signals are modeled to ensure that the filter is able to remove them.

A. Preservation of AUT Properties

Figure 4 shows the elevation principal plane of the AUT’s far-field pattern (“Truth”) as the thick black trace, and the ESSL between the filtered and Truth as the dashed red line. If we apply the filter to PNF data with no stray signals modeled as is the case in Figure 4, then we expect a good overlay among measured, filtered, and Truth, except for any scan-plane truncation effects. We see that desired overlay in Figure 4, suggesting that this filter achieves its primary goal of preserving AUT characteristics.

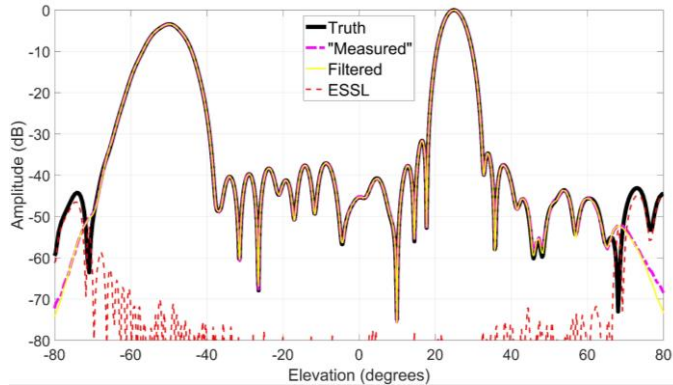


Figure 4 Overlays with no stray signals.

B. Rejection of Signals from Outside the AUT Bounds

Next, we modified the simulation to include a pair of attenuated specular reflections from a notional pair of horizontal flat plates. (These are conceptually similar to floor and ceiling bounces, but for illustration purposes those flat plates were located within the vertical range of probe travel.) Figure 5 shows the Z-Y hologram before and after the filtering. The left plot shows the two stray signals arriving at apparent aspects of $+50^\circ$ and -25° , but originating outside the assumed AUT volume. The right plot shows that nearly all of that stray content has been

removed, though the conservative taper function in Figure 2 left a little of the $+50^\circ$ stray signal in the filter output.

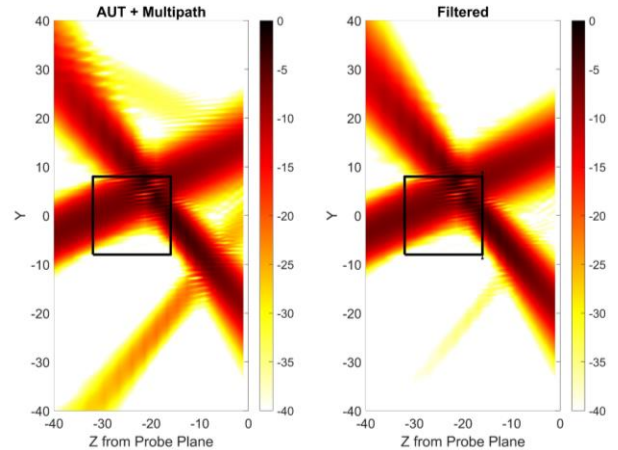


Figure 5 Z-Y holograms with stray signals injected.

Figure 6 again shows the overlay of Truth, unfiltered, filtered, and ESSL in the centerline elevation cut. We see that the two strong injected stray signals were there in the unfiltered data and are no longer visible in the filtered data. More significantly, the ESSL between filtered data and Truth appears unchanged with the addition of the strong stray signals in the filter input. This indicates that the filter met both competing goals of AUT-signal preservation and stray-signal removal, at least for the signals modeled.




Figure 6 Overlays with stray signals injected.

IV. MEASURED DATA

The measured data evaluated here have no shortage of stray-signal contributions. The AUT was an upward-looking gain horn beneath a large horizontal PNF scanner. These were located on the factory floor during early scanner integration with no absorber treatment anywhere. Known stray signals included the floor, ceiling, scanner structure, scissor lifts, metal shelving units, and multiple leakage sources. Figure 7 shows the PNF amplitudes before and after applying the filter.

Figure 8 shows the overlaid principal planes of those transformed PNF data. The ESSL trace shows the difference between the measured and filtered results, and is not referenced to Truth because the true pattern of this AUT as mounted is not

known. However, the filtered output does seem much closer to an expected gain-standard pattern than is the measured output.

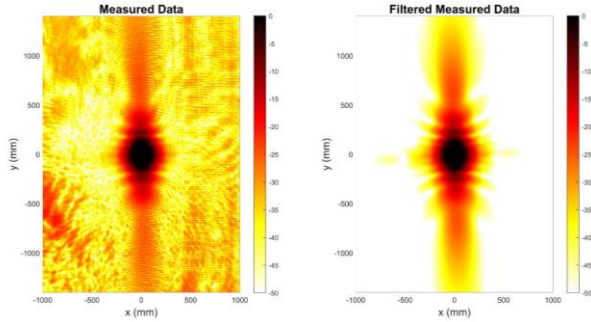


Figure 7 Filtered near-field amplitudes.




Figure 8 Filtered far-field principal planes.

V. COMPARISON TO OTHER FILTERING TECHNIQUES

A. MARS/IsoFilter™

The mathematical absorber reflection suppression (MARS) [2][4] and *IsoFilter™* [3] techniques are very similar to each other, and have the same purpose as the holographic PNF filter introduced here. Rather than altering coefficients in the spectral or modal-coefficient domain as MARS and *IsoFilter™* do, however, the holographic PNF filter alters the content of the spatial domain.

The processing in MARS (usually) and *IsoFilter™* (always) operates on the coefficients of spherical modes. Figure 9 illustrates several of the spherical modes[5] that those coefficients multiply. The lower-order modes (at the bottom rows and then the center columns in Figure 9) represent fields possible from antennas of small electrical size located close to the far-field origin. Larger antennas and/or larger separation of radiators from the origin require higher-order (in addition to lower-order) modes in their characterization, usually with the highest mode number according to the common formula KR_0+10 . By translating the far-field origin to the AUT and then attenuating and/or truncating higher-order modal coefficients, signals originating from outside the corresponding minimum sphere R_0 have their contributions attenuated.

Spherical modes and their coefficients are not native outputs of standard PNF processing, and are typically produced by a partial spherical near-field (SNF) transformation of the PNF far-field transform output. In order to fully capture all potential stray signals in the measured PNF data, the SNF-input data must be interpolated to a regular spherical grid with increments

corresponding to a minimum-sphere (and minimum-cylinder) radius equal to the maximum distance from probe to AUT center during the PNF acquisition[5]. This interpolation is often denser, and thus more time-consuming, than what the PNF transform would normally do.




Figure 9 Spherical mode visualization[5].

The planar “modes” are very different from the spherical modes. Each planar mode is a constant-amplitude phase gradient that is periodic over the zero-padded spatial extent, and higher-order modes have more periods of phase over that extent than the lower-order modes. Unlike the SNF modes, there is no correlation between AUT size or location and PNF modal content. Instead, a particular PNF modal coefficient is directly related to pattern strength in a particular direction (the direction of the one plane wave that matches the mode’s phase gradient). Thus, direct manipulation of the PNF modal (or spectral) coefficients (other than zeroing those corresponding to imaginary angles) to reject stray signals is inappropriate.

Another fundamental difference between the MARS/*IsoFilter™* approach and the holographic PNF filter is the dataset that gets filtered. MARS and *IsoFilter™* both filter the far-field output after the PNF transformation, and the holographic filter produces filtered PNF data that would then serve as the input to the PNF transformation. This distinction could be important in the performance comparison if the same PNF data are to be transformed multiple times, or especially if an aperture distribution is the only required output.

A direct comparison of timing between the modal and holographic methods is difficult because it depends greatly on several things, including scan-plane size, transform output desired, and how tightly the filtering logic is embedded within the PNF transform. In general the holographic PNF filtering is expected to be much faster. In a small set of trials, the addition of PMARS to PNF processing took over ten times longer than adding holographic PNF filtering at its input. Figure 10, when compared to Figure 6, indicates that the holographic PNF filter is equivalent to the spherical modal-filtering approach in terms of both stray-signal rejection and AUT-signal preservation.




Figure 10 Stray-signal filtering with spherical modes.

Note that the Truth trace in Figure 10 is different than in the other figures. Figure 10's Truth trace is the dataset with no stray signals injected through SNF with no filtering. Even so, its ESSL is slightly higher than that of the holographic PNF filter.

B. Single-Slice Holographic Filtering[6]

Figure 11 evaluates the AUT-signal preservation of the synthesized AUT-only PNF through a single-slice holographic filter, which is notably degraded relative to the multislice holographic PNF filter in Figure 4. The preservation can be improved (at the expense of stray filtering) by feeding a larger value in for the AUT size, but the user would have little guidance on the "proper" values to enter for AUT size and the one Z location.




Figure 11 Single-slice holographic filter, no stray signals.

The AUT-height single-slice filter is at least as good at rejecting stray signals as the multi-slice filter. This is to be expected because the multislice filter reduces the total power to be rejected in each slice. The single-slice filter is predictably faster than the multislice filter, with each slice performing a pair of 2D FFTs plus a phase propagation. With a 2Kx2K padded PNF matrix for one frequency and polarization component, the 5-slice filter took 0.9 seconds and the single-slice filter took 0.6 seconds. Its AUT-signal preservation, however, was poor.

C. Direct Low-Pass Spatial Filtering

The initial concept for this PNF-filtering effort was to exploit the predictably limited spatial bandwidth of a bounded set of radiators[1] by applying a band-centering phase modulation[1], using a digital low-pass filter in the spatial domain to reject content outside that local bandwidth, and then demodulating. This approach was quickly abandoned when the holographic approach was made to work, due to numerous perceived filtering difficulties that the holography addressed without issue:

- The filter's cutoff and number of filter taps could vary continuously with x-y probe position.
- With the varying bandwidth, the filter input's spectrum would not be stationary (and thus not shift-invariant).
- Processing time was expected to be onerous.

VI. FUTURE WORK

There are parameters in the holographic PNF filter's implementation that have not yet been thoroughly explored, notably:

A. Blanking-Taper Width and Shape

Figure 2 shows the blanking taper in use in the examples. Its shape is a raised half-period cosine, and its width on each side of the AUT area is the same width as the AUT area in that dimension. Smaller widths would improve rejection of close-in scatterers, but due to increased truncation effects might degrade the AUT-signal preservation.

It is also not known whether the ideal taper width is a percentage of AUT width, a number of wavelengths, or a combination of the two.

B. Step Size in Z

The modeling in Section III used five slices or Z positions. Figure 3 suggests that this was a good choice for the geometry modeled. The ideal number might be a function of AUT size in each dimension, probe separation, critical angle, and/or wavelength. Specifying a larger number of slices will increase processing time (about 80 ms per slice in limited trials). Specifying fewer slices risks treating AUT signal as stray signal and thus perturbing it.

VII. CONCLUSIONS

A new holographic PNF filter algorithm has been presented for filtering PNF data based solely on a known 3D AUT boundary. The filter is extremely fast because its primary operations are 2D FFTs. The new filter solves the problem of holographic defocusing when the AUT has radiators at multiple depths. The holographic PNF filter produces filtered PNF data suitable for input to the PNF transform, such that the filtering does not need to be repeated if multiple transformations are needed.

Synthesized data for a nontrivial antenna pattern in Section III were used to evaluate the filter's goals of AUT-signal preservation and then stray-signal attenuation. Both of those goals were met extremely well, and matched or exceeded the performance of other methods tested. Those combined with its speed and ease of implementation make the holographic PNF filter an attractive option.

REFERENCES

- [1] O.M. Bucci, C. Gennarelli, and C. Savarese, "Representation of electromagnetic fields over arbitrary surfaces by a finite and nonredundant number of samples," *IEEE Trans. Antennas Propag.*, vol. 46, no. 3, pp. 351-359, Mar. 1998.
- [2] S. Gregson, A. Newell, G. Hindman, M. Carey "Extension of the mathematical absorber reflection suppression technique to the planar near-field geometry," AMTA 2010, Atlanta, GA.
- [3] D.W. Hess, "The *IsoFilter*TM technique: isolating an individual radiator from spherical near-field data measured in a contaminated environment," Post-Deadline Paper, AMTA 2006, Austin TX.
- [4] S. Gregson, A. Newell, G. Hindman, "Advances in planar mathematical absorber reflection suppression," AMTA 2011, Englewood, CO.
- [5] S.T. McBride, "Stray-signal reduction using near-field modal filtering," presentation at AMTA 2015 Regional Event, Boulder, CO.
- [6] F. J. Cano-Fácila, S. Burgos, F. Martín, and M. Sierra-Castañer, "New reflection suppression method in antenna measurement systems based on diagnostic techniques," *IEEE Transaction on Antennas and Propagation*, Vol. 59, no. 3, pp. 941-949, Mar. 2011.

