Exploration of a More Conformal AUT Volume for Wide-Mesh PNF Sampling

Scott T. McBride, FELLOW, AMTA
NSI-MI Technologies, Suwanee, GA, USA

***Abstract*— The approach of non-redundant near-field sampling has been available for many years. A general and automated approach that yields the expected time reduction for an arbitrary antenna volume, however, has been elusive. One of the more practical approaches is the "PNF wide-mesh" sampling, where the probe grid is separable in x and y, and this approach is the one explored in this paper.**

A fundamental step in non-redundant sampling is to identify a volume that fully contains the AUT. Constraints imposed by theory have typically led this volume to be rotationally symmetric about a z-oriented line, and often also require that the volume be more spherical (less oblate) than a volume circumscribing the AUT. That larger volume generally results in more acquisition time than would a conformal volume, but allows those samples to be readily interpolated to the conventional half-wavelength PNF grid. This paper examines the impacts of relaxing those constraints in order to further reduce the required sampling time for a box-shaped AUT. It then looks for ways to reduce or remove those impacts.

The implementation of this algorithm involved a minor reformulation, specific to the PNF (or linear-axis) geometry, of the underlying non-redundant sampling theory. That reformulation is briefly described herein. A new family of tunable AUT-volume edge treatments similar to the existing "double-bowl" is also described.

The paper will show minor reductions in predicted acquisition time compared to non-redundant sampling with a circular double-bowl volume. Each non-redundant approach typically offers a 40-60% reduction with a rectangular AUT volume compared to a full conventional scan. A more notable advantage of the new approach is a significant reduction in preacquisition activity defining the several parameters that govern the non-redundant acquisition and processing.

Index Terms*— *non-redundant, PNF, sampling, spatial bandwidth, superellipse, wide-mesh.

I. INTRODUCTION

Planar wide-mesh scanning (PWMS) [1] is an approach that should be applicable to the vast majority of PNF systems with x-y scanners. When stepping one axis with non-redundant sampling [2] and scanning the other at full resolution as is assumed herein, acquisition time (nearly proportional to the number of scans) is typically cut by 40-60% relative to the conventional approach with full-resolution steps. Since each scan samples at a constant increment along the moving axis, the only unconventional requirement on the acquisition system is moving the step axis through a prescribed list of positions.

A. *Motivations*

We recently investigated using PWMS with a box-shaped AUT and a small scanner. The virtual addition of a rotationally symmetric (double-bowl) expansion of the AUT volume pushed that volume close to the scanner edges, such that the time savings from PWMS would be minimal. This investigation sparked an interest in seeing whether or not a rotationally symmetric volume is truly necessary and/or seeing if there are ways to mitigate any errors induced by a more conformal volume.

Another motivation is to increase the understanding and acceptance of non-redundant sampling, and especially PWMS because it is so broadly applicable and beneficial. One goal of this effort is to provide automation of and/or recommendations for parameter values that balance time savings and fidelity.

B. *Approach*

The approach described herein is specific to a box-shaped AUT. The acquisition will be sparse only along the step axis x, and the y axis will be conventionally scanned with increments close to $\lambda/2$.

Because the set of x positions will be the same for all y, those x increments must be small enough for use at any y. Toward that end we need to identify the x-y path with the worst-case bandwidth $w(x, y_{\max}(x))$.

A simulation campaign was undertaken to ensure proper sampling and reconstruction and to experiment with controlling-parameter values to gauge their effects on acquisition time and reconstruction fidelity. Efforts have also been made to automate and/or recommend parameter values.

C. *Outline*

Nyquist does not say "thou shalt sample at $\lambda/2$," but rather that "thou shalt sample at a rate higher than the local bandwidth." Thus, the notion of spatial bandwidth is core to the non-redundant approach. Section II offers an alternative description of spatial bandwidth in the PNF geometry.

PWMS traditionally calls out a rotationally symmetric fictitious volume surrounding the physical AUT volume. For a rectangular AUT, that can notably increase the computed bandwidth and the corresponding number of scans. A primary goal of this effort was to relax that symmetry requirement while still achieving adequate fidelity. Section III describes the form of the replacement "bumper" shape used here.

Section IV describes how the coarse step-axis (x) positions are computed once the spatial bandwidth has been quantified per Sections II and III.

Once the data are acquired on the PWMS grid, they need to be made ready for input to the PNF transform. Section V describes that interpolation to a regular grid.

Section VI describes a simulation campaign undertaken to implement, test, and quantify the new approaches herein described.

II. SPATIAL BANDWIDTH DESCRIBED FOR A LINEAR AXIS

When sampling in the time domain, it is well known that the sampling rate must be ≥ the temporal-frequency bandwidth or aliasing can occur[3]. Similarly, sampling in the spatial domain must use a sample rate ≥ the spatial-frequency bandwidth. What, then, is spatial-frequency bandwidth? And better yet, how does one quickly understand it without dozens of equations? We'll start with a couple simple equations and then offer examples in what might be more familiar scenarios.

The equation for spatial bandwidth $w(x,y,z)$ [2] (with slight modification herein) is defined as (1):

$$w(x,y,z,\hat{t}) = \max\left(\hat{R} \bullet \hat{t}\right) - \min\left(\hat{R} \bullet \hat{t}\right), \tag{1}$$

where

(x,y,z) is the probe location,

$\hat{t}$ is the probe-travel unit vector (always equals $\hat{x}$ or [1, 0, 0] in our chosen geometry),

$\hat{R}$ is the collection of unit vectors from the probe location to candidate points on the AUT volume, and

$\hat{R} \bullet \hat{t} = \delta\phi/\delta x = R_x/|R| = K_x/K$ (stated without proof) and has values in the range ±1.

We see that $\hat{R} \bullet \hat{t}$ has multiple equivalent interpretations. For one dot product with a value within ±1, either phase slope $\delta\phi/\delta x$ or K_x/K in units of "periods of phase per wavelength traveled" seems descriptive. However, the difference w between max and min will be a value between 0 and 2, and more than one phase period per wavelength is physically implausible. It seems less confusing to avoid assigning units to w, and instead think of w as the span of K_x/K.

A. Behavior through the FFT, sample spacing ⇔ bandwidth

If time-domain data are sampled at constant interval Δt in seconds and passed through an FFT, then the FFT output will represent a temporal-frequency span or maximum reported bandwidth in Hz (cycles/second) equal to $1/\Delta t$.[3].

If spatial-domain data are sampled at constant interval Δx in linear units and passed through an FFT, then the FFT output (with each value quantifying the contribution of a particular phase slope $\delta\phi/\delta x$) will represent a spatial-frequency (K_x/K) span or maximum reported bandwidth w equal to $1/(\Delta x/\lambda)$ or $\lambda/\Delta x$. Inverting that relationship, we see that we can sample at a rate ≤ local bandwidth w by following (2):

$$\Delta x \leq \frac{\lambda}{w}. \tag{2}$$

Just as the bandwidth of a time-based signal is defined as the width of the temporal-frequency spectrum where power is present, so is the spatial bandwidth defined as the width of the spatial-frequency spectrum where power is or can be present.

B. Geometric description

Figure 1. attempts to illustrate the determination of worst-case spatial bandwidth of the conformal box at three probe positions (over the aperture, just beyond the aperture, and well beyond the aperture) using simple dot products[2]. In Figure 1. the (black) rectangles represent the front (z=0) and back (-z) of the AUT box, the long (green) line represents probe travel parallel to x with a (green) dot at the illustrated probe position, the skinny (red) cone shows first contact with the box expanding outward from behind the probe (min $\hat{R} \bullet \hat{t}$), and the larger (blue) cone does the same expanding outward from ahead of the probe (max $\hat{R} \bullet \hat{t}$). Multi-color dots are shown at the two contact points for each plot. Local bandwidth $w(x)$ is then the difference between the max and min dot products consistent with (1).

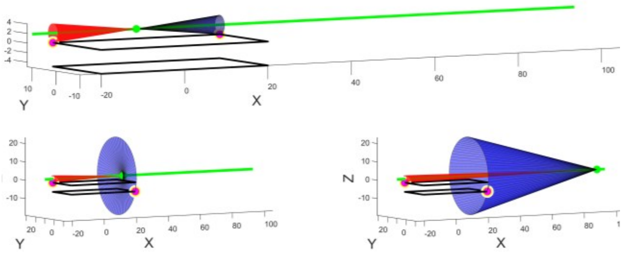

Figure 1. Min and max dot products

Figure 2. shows with its top trace a typical local spatial bandwidth $w(|x|)$ for a box-shaped AUT volume. Here we see the high bandwidth when centered over the aperture, diving down to unity at the aperture edge, and the slower decay after the probe starts seeing the back of the AUT as $\max(\hat{R} \bullet \hat{t})$. The discontinuity in the bandwidth's first derivative is also labeled and is the reason for artificially expanding the AUT volume as discussed in Section III.

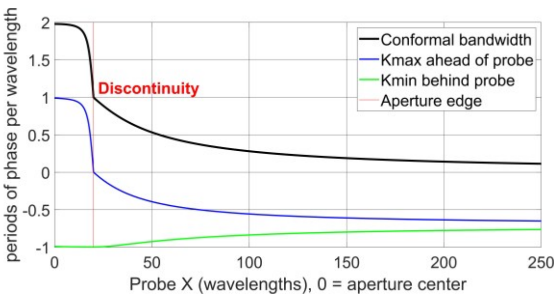

Figure 2. Spatial bandwidth of box shape

In summary, the "spatial bandwidth" of a field to be sampled along a linear axis is the width, range 0 to 2, of the spatial-frequency domain K_x/K where power is possible. Thus (2) provides a straightforward relationship between spatial bandwidth and the sample increment. If one is using a fixed increment at all x, and if the probe passes close to the aperture center where local bandwidth w is close to 2.0, then that fixed sample increment should be, per (2), $\Delta x \leq \lambda/2$.

C. *Accounting for Probe Y*

In PWMS, the same set of x positions will be sampled for all probe y positions. It is thus important to evaluate bandwidth $w(x, y_{\max}(x))$ as the worst-case locus through the x-y probe plane, where $y_{\max}(x)$ is the probe-y location at each probe x that sees the maximum local bandwidth with $\hat{t} = \hat{x}$. (Note that we continue to use $\hat{x}$ because subsequent sampling and interpolation will be done along x.) Figure 3. shows conformal $w(x,y)$ computed for three different AUT-box aspect ratios, with $w(x,y)$ depicted as colorfill and contours, AUT box as a magenta rectangle, and $\pm\,|y_{\max}(x)|$ as nearly diagonal black on yellow lines. Figure 3. indicates that if $y(x)$=0 (where conformal bandwidth is at its minimum) were used to find sample spacing, then the field along x at nonzero y would be undersampled. An early hypothesis was that $y_{\max}(x)$ would be found along the diagonals $y_{\max}(x) = \pm|\,x\,|$, but Figure 3. shows that to be true only when the AUT box is a square. The approach taken here is to tabulate $y_{\max}(x)$ from the 2D conformal bandwidth and use that later to find the (wider) 1D sampling bandwidth.

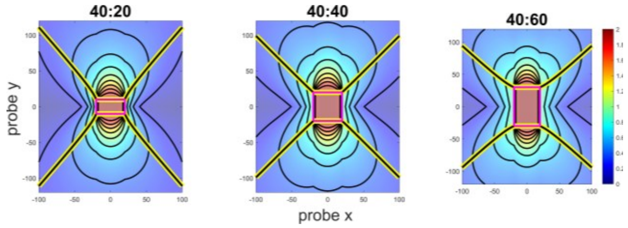

Figure 3. Conformal spatial bandwidth vs. x and y (with $\hat{t} = \hat{x}$)

Figure 2. shows the local spatial bandwidth $w(x)$ when we follow $y_{\max}(x)$ in that locus. This plot has some features worthy of note (Figure 1. might be helpful in visualizing these notes):

- Small $|x|$ (probe over aperture center) has $w(x)$ very close to 2, so the local Nyquist increment is close to $\lambda/w = \lambda/2$.
- x at the aperture edge has conformal $w(x)$ very close to 1, so the local Nyquist increment is close to λ.
- As $|x|$ passes beyond 20 (the edge of our conformal AUT box), something bad happens. At x=19.999 there is one distinct forward point of first contact at Δx=0.001, Δy=0, and Δz=probe separation. At x=20.001, that point of first contact jumps to the far corner of the x=+20 side, suddenly increasing the $|R|$ in the maximum dot product $R_x/|R|$.

The sudden change is 'bad' because a discontinuity got introduced into our $w(x)$ curve's first derivative, annotated in Figure 2. We'll be doing bandlimited interpolation later in the process, and that requires that our samples be equally spaced along a bandlimited path. Discontinuities at least in early derivatives of $w(x)$ tend to unacceptably degrade the approach. This bad behavior is typically avoided[1][2] by surrounding the conformal AUT volume with a less-conformal but better-behaved shape. Section III discusses the less-conformal shape introduced here.

III. A NEW AUT-BOUNDING VOLUME

There are multiple competing goals when choosing a bounding volume:

- Provide a smooth curve $w(x)$ everywhere, especially at the aperture edge.
- Yield a local-bandwidth curve $w(x)$ that is everywhere ≥ the conformal curve.
- Yield a $w(x)$ curve that is close to the conformal curve so that it minimizes the number of samples along x.

The conventional approach [1][2] to this bounding shape is a rotationally symmetric analytical volume. These have included spheres, ellipsoids, and double bowls. Figure 4. shows cross sections of the double bowl[1] (top) and the new bumper (bottom).

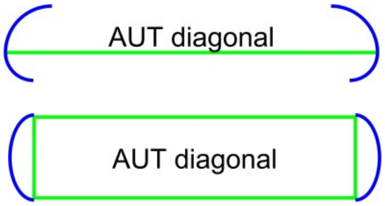

Figure 4. Cross sections of double bowl (top) and new bumper (bottom)

The double bowl's cross section is rotated through a circle to form its volume. The new bumper's extrusion is along the locus of a superellipse[4], whose equation is (3):

$$\left(\frac{x}{a}\right)^N + \left(\frac{z}{b}\right)^N = 1\,, \tag{3}$$

where

a is the value of x when z=0,

b is the value of *z* when *x*=0, and

N is the real-valued order of the superellipse, ≥ 2.0.

Figure 5. shows superellipses orders 2, 3, and 6 (order 2.0 is a regular ellipse) scaled to circumscribe the same rectangular box. The superellipse was chosen to allow controlled growth of the new volume away from the conformal box, with the assumption that excessive growth would increase the bandwidth and thus the number of required samples. We see that the higher orders become more and more conformal (fewer scans) while still yielding a convex shape, though the derivative content is becoming large (possible bandwidth expansion).

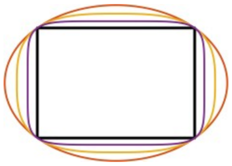

Figure 5. Order 2, 3, 6 superellipses circumscribing a box

Figure 6. then shows the elliptical bumper extruded around a superelliptical path. In this example, the aspect ratio of the superellipse has been chosen to be different from that of the box it circumscribes.

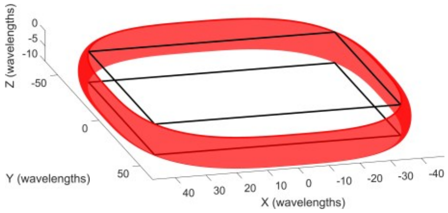

Figure 6. Illustration of bumper for *w*(*x*)

Figure 7. illustrates the loci of min (left) and max (right) dot products $\hat{R} \bullet \hat{x}$ as the probe moves along the quadrant-1 diagonal line (with *y*(*x*) chosen to maximize the bandwidth *w(x)*). Only the portions of the bumper where initial contact is possible are shown. Blue dots are shown at those initial-contact points, and thin lines connect (at coarse intervals) probe positions and corresponding blue dots. Geometric symmetry is exploited, reusing the quadrant-1 values to populate the other three quadrants.

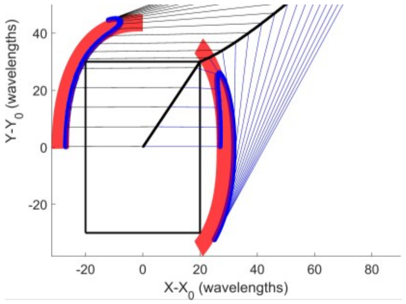

Figure 7. Loci of dot-product hits on the new bumper

Figure 8. again shows the conformal bandwidth and adds an overlay of the smoother bandwidth computed with the bumper, both from the geometry in Figure 7.

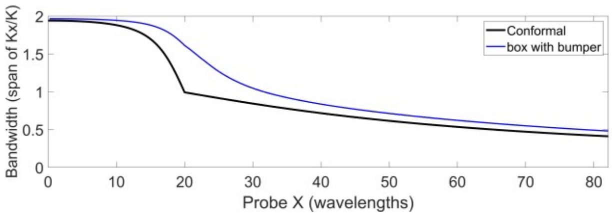

Figure 8. Spatial bandwidth *w*(*x*) with bumper

IV. SAMPLING PER THE SPATIAL BANDWIDTH

A key feature of non-redundant sampling[2] is its recognition that the local spatial bandwidth changes with probe position and that the local sample increment can be made larger as local bandwidth *w*(*x*) decreases.

When periodically sampling a signal in the time domain, Nyquist typically requires that the sample rate (in samples/second) be ≥ the signal's bandwidth *w*(*t*)[3] (in cycles/second). The same principle applies to spatial sampling, where the sample rate (in samples/wavelength) must be ≥ the signal's local spatial bandwidth *w*(*x*).

We are going to use bandlimited interpolation when processing the sampled data, which benefits greatly from a constant increment along some bandlimited domain ξ(*x*). The purpose of the virtual edge treatment around the AUT volume is to limit any bandwidth expansion due to the shape of *w*(*x*). With the conformal box shape, the bandwidth expansion to be mitigated results from the derivative discontinuity annotated in Figure 2.

In defining ξ(*x*) we are going to integrate sample rate (bandwidth *w*(*x*)) per (4):

$$\xi(x_f) = \int_0^{x_f} \frac{\chi w(x)}{\lambda_{min}}\, dx, \tag{4}$$

where χ is an oversampling factor > 1.0 that causes the sample rate to be greater than the bandwidth.

Our variable spacing along x will correspond to uniform spacing along ξ(*x*). For simplicity, we choose to scale ξ(*x*) to represent "sample numbers", and every integer value of ξ(*x*) will correspond to a sample location. Toward that end, we first define a vector x_f of finely spaced *x*, compute bandwidth w(x_f) per (1), scale by χ/λ (at max frequency), and integrate per (4) to get $\xi_f = \xi(x_f)$.

The integral in (4) is accomplished via simple summation. Because vectors ξ_f and x_f are each ever-increasing, we can invert and interpolate to obtain the unequally spaced x_s sampling values with (5):

$$x_s = \text{interp}(\xi_f,\ x_f, 1{:}\,\xi_{max}), \tag{5}$$

where

ξ_{max} = int(max($\xi(x_f)$)), and

x_s contains sample positions at +x with a sample at and odd symmetry about the aperture center.

Later, during interpolation to the conventional grid, we will need to know the ξ_{out} values for each of the equally spaced output locations x_{out}. Those values can be found using (6):

$$\xi(x_{out}) = \text{interp}(x_f, \xi(x_f),\ x_{out}). \tag{6}$$

V. INTERPOLATION TO CONVENTIONAL GRID

The implementation described herein is sparse only in x, with conventional scans in y with Δy near $\lambda/2$. Each row of the acquired matrix shares the same set of x_s sample locations, and each row needs to be interpolated to the conventional grid for transformation[2]. The interpolation has three fundamental stages: geometric phase modulation to center the bandwidth at zero; bandlimited interpolation along ξ from integer sample locations to $\xi(x_{out})$; and geometric phase demodulation to undo the earlier modulation.

A. Phase modulations (the magic that enables non-redundant sampling)

When defining the x-axis step positions in Section IV, $\Delta x(x)$ was based on local $\max(\delta\phi/\delta x) - \min(\delta\phi/\delta x)$, and that almost never yields a bandwidth centered at $\delta\phi/\delta x = 0$. This yields $\Delta x >> \lambda/2$ outside the aperture bounds, which confuses the sequence of phase measurements vs. x in a very predictable way such that measured I and Q cannot be directly interpolated. The theory[2] exploits the predictability of this problem, applying a phase modulation proportional to the integrated local slope $\psi(x) = (\max(\delta\phi/\delta x) + \min(\delta\phi/\delta x))/2$ for each acquired x-y position prior to any interpolation. A corresponding demodulation is applied after interpolation based on the equally spaced output locations in x and y.

PWMS acquires data on a separable grid of x and y, such that the spatial bandwidth and its resulting x-axis increments are appropriate for the worst-case y. That restriction does not apply to the phase modulations, which can and should be computed vs. x and y.

The algorithm described herein computes the phase modulations based on a simplified bumper to maximize processing speed. Using a conformal box for phase yielded (as expected) reconstruction errors near the aperture edges due to the derivative discontinuities shown in Figure 2. The simplified phase bumper, which heavily constrains the locus of dot-product hits, is illustrated at 5 y positions in Figure 9. The phase bumper is a single tilted half sine (magnitude = 3D bumper radius) redefined at each y position to connect min or max x at the box front (green dot) to the far back corner (blue dot). The modulating phases are computed from the average of min and max dot products $\hat{R} \bullet \hat{x}$ over x travel[2] at each value of y. Near-field reconstruction using this phase-modulation approach yielded a notable reduction in ESSL vs. the initial conformal approach. The per-acquisition phase-computation time with this approach was consistently below one second. Use of the bumper in Figure 7. to compute phase modulations has not been tried, but was crudely estimated to take several minutes.

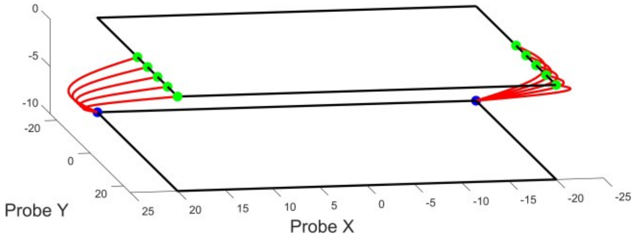

Figure 9. Illustration of phase bumper for $\psi(x)$

B. Bandlimited interpolation

The equally spaced x positions in the conventional grid will rarely correspond to the integer values along ξ where data were acquired such that interpolation is required. Following the phase modulation, we have a bandlimited signal ($\pm w(x)/2$) equally spaced at integer values along ξ. Bandlimited interpolation to ξ between the samples is done as a phase taper through the spectral domain[3][2]. That interpolation works best when the two center samples of the fft's (or equivalent summation's) input straddle the desired output ξ. The number of measured samples is limited, and windowing[5] of the interpolation input is required[2], along with windowing compensation of the interpolation output.

In this implementation the number 2P of samples in each interpolation dynamically changes with each ξ position. At the min and max x positions of the output grid, P equals the number of guard columns (typically 2). As $\xi(x)$ passes the next integer value, P gets incremented by one up to a user-specified maximum, or reduced by one if there are now fewer than P samples ahead of the current x.

After the data are interpolated to the conventional grid, the phase slope that was modulated out at the sample locations is demodulated back in at the output locations.

C. Spatial filtering

As a final step in estimating the conventionally spaced samples, the interpolated data are run through a single pass of the holographic PNF filter[6]. The filter's most notable impact was in lowering the near-field ESSL at the aperture edges, typically by about 5-8 dB.

VI. SIMULATION RESULTS

A. Modeled AUT

A highly unorthodox AUT model (very similar to one used in testing the holographic PNF filter) was chosen as a challenging test of non-redundant sampling. One 2D dipole array steered in azimuth at +70° fills the front (z=0) surface of the AUT box. A second 2D dipole array steered at -50° fills the rear surface (z=-12λ) of the box. Excitation windows and randomization are applied to each array to provide a nontrivial pattern to reconstruct. No mutual coupling or occlusion is modeled among any of the dipoles. The probe is modeled as another dipole, and no probe correction is performed. The zeroelevation pattern (as transformed from the conventional grid with a cosθ factor) is shown as the Truth trace in Figure 10.

B. *Truncation mitigation*

Non-redundant sampling generally requires extra “guard samples” at each end of any sparse row or column to enable accurate interpolation near those ends. This simulation campaign found that two guard samples at each end were sufficient provided that those ends were at least 15 wavelengths beyond the aperture edge.

C. *Reconstruction fidelity*

Figure 10. shows expected fidelity with a nearly conformal bumper (order 10.0, radius 0.1λ). A conformal volume is expected to yield best-case time savings and worst-case fidelity. Scan count (and thus acquisition time) is 69% lower than the equivalent conventional acquisition. This fidelity might be sufficient for many applications, but there is certainly room for improvement. The mean ESSL voltage magnitude for the one displayed centerline scan is -63 dB.

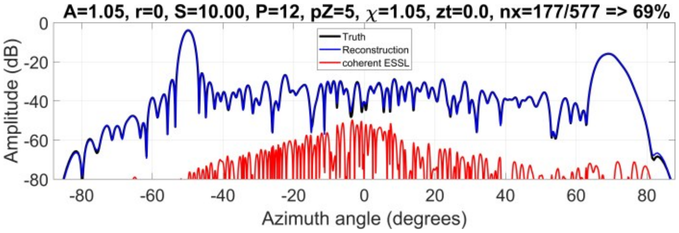

Figure 10. FF centerline, volume nearly conformal

Figure 11. shows the same analysis when a much more relaxed bumper (order 2.0, radius 6λ) is specified. The scan count is 64% less than conventional, though 32 (or 18%) more scans than in Figure 10. The mean ESSL for the one displayed centerline scan is -82.9 dB.

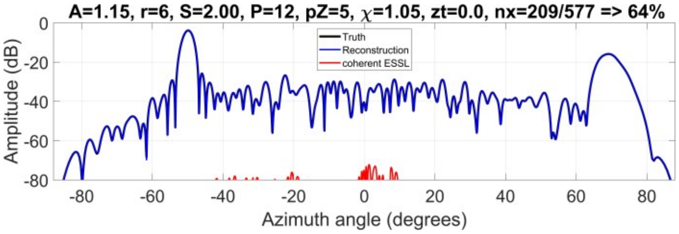

Figure 11. FF centerline, volume nearly double-bowl

The primary trade space is between a low number of y scans (proportional to acquisition time) and low ESSL. There are four main algorithm parameters that control the locus through that trade space: the oversampling factor χ, the elliptical bumper’s radius (typically its minor axis), the order of the superellipse upon which the elliptical bumper is extruded, and the axial ratio of the superellipse. Increases in χ increase the bandwidth estimate everywhere, and one goal is to achieve more targeted control, so χ was always set to 1.05 herein. The optimal axial ratio varies with selected critical angle, and is automatically set to minimize the number of scans. Thus the primary parameters are reduced to the bumper radius and order.

Figure 12. quantifies the PWMS results from surrounding one hypothetical AUT (z depth = 12λ) with each of 10 bumpers. Three ellipse radii (3, 6, and 9λ) and three superellipse orders (2, 4, and 6) were modeled, plus the ‘conformal’ bumper from Figure 10. with order 10 and radius 0.1 λ. Here we see trends of azimuth-pattern fidelity improvement with increased bumper radius and/or decreased superellipse order, and that higher fidelity generally requires more step positions or y scans. Findings from this effort are shown in Table I.

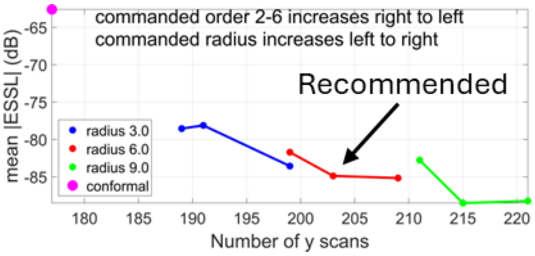

Figure 12. Uncertainty and time vs. order and radius

TABLE I. SELECTED FINDINGS

Topic or Parameter	Finding or Recommendation
Rotational symmetry	Not required
Order of superellipse	4.0
Bumper radius	Max(3λ, ½ AUT depth in z), same size for *w* and ψ
#guard columns	2, provided truncation > 15λ from aperture edges
Superellipse axial ratio	Automatically determined to minimize #scans
Oversample-factor $\chi\chi$’	1.05 sufficient, < 1.00 fails
Probe Z	More is better until critical angles > limits
Time savings	Slight improvement vs. rotationally symmetric approach. Percentage grows in either approach as critical angles get larger.

VII. CONCLUSIONS

PWMS is an approach that should be applicable to the vast majority of PNF systems with x-y scanners. This paper has introduced new fictitious AUT volumes that offer additional advantage for antennas without rotational symmetry. Several controlling parameters are available, with recommendations supplied in Table I. for each. This approach is expected to lower the number of scans by 40-60% vs. conventional sampling with typically specified critical angles.

REFERENCES

[1] F. D'Agostino, F. Ferrara, C. Gennarelli, R. Guerriero, G. Riccio, “An effective antenna modelling for the NF-FF transformation with planar wide-mesh scanning,” AMTA 2005.

[2] 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 Propagat., vol. 46, pp. 351-359, 1998.

[3] A.V.Oppenheim and R.W.Schafer, Discrete-Time Signal Processing. New Jersey: Prentice Hall, 1989.

[4] Wikipedia contributors. "Superellipse." Wikipedia, The Free Encyclopedia, 25 Dec. 2024.

[5] F.J. Harris, “On the use of windows for harmonic analysis with the discrete Fourier transform,” Proceedings of the IEEE, vol. 66, no. 1, 1978, pp. 51-83.

[6] S.T.McBride, “Truncation mitigation using the holographic PNF filter,” AMTA 2024, Cincinnatti, OH.