Reiterative MMSE using Feed-Forward Prior Estimates for Improved Direction Finding

Logan Satterfield[1], Jonathan Owen[1], Alex Bouvy[2], Benjamin Kirk[2], Patrick McCormick[1], Shannon Blunt[1]
[1] Radar Systems Lab, University of Kansas, Lawrence, KS
[2] Army Research Laboratory, Adelphi, MD

***Abstract*— Reiterative super-resolution (RISR) is a variant of the reiterative minimum mean squared error (RMMSE) algorithm class, originally developed for adaptive direction finding. RISR was recently experimentally demonstrated to enhance open-air direction-of-arrival estimation while providing robustness to non-ideal calibration errors via practical modeling and incorporation of a gain constraint. RISR is generally initialized with a standard (non-adaptive) beamforming estimate.**

Here, the recycled estimate (RE)-RISR variant is proposed that instead uses RISR angle estimates from recent snapshots as an initialization, thereby avoiding the need for complete reconvergence at each snapshot. Given sufficient (yet modest) stationarity, RE-RISR significantly reduces the number of required iterations (and therefore computational cost). Simulated performance for a 6-element uniform circular array (UCA) with stationary and nonstationary incoming signals reveals robustness even in dynamic scenarios.

Keywords—array processing, direction finding, uniform circular array, minimum mean squared error (MMSE)

I. Introduction

Direction Finding (DF) is the process of determining the spatial angles from which received signals impinge on a sensor array. Several signals may arrive simultaneously, and some may possess high temporal correlation (e.g. multipath). Classical DF techniques such as Capon beamforming, MUSIC, and ESPRIT [1]-[6] rely on the determination of a sample covariance matrix (SCM) from a set of spatial snapshots, which implicitly assumes a degree of stationarity.

The RISR algorithm [7], [8] employs the reiterative minimum mean-square error (RMMSE) framework to perform adaptive DF via a structured covariance matrix that avoids an SCM, thus making it suitable when low sample support is required due to nonstationarity. The partially constrained RISR (PC-RISR) form [9] was subsequently considered to realize super-resolution in the low SNR regime by incorporating tunable selectively between standard RISR and a gain constrained version [10]. The PC-RISR form has been experimentally demonstrated to perform well on open-air data for both DF and fast time-frequency estimation [11], [12].

The recycled estimate RISR (RE-RISR) variant is a modification of PC-RISR that replaces the standard (non-adaptive) beamforming initialization with an adaptive DF estimate from the previous spatial snapshot in time. Consequently, convergence can be achieved with far fewer iterations per snapshot in stationary conditions. In quasi-stationary or nonstationary cases, only mild degradation is observed that is proportional to the degree of nonstationarity. A collection of signals having modest stationarity in received angle over time yields DF performance that is indistinguishable from PC-RISR, achieved with far lower computational cost. Conceptually, dynamic modification of the number of iterations per snapshot for RE-RISR (i.e. only when needed due to perceived changes) would also further improve tolerance to nonstationarity, with the caveat of increased computational cost on an as-needed basis.

The uniform circular array (UCA) provides a simple geometry to achieve 360° DF with limited antenna elements [13], though the associated array manifold is less advantageous due to high sidelobes. Here, both the RISR and RE-RISR forms are applied in simulation to a UCA having six antenna elements, demonstrating significant improvements in angle estimation over standard beamforming. RE-RISR estimation is compared to that of RISR, both qualitatively and quantitatively, for different signal arrangements.

II. UCA Steering Matrix

Consider the array factor for an arbitrary array geometry

$$\mathrm{AF}(\theta,\phi) = \sum_n w_n \mathrm{e}^{\,\mathrm{j}\mathbf{c}_n^T\mathbf{k}(\theta,\phi)} \tag{1}$$

where w_n is the weight of the n^{th} antenna element, and θ,ϕ are elevation and azimuth, respectively. Given a narrowband spectrum assumption, the wave number (or *K*-space) spherical-to-Cartesian transform vector is

$$\mathbf{k}(\theta,\phi) = \frac{2\pi}{\lambda_\mathrm{c}}\begin{bmatrix}\sin\phi\cos\theta\\ \cos\phi\cos\theta\\ \sin\theta\end{bmatrix} \tag{2}$$

and Cartesian coordinates of the n^{th} antenna element are

$$\mathbf{c}_n = \begin{bmatrix}x_n\\ y_n\\ z_n\end{bmatrix}. \tag{3}$$

Setting $z_n = 0$ (for a UCA on the *x-y* plane), the steering vector for an incoming signal arriving from angles θ,ϕ is represented by

$$\mathbf{s}(\theta,\phi) = \mathrm{e}^{\,\mathrm{j}\mathbf{C}^T\mathbf{k}(\theta,\phi)} \tag{4}$$

where $\mathbf{C} = [\mathbf{c}_1,\cdots,\mathbf{c}_N]$ is the matrix of Cartesian array positions. For all M look angles of interest (where $M \gg N$), the steering vectors $\mathbf{s}(\theta,\phi)$ are compiled into the columns of the $N \times M$ steering matrix $\mathbf{S}$ as

This work was supported by the Army Research Laboratory subcontract to Alion S&T for prime contract #FA8075-18-D-0002.
DISTRIBUTION STATEMENT A. Approved for public release.

$$\mathbf{S}(\theta,\phi) = [\mathbf{s}(\theta_1,\phi_1) \quad \mathbf{s}(\theta_2,\phi_2) \quad \cdots \quad \mathbf{s}(\theta_M,\phi_M)]. \tag{5}$$

III. Reiterative Super-Resolution Review

Consider the single snapshot receive model for an N element arbitrary array at discrete time index ℓ, denoted as

$$\begin{aligned}\mathbf{y}(\ell) &\triangleq \left(\mathbf{Sx}(\ell)\right) \odot \mathbf{z} + \mathbf{v}(\ell) \\ &= \mathbf{Sx}(\ell) + \mathbf{v}(\ell) + \mathbf{v}_{\mathrm{z}}(\ell),\end{aligned} \tag{6}$$

where $\mathbf{x}(\ell)$ defines the $M \times 1$ signal arriving from each possible spatial direction, $\mathbf{v}(\ell)$ is an $N \times 1$ vector of additive noise, and $\mathbf{y}(\ell)$ is the received signal on the antenna array for a given time snapshot. The vector $\mathbf{z}$ accounts for modeling errors and has nth element represented as

$$z_n = \left[1 + \Delta_{\mathrm{a},n}\right] \mathrm{e}^{\mathrm{j}\Delta_{\varphi,n}}, \tag{7}$$

where $\Delta_{\mathrm{a},n}$ is an amplitude error and $\Delta_{\varphi,n}$ is a phase error, both with arbitrary distributions. This formulation permits calibration tolerances to be accommodated. It is shown in [8] that minimizing the MMSE cost function

$$J = \mathbb{E}\{\|\mathbf{x}_k(\ell) - \mathbf{W}_k^H(\ell)\mathbf{y}(\ell)\|_2^2\} \tag{8}$$

yields the $N \times M$ filter bank $\mathbf{W}_k(\ell)$, having the mth column

$$\mathbf{w}_{m,k}(\ell) = \mathbf{P}_{m,m,k}(\ell)\left(\mathbf{SP}_k(\ell)\mathbf{S}^H + \mathbf{R} + \mathbf{R}_{\mathrm{z},k}(\ell)\right)^{-1}\mathbf{s}_m \tag{9}$$

The $M \times M$ diagonal matrix $\mathbf{P}_k(\ell)$ contains the incident powers for the M spatial directions for the ℓ^{th} snapshot, the $N \times N$ matrix $\mathbf{R}$ is the noise covariance, and the $N \times N$ matrix $\mathrm{R}_{\mathrm{z},k}(\ell)$ is the calibration error covariance defined in [8] as

$$\mathbf{R}_{\mathrm{z},k}(\ell) = \mathbf{I}\sigma_{\mathrm{z}}^2 \odot [\mathbf{SP}_k(\ell)\mathbf{S}^H], \tag{10}$$

for $\mathbf{I}$ an $N \times N$ identity matrix and σ_z^2 the calibration error variance. The terms $\mathbf{s}_m$ and $\mathbf{P}_{m,m,k}(\ell)$ specify the mth column of $\mathbf{S}$ and the mth diagonal element of $\mathbf{P}_k(\ell)$, respectively. Because the values in $\mathbf{P}_k(\ell)$are not initially known, they must be estimated. Traditionally, RISR applies the standard (non-adaptive) beamformer

$$\hat{\mathbf{x}}_0(\ell) = \mathbf{S}^H\mathbf{y}(\ell), \tag{11}$$

to obtain the matrix of power estimates

$$\widehat{\mathbf{P}}_k(\ell) = [\hat{\mathbf{x}}_k(\ell)\hat{\mathbf{x}}_k(\ell)] \odot \mathbf{I}. \tag{12}$$

Using (9), the adaptive filter bank can be computed and applied to revise the estimates as

$$\hat{\mathbf{x}}_k(\ell) = \mathbf{W}_k^H(\ell)\mathbf{y}(\ell) \tag{13}$$

that are subsequently used to update (12) for the k^{th} of K iterations. This process of applying (9), (13), and (12) is repeated until acceptable convergence is achieved. The partial gain constraint from [9] can be incorporated to enable selectivity via

$$\mathbf{w}_{\mathrm{PC},m,k}(\ell) = \left[\left(\frac{1}{\mathbf{s}_m^{\mathrm{H}}\mathbf{D}_k(\ell)\mathbf{s}_m}\right)^{\alpha}\left(\mathbf{P}_{m,m,k}(\ell)\right)^{1-\alpha}\right]\mathbf{D}_k(\ell)\mathbf{s}_m\,, \tag{14}$$

in which

$$\mathbf{D}_k(\ell) = \left(\mathbf{SP}_k(\ell)\mathbf{S}^{\mathrm{H}} + \mathbf{R} + \mathbf{R}_{\mathrm{z},k}(\ell)\right)^{-1}, \tag{15}$$

for exponent $0 \le \alpha \le 1$ the weighting factor that controls the balance between unconstrained RISR ($\alpha = 0$) and fully constrained RISR ($\alpha = 1$).

IV. Recycled Estimate RISR (RE-RISR)

The modification to implement RE-RISR is straightforward. Rather than initializing the adaptive filter bank with the standard beamformer in (11) for each new snapshot in time, we can instead initialize using the adaptive filter bank from the most recent snapshot as

$$\hat{\mathbf{x}}_0(\ell) = \mathbf{W}_K^H(\ell-1)\mathbf{y}(\ell). \tag{16}$$

Provided that snapshot $\mathbf{y}(\ell)$ is sufficiently similar to snapshot $\mathbf{y}(\ell-1)$, this initialization eliminates redundant processing that would otherwise occur in successive snapshots. Note that this initial estimate of $\hat{\mathbf{x}}_0(\ell)$ also updates $\widehat{\mathbf{P}}_0(\ell)$ and therefore $\mathbf{SP}_0(\ell)\mathbf{S}^H$, improving upon the initial covariance estimate. Further, the power estimation matrix in (12) is modified for all iterations to include a perturbation term as

$$\widehat{\mathbf{P}}_k(\ell) = [\hat{\mathbf{x}}_k(\ell)\hat{\mathbf{x}}_k(\ell) + \varepsilon] \odot \mathbf{I}. \tag{17}$$

The term ε is a small perturbation to avoid passing forward sparse a-priori estimates (which may occur for small partial constraints $\alpha \le 0.5$) that would otherwise cause incorrect convergence for subsequent iterations.

Parallels may be drawn between RE-RISR and the Kalman filter, both of which rely on MMSE and an information feed-forward mechanism [14]. The incorporation of the perturbation term ε is inspired by Kalman filter process noise, which prevents overconfidence in the model [14]. That said, distinctions exist, most critically the benefit of reiterative processing performed in the RMMSE framework to refine the current snapshot estimate without additional observations.

V. Simulated Results

To examine the behavior of RISR (initialized with the standard beamformer) versus RE-RISR, a 6-element UCA with element spacing $d = 16$ cm and radius $r = 16$ cm is simulated for two scenarios. The selected UCA phase center locations are illustrated in Fig. 1. The adjacent element spacing corresponds to $\lambda/2$ spacing for $\lambda = 32$ cm or frequency $f_{\mathrm{a}} = 937.5$ MHz, which sets the array effective max operating frequency. While $\lambda/2$ spacing meets the spatial Nyquist criterion for ULAs, the UCA phase centers inherently lie on a non-uniform positional grid for Cartesian coordinates. Consequently, $\lambda/2$ spacing for a UCA implies a lower-bound on the maximum operable frequency of the array for all look angles assuming a planar incoming wavefront.

All incoming signals are selected to arrive with $\theta = 0°$ such that only azimuth is estimated in all directions. In practice, non-zero elevation θ can only be discriminated within the limitations of the array factor in that dimension. The beampattern of the UCA is modestly invariant across azimuthal look angles, such that ensuing evaluation of DOA estimation performance is generalizable. Further, due to the rotational symmetry of the array, the beampattern repeats every 60° in azimuth under ideal calibration conditions [15].

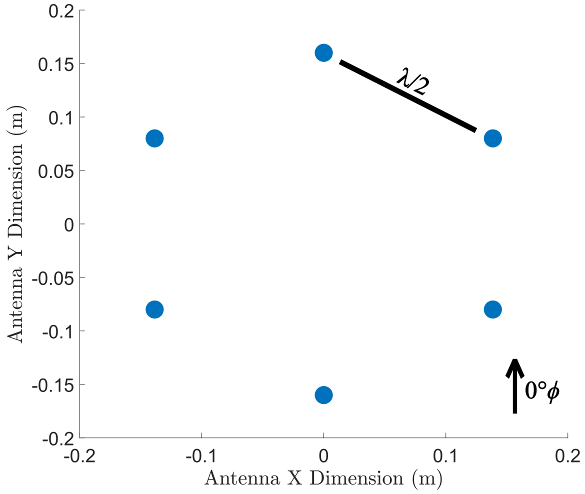

Fig. 1. Six-element UCA antenna pattern, with element spacing $d = \lambda/2 = 16$ cm and radius $r = 16$ cm. Note that $\phi = 0°$ corresponds to the positive y direction.

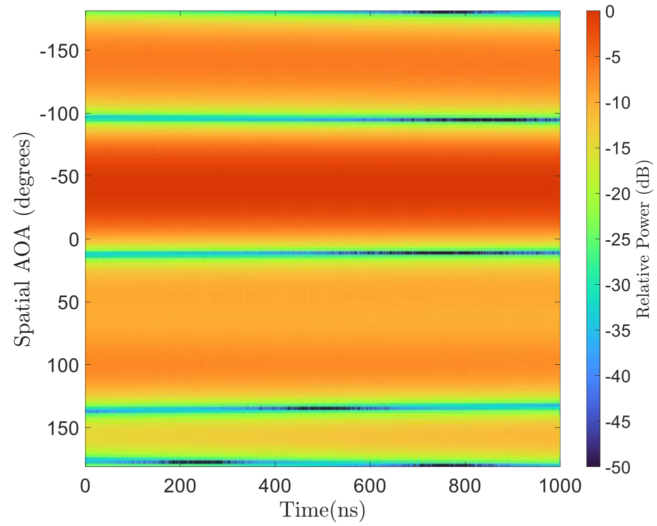

Fig. 2. Case A; Magnitude of the standard matched beamforming estimate over time. The chirp waveform direction-of-arrival is properly estimated over time, but significant spatial mainlobe roll-off and sidelobes are present.

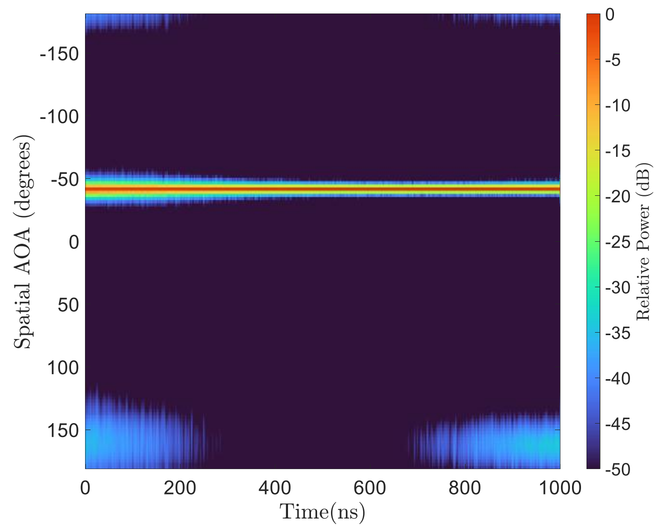

Fig. 3. Case A; Magnitude of RISR estimate over time, with 10 iterations per snapshot. The signal direction-of-arrival is clearly defined at the true incoming angle, with minimal error.

The first scenario (Case A) is intended to demonstrate baseline performance, where an LFM chirp is incident at -40 degrees with 40 dB signal-to-noise ratio (SNR). The second scenario (Case B) demonstrates robustness to severe nonstationary effects and modest SNR discrepancies between incoming signals, where a pulsed chirp waveform with 50% duty cycle is incident at -40 degrees with 40 dB SNR and an OFDM signal arrives at 30 degrees with 10 dB SNR. The DF estimates produced by standard beamforming, RISR with 10 iterations, RISR with 3 iterations, and RE-RISR with 3 iterations are shown for qualitative comparison. Because the RMMSE framework attempts to minimize mean squared error (MSE), it is the natural metric for relative quantitative performance evaluation. All incoming signals are simulated with a center frequency $f_c = 800$ MHz, bandwidth of $B = 50$ MHz, and max frequency $f_c + B/2 < f_a$. For both the RISR and RE-RISR cases, the partial gain constraint is set to $\alpha = 0.65$.

For Case A, the DF estimates over time for each algorithm are shown in Figs. 2-5 and their achieved MSE relative to ground truth over time are shown in Fig. 6. Specifically, Fig. 2 demonstrates the standard beamforming estimate for a single incident LFM chirp. The signal's incoming spatial angle is correctly estimated, but the relatively broad mainlobe makes precise estimation of the true angle difficult. Additionally, the significant spatial sidelobes would easily overpower any modestly lower energy signals that could otherwise be present. Standard spatial windowing functions for circular arrays may be applied to suppress spatial sidelobes [16], at the cost of further mainlobe broadening and SNR mismatch loss. The ambiguity observed due to mainlobe rolloff and sidelobes is indicated by significant MSE (the purple trace) in Fig. 6.

Applying RISR with 10 iterations produces the DF response in Fig. 3, which well estimates the incoming signal and provides a significant decrease in mainlobe width relative to the standard beamformer. The achieved super-resolution is a consequence of an "angle adaptive tradeoff" that minimizes MSE. In other words, at high observed SNR values (40 dB for Case A), the RISR algorithm produces a filter bank $\mathbf{W}$ that exchanges SNR mismatch loss in order to improve MSE.

The decrease and subsequent increase in MSE over time for 10 iteration RISR (the red trace) in Fig. 6 is due to frequency-angle coupling bias for lower chirp frequencies between 0 to 300 ns and higher chirp frequencies between 700 to 1000 ns. The coupling is also observed in Fig. 3, where a coarser DOA estimate is present between 0 to 300 ns. Additionally, remnant spatial sidelobes appear in these temporal spans around the 150° azimuth look-angle with -40 dB suppression.

As evidenced in Fig. 4, reducing the iterations per snapshot of RISR increases the mainlobe width because the algorithm is not yet fully converged. The aforementioned frequency-angle coupling is also now more apparent, where the estimated spatial angle is slightly biased towards negative values between 0 to 300 ns and towards positive values between 700 to 1000 ns. A clear tradeoff exists between RISR estimation performance and the computational cost of additional iterations.

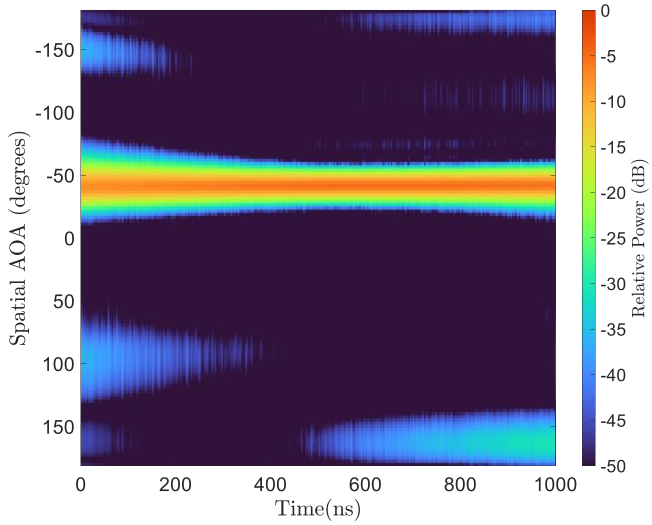

Fig. 4. Case A; Magnitude of RISR estimate over time, with 3 iterations per snapshot. The signal direction-of-arrival is coarser when RISR is performed with fewer iterations.

Now consider 3 iterations of RE-RISR, which yields the response in Fig. 5. Recycling of the filter bank from the previous snapshot in time achieves an improved estimate relative to 10-iteration RISR (Fig. 3) with the same computational cost as 3-iteration RISR (Fig. 4). The information fed forward permits continued refinement of the true spatial angle due to the near complete stationarity of the scene. In contrast to the slight mainlobe coarseness observed for 10-iteration RISR from 0 to 300 ns in Fig. 3, RE-RISR further super-resolves the estimate. However, the remnant spatial sidelobes produced by RE-RISR at the 150° azimuth look-angle appear slightly coarser. These qualitative similarities indicate that 3-iteration RE-RISR yields a 70% computational cost reduction relative to 10-iteration RISR without estimation degradation.

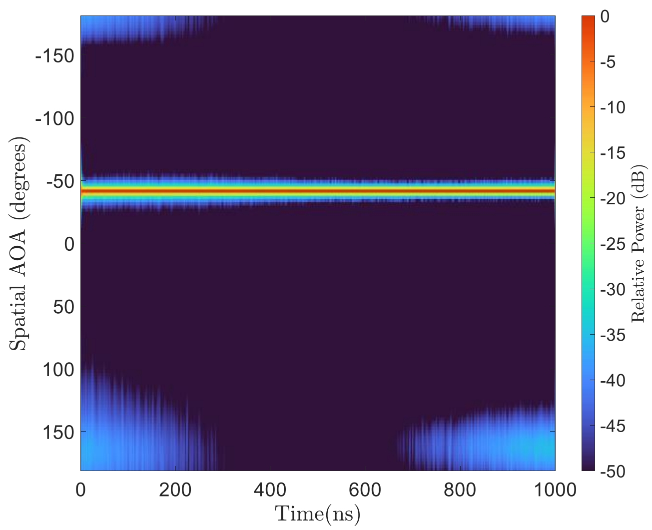

Fig. 5. Case A; Magnitude of the RE-RISR estimate over time, with 3 iterations per snapshot. RE-RISR with 3 iterations is able to qualitatively achieve the same estimation performance as RISR with 10 iterations (Fig. 3).

From Fig. 6, 3-iteration RE-RISR is able to converge to a solution with less or equal error than 10-iteration RISR for the Case A scenario. It should be noted that the error observed for RISR and RE-RISR having values < -40 dB are largely due to frequency-angle coupling caused by the fast-time signal phase characteristics. Importantly, the zoomed inset of Fig. 6 indicates that RE-RISR has a "learning rate", requiring nearly 6 temporal snapshots to fully converge to a minimized MSE of < -40 dB due to the feed-forward mechanism using fewer iterations. Of course, more iterations of RE-RISR can be performed per snapshot (with diminishing improvements in estimation). Doing so would be more computationally costly, but reduce the number of snapshots to achieve minimized MSE convergence.

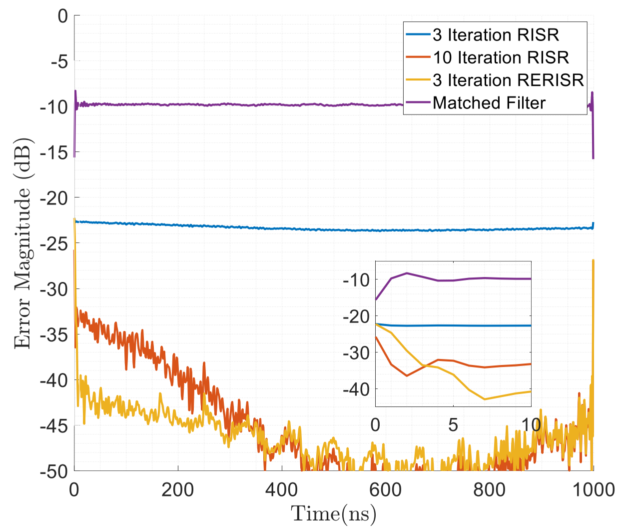

Fig. 6. Mean-squared error over time for Case A for each algorithm.

Case B is intended to stress the RE-RISR frequency-angle stationarity assumption with two signals. The pulsed chirp waveform illustrates the RE-RISR learning rate on slow-time scales. The OFDM waveform demonstrates rapidly varying fast time-frequency characteristics to additionally stress the learning rate. Lastly, the dynamic range between the chirp and OFDM waveforms is set to a 30 dB difference to examine small signal estimation.

In Fig. 7 the standard beamformer is applied to case B, with the pulsed chirp arriving from -40 degrees and the OFDM signal arriving from 30 degrees. When the chirp is active, the ensuing spatial sidelobes mask the OFDM signal, preventing it from being detected. Application of RISR with 10 iterations (Fig. 8) or 3 iterations (Fig. 9) allows the OFDM signal to be resolved when the pulsed chirp is active. For 3-iteration RISR, significant smearing of the OFDM signal occurs at the rise/fall times of the chirp waveform. By applying 10-iteration RISR, this effect is reduced but not eliminated. The sidelobes existing at $\phi = 150°$ present an estimation challenge for identifying small signals amidst a large signal error floor [17], [18].

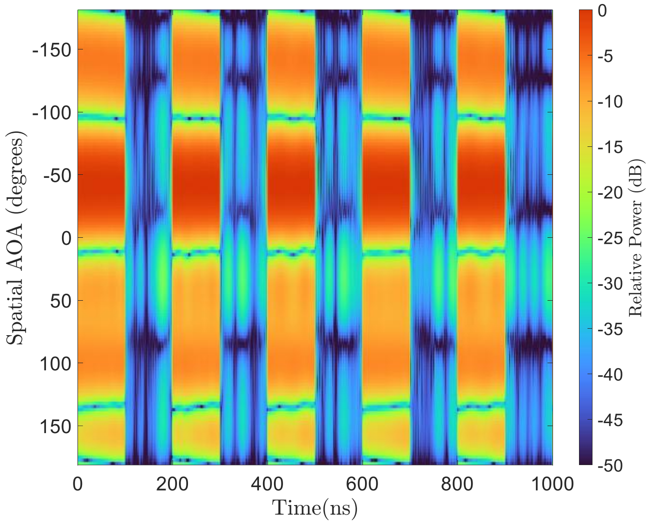

Fig. 7. Case B; Magnitude of the standard matched beamforming estimate over time. The low SNR OFDM signal at $\phi = 30°$ is masked by spatial sidelobes of the high SNR chirp signal at $\phi = -40°$. Poor resolution and sidelobes mask direction-of-arrival estimation.

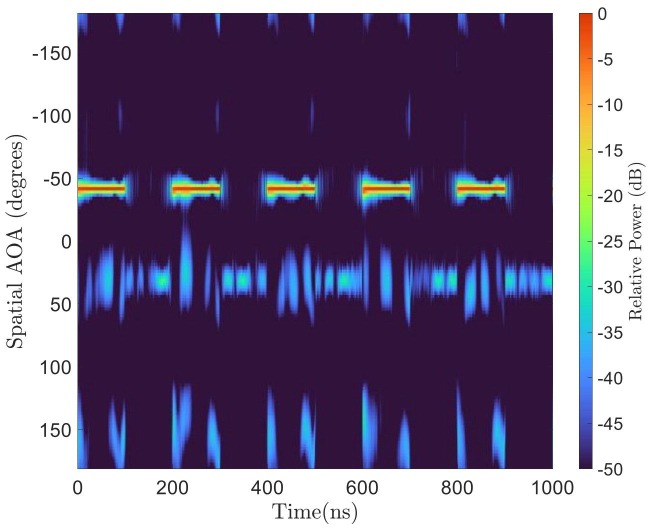

Fig. 8. Case B; Magnitude of RISR estimate over time, with 10 iterations per snapshot. Both signals are detected, although sidelobes at $\phi = 150°$ are not fully suppressed.

Fig. 10 shows the estimate for 3-iteration RE-RISR for the case B scenario, which is visually similar to that of 10-iteration RISR in Fig. 8. The primary difference is the focusing period when the pulsed chirp switches states. For the first snapshot after the chirp switches to on or off, the feed-forward initialization is incorrect, as it contains information about the chirp in its opposite state. Consequently, the RE-RISR estimate is less focused for transition snapshots. For the subsequent few snapshots, the algorithm re-focuses to a result that matches or surpasses 10-iteration RISR performance. The OFDM signal estimated by RE-RISR appears less focused than that of 10-iteration RISR due to its fast frequency-angle coupled nature. However, RE-RISR does better at estimating the OFDM signal when the chirp waveform is active than that of 3-iteration RISR. This result demonstrates that RE-RISR is able to contend with moderate nonstationarities on both slow-time and fast-time scales.

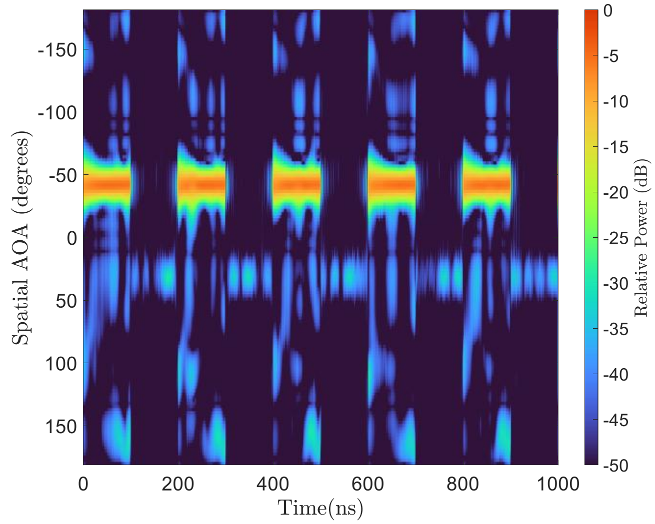

Fig. 9. Case B; Magnitude of RISR estimate over time, with 3 iterations per time snapshot. Both signals are detected, although estimates are significantly coarser.

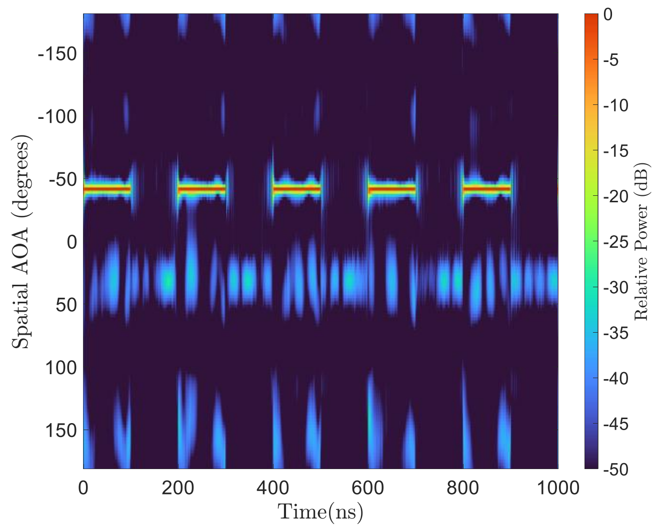

Fig. 10. Case B; Magnitude of the RE-RISR estimate over time, with 3 iterations per time snapshot. Additional error is present during nonstationary pulsed rise/fall times. The DOA for both signals are properly determined.

If all observed signals are nonstationary in fast time, but stationary in angle, RE-RISR will defocus while still providing accurate angle estimation. RE-RISR may be challenged if the scene is highly nonstationary in angle (on the order of a few snapshots) where the algorithm cannot converge. Resilience to nonstationary incoming angle can be mitigated by increasing the number of iterations per snapshot, at the cost of increased computational cost.

It can be seen in Fig. 11 that the error magnitude for case B is much higher when the pulsed chirp is on, due to the significant increase in total power when the chirp is active. The error in Fig. 11 is predominately caused by the higher power of the chirp signal. It then follows a similar pattern to the error in Fig. 6, on a shorter time scale for each pulse.

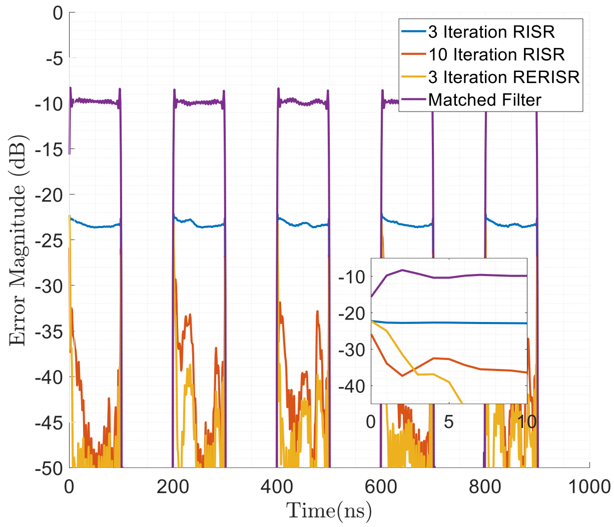

Fig. 11. Mean-squared error versus time for Case B for each algorithm.

VI. Conclusions

The RE-RISR form is a modification of the RISR algorithm that has been demonstrated in simulation to improve angle estimation performance and lower computational cost in scenarios where modest stationarity can be assumed. The feed-forward recycling of prior filter bank estimates, when at least partially accurate, greatly speeds up convergence, thereby allowing RE-RISR to reach an equivalent or superior result in fewer iterations when compared to RISR. This feed-forward technique can be applied to RMMSE estimators in other domains as well, given that the same quasi-stationary assumptions can be made. For instance, the spectrogram estimator demonstrated in [12] would potentially experience a significant reduction in computational cost through recycled estimation.

References

[1] M. Wax and J. Sheinvald, "Direction finding of coherent signals via spatial smoothing for uniform circular arrays," in *IEEE Trans. on Antennas and Propagation*, vol. 42, no. 5, pp. 613-620, May 1994.

[2] C. P. Mathews and M. D. Zoltowski, "Eigenstructure techniques for 2-D angle estimation with uniform circular arrays," in *IEEE Trans. on Signal Processing*, vol. 42, no. 9, pp. 2395-2407, Sept. 1994.

[3] P. Ioannides and C. A. Balanis, "Uniform circular and rectangular arrays for adaptive beamforming applications," in *IEEE Antennas and Wireless Propagation Letters*, vol. 4, pp. 351-354, 2005.

[4] R. Goossens and H. Rogier, "A Hybrid UCA-RARE/Root-MUSIC Approach for 2-D Direction of Arrival Estimation in Uniform Circular Arrays in the Presence of Mutual Coupling," in *IEEE Trans. on Antennas and Propagation*, vol. 55, no. 3, pp. 841-849, March 2007.

[5] B. R. Jackson, S. Rajan, B. J. Liao and S. Wang, "Direction of Arrival Estimation Using Directive Antennas in Uniform Circular Arrays," in *IEEE Trans. on Antennas and Propagation*, vol. 63, no. 2, pp. 736-747, Feb. 2015.

[6] C. Degen, "On single snapshot direction-of-arrival estimation," *2017 IEEE Int'l Conference on Wireless for Space and Extreme Environments*, Montreal, QC, Canada, 2017, pp. 92-97.

[7] S.D. Blunt, T. Chan, and K. Gerlach, "A new framework for direction-of-arrival estimation," IEEE Sensor Array & Multichannel Signal Processing Workshop, 21-23 July 2008.

[8] S.D. Blunt, T. Chan, and K. Gerlach, "Robust DOA estimation: the reiterative super resolution (RISR) algorithm," IEEE Trans. Aerospace & Electronic Systems, vol. 47, no. 1, pp. 332 – 346, Jan. 2011.

[9] E. Hornberger, S. D. Blunt and T. Higgins, "Partially constrained adaptive beamforming for super-resolution at Low SNR," *2015 IEEE 6th Int'l Workshop on Comp. Advances in Multi-Sensor Adaptive Processing*, Cancun, Mexico, 2015, pp. 129-132.

[10] T. Higgins, S. D. Blunt and K. Gerlach, "Gain-constrained adaptive pulse compression via an MVDR framework," *2009 IEEE Radar Conference*, Pasadena, CA, USA, 2009, pp. 1-6

[11] C. C. Jones *et al.*, "Development & Experimental Assessment of Robust Direction Finding and Self-Calibration," *2022 IEEE Radar Conf.*, New York City, NY, USA, 2022.

[12] C. C. Jones, Z. E. Gannon, S. D. Blunt, C. T. Allen and A. F. Martone, "An Adaptive Spectrogram Estimator to Enhance Signal Characterization," *2022 IEEE Radar Conf.*, New York City, NY, USA, 2022.

[13] R. Mital, M. Parent, A. Stumme, "Circular Array Beamforming Using Phase Modes", NRL Memorandum Report, NRL/MR/5317--19-9908, Dec. 2019.

[14] S. Kay, *Fundamentals of Statistical Siganl Processing: Estimation Theory*. Prentice Hall, 1993.

[15] K. . -L. Du, "Pattern analysis of uniform circular array," in *IEEE Trans. on Antennas and Propagation*, vol. 52, no. 4, pp. 1125-1129, April 2004.

[16] C. Suarez, M. Ferrando-Bataller and A. Valero-Nogueira, "Pattern synthesis of uniform circular arrays with directive elements," *IEEE Antennas and Propagation Society Symposium, 2004.*, Monterey, CA, USA, 2004, pp. 2812-2815.

[17] C. D. Richmond, "Performance of a class of adaptive detection algorithms in nonhomogeneous environments," in *IEEE Trans. on Signal Processing*, vol. 48, no. 5, pp. 1248-1262, May 2000.

[18] F. Bandiera, A. Farina, D. Orlando and G. Ricci, "Detection Algorithms to Discriminate Between Radar Targets and ECM Signals," in *IEEE Trans. on Signal Processing*, vol. 58, no. 12, pp. 5984-5993, Dec. 2010.