Computationally Efficient Joint-Domain Clutter Cancellation for Waveform-Agile Radar

Christian Jones[1,3], Brandon Ravenscroft[1], James Vogel[2], Suzanne M. Shontz[1], Thomas Higgins[3], Kevin Wagner[3], Shannon Blunt[1]
[1]Radar Systems Lab (RSL), EECS Dept., University of Kansas, Lawrence, KS, USA
[2]Systems & Technology Research, Woburn, MA, USA
[3]Naval Research Laboratory - Radar Division, Washington, DC, USA

Abstract—**Waveform agility provides greater design freedom through the generation of a coherent processing interval (CPI) of nonrepeating waveforms. However, doing so introduces coupling of the range and slow-time Doppler dimensions that hinders clutter cancellation if not addressed. Non-identical multiple pulse compression (NIMPC) is a joint-domain approach that solves this problem, though the high dimensionality incurs a prohibitive computational cost. Inspired by recent work that exploits the Toeplitz structure of NIMPC, here we take that notion even further, demonstrating the processing of measured data at a cost significantly lower than direct NIMPC application.**

Keywords—moving target indication, waveform agility, joint-domain processing, computational efficiency

I. Introduction

Standard moving target indication (MTI) radar involves illumination by a repeated waveform, such as the commonly used linear frequency modulated (LFM) chirp. The receive-captured reflections are then pulse compressed followed by (slow-time) Doppler processing to separate movers from stationary clutter. For this repeated-waveform framework, the range and slow-time Doppler domains are decoupled; thereby permitting separate, sequential processing according to the dimensionality of each distinct domain [1].

The introduction of variability among the waveforms within the CPI via *waveform agility* (or *pulse agility*) greatly increases design freedom [2], but also causes the range and slow-time Doppler domains to become coupled. The result is variability of the sidelobe structure on a pulse-to-pulse basis, inducing the phenomenon known as range sidelobe modulation (RSM) [3]. While RSM provides a benefit in terms of incoherent sidelobe averaging, it also introduces a nonstationary effect that can hinder clutter cancellation when standard sequential (decoupled) processing is performed [4].

Specifically, Doppler processing of a CPI containing RSM yields a smearing of clutter (at/near zero Doppler) across the entire Doppler domain. Moreover, significant changes in the waveform spectrum over the CPI (such as in dynamic spectral notching [5]) can also modulate the pulse compression mainlobe, further exacerbating this smearing effect.

There are two general philosophies when it comes to compensating for RSM for clutter cancellation. One seeks to "homogenize" on a pulse-to-pulse basis so that processing can still be performed in a sequential manner. Early efforts designed mismatched filters (MMFs) that produce sidelobe responses with sufficient similarities to minimize RSM [2,6]. However, because doing so tends to expend degrees of freedom that would be better served suppressing sidelobes, it was realized (see [7]) that MMFs focused solely on sidelobe suppression (and not homogenization) are generally a better alternative for arbitrary nonrepeating waveforms (see [4]).

The other philosophy involves jointly processing the range and slow-time Doppler domains, which is appropriate given their inherent coupling in this context. These approaches model delay-shifted versions of each pulsed waveform to account for the range domain, combined with the slow-time phase progression across the CPI for each Doppler bin. While some adaptive joint estimation methods have been conceived [8-10], they do not address clutter cancellation.

In contrast, NIMPC [11] was formulated to have a structured (clutter + noise) covariance matrix within this joint-domain context. While not adaptive in the classical sense (i.e. not constructed from snapshots of sampled data), NIMPC can be applied in a manner akin to maximum SINR filtering [12] to cancel clutter and RSM, as shown with recent experimental demonstration [13].

Of course, a limiting factor for NIMPC is the computational cost of inverting an $NM \times NM$ joint-domain matrix, where M is the number of pulses in the CPI and N is the number of samples in each discretized waveform (assuming consistent time-bandwidth product (BT)). To achieve sufficient fidelity for receive processing, over-sampling relative to waveform 3-dB bandwidth is required depending on the spectral roll-off, e.g. LFM is quite sharp while nonrepeating random FM (RFM) waveforms [4] tend to be more gradual. Noting that NIMPC incurs a computational cost of $O((MN)^3)$ for the matrix inverse alone, in [14] this covariance matrix was reformulated as block-Toeplitz, thereby reducing the matrix inversion cost by a factor of N.

Here we take the idea in [14] a step further, facilitating efficient application of linear conjugate gradient (CG) methods for NIMPC by leveraging the block-Toeplitz structure to realize a frequency-domain implementation of circulant matrix multiplication. The performance of these efficient solutions is assessed using open-air measurements.

II. NIMPC Model

Denote an $N \times 1$ collection of contiguous fast-time received samples induced by the mth pulse as

$$\mathbf{y}_m(l) = \sum_{\omega} \mathbf{S}_{m,\omega}\, \mathbf{x}_{\omega}(l) + \mathbf{n}_m(l)\,, \tag{1}$$

This work was supported by the Office of Naval Research under Contract #N00014-20-C-1006.

where l indicates the range cell index and $\mathbf{n}_m(l)$ contains the N associated additive noise samples. The $N \times (2N-1)$ matrix $\mathbf{S}_{m,\omega} = \mathbf{S}_m \, e^{j(m-1)\omega}$ accounts for the phase shift associated with Doppler frequency ω. The Toeplitz structure of

$$\mathbf{S}_m = \begin{bmatrix} s_{m,N} & \cdots & \cdots & s_{m,1} & 0 & \cdots & 0 \\ 0 & \ddots & & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & & \ddots & 0 \\ 0 & \cdots & 0 & s_{m,N} & \cdots & \cdots & s_{m,1} \end{bmatrix} \tag{2}$$

models the convolution of N discretized samples of waveform $s_m(t)$, denoted as vector $\mathbf{s}_m = [s_{m,1} \; s_{m,2} \; \cdots \; s_{m,N}]$, with $2N-1$ discretized range samples of complex scattering in $\mathbf{x}_\omega(l)$.

The response in (1) for each pulse is then collected into the $N \times M$ matrix

$$\mathbf{Y}(l) = [\mathbf{y}_1(l) \;\; \mathbf{y}_2(l) \; \cdots \; \mathbf{y}_M(l)] \,, \tag{3}$$

which is organized in terms of fast-time along each column and slow-time along each row. This range-Doppler coupled form can likewise be expressed as an $NM \times 1$ vector by concatenating (or *vectorizing*) the columns of (3) as

$$\bar{\mathbf{y}}(l) = [\mathbf{y}_1^T(l) \;\; \mathbf{y}_2^T(l) \; \cdots \; \mathbf{y}_M^T(l)]^T \,, \tag{4}$$

where $(\bullet)^T$ is the transpose operator. The coupled equivalent to standard sequential range-Doppler processing at the lth range bin and Doppler frequency ω can then be realized via

$$\hat{x}(l,\omega) = \frac{\bar{\mathbf{b}}_\omega^H \, \bar{\mathbf{y}}(l)}{\bar{\mathbf{b}}_\omega^H \bar{\mathbf{b}}_\omega} \,, \tag{5}$$

where $(\bullet)^H$ is the Hermitian operator,

$$\bar{\mathbf{b}}_\omega = [\mathbf{s}_1^T \;\; e^{j\omega}\mathbf{s}_2^T \;\; \cdots \;\; e^{j(M-1)\omega}\mathbf{s}_M^T]^T \tag{6}$$

is a Doppler-shifted waveform sequence, and normalization in (5) is included for consistent scaling.

Incorporating (1) into (4) yields

$$\bar{\mathbf{y}}(l) = \sum_\omega \bar{\mathbf{S}}_\omega \, \mathbf{x}_\omega(l) + \bar{\mathbf{n}}(l) \,, \tag{7}$$

where the $NM \times (2N-1)$ concatenated matrix is

$$\bar{\mathbf{S}}_\omega = [\mathbf{S}_{1,\omega}^T \;\; \mathbf{S}_{2,\omega}^T \; \cdots \; \mathbf{S}_{M,\omega}^T]^T \,, \tag{8}$$

which subsumes Doppler as noted after (1), and the NM noise samples are collected into $\bar{\mathbf{n}}(l) = [\mathbf{n}_1^T(l) \;\; \mathbf{n}_2^T(l) \; \cdots \; \mathbf{n}_M^T(l)]^T$.

The NIMPC clutter cancellation framework [11] defines a joint-domain filter for each Doppler bin and suppresses the clutter response from a collection of specified Doppler values (i.e., around zero). To do so, the $NM \times NM$ clutter-plus-noise structured correlation matrix

$$\bar{\mathbf{R}} = \sum_{\omega \in \Omega} (\bar{\mathbf{S}}_\omega \, \bar{\mathbf{S}}_\omega^H) + \bar{\mathbf{R}}_{\text{nse}} \tag{9}$$

is formed, where the noise component $\bar{\mathbf{R}}_{\text{nse}} = \sigma_{\text{nse}}^2 \, \mathbf{I}_{NM \times NM}$ is an identity matrix scaled by noise power σ_{nse}^2 in the case of white noise. The set Ω comprises the particular Doppler values associated with clutter (requires prior determination) and each term $(\bar{\mathbf{S}}_\omega \bar{\mathbf{S}}_\omega^H)$ in the summation is composed of $M \times M$ blocks, where each block is an $N \times N$ Toeplitz matrix.

In the efficient implementation that follows it is convenient to express the correlation matrix in a more compact form while also explicitly specifying the cardinality of Ω. Consequently, we can also pose (9) as

$$\bar{\mathbf{R}} = \bar{\mathbf{S}}_\Omega \, \bar{\mathbf{S}}_\Omega^H + \bar{\mathbf{R}}_{\text{nse}} \tag{10}$$

in which the Doppler interval of the clutter is discretized into K bins (having sufficient granularity). Thus,

$$\bar{\mathbf{S}}_\Omega = [\bar{\mathbf{S}}_{\omega_1} \;\; \bar{\mathbf{S}}_{\omega_2} \; \cdots \; \bar{\mathbf{S}}_{\omega_K}] \tag{11}$$

is an $NM \times (2N-1)K$ matrix, with ω_k for $k = 1, \cdots, K$ the indexed Doppler values.

The NIMPC filter [11] for Doppler ω, expressed using the nomenclature above, is therefore

$$\bar{\mathbf{w}}_\omega = \bar{\mathbf{R}}^{-1} \bar{\mathbf{b}}_\omega \tag{12}$$

and is applied by replacing $\bar{\mathbf{b}}_\omega$ with $\bar{\mathbf{w}}_\omega$ in (5). However, this approach requires inversion of an $NM \times NM$ matrix, which is computationally prohibitive. Alternatively, consider the form

$$\bar{\mathbf{R}} \, \bar{\mathbf{w}}_\omega = \bar{\mathbf{b}}_\omega \tag{13}$$

to solve for $\bar{\mathbf{w}}_\omega$ more efficiently by leveraging redundant signal structure in $\bar{\mathbf{b}}_\omega$ and $\bar{\mathbf{R}}$. It was shown in [14] that $\bar{\mathbf{b}}_\omega$ from (6) can be equivalently expressed as

$$\bar{\mathbf{b}}_\omega = \begin{bmatrix} \mathbf{s}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{s}_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{s}_M \end{bmatrix} \begin{bmatrix} 1 \\ e^{j\omega} \\ \vdots \\ e^{j(M-1)\omega} \end{bmatrix} = \mathbf{C}\mathbf{v}_\omega \,, \tag{14}$$

where $\mathbf{C} = [\mathbf{c}_1 \;\; \mathbf{c}_2 \; \cdots \; \mathbf{c}_M]$ is $NM \times M$, each $\mathbf{0}$ in the matrix is an $N \times 1$ vector of zeros, and $\mathbf{v}_\omega$ is simply an $M \times 1$ Doppler steering vector. Ignoring the $\mathbf{v}_\omega$ term for a moment, (13) can otherwise be rewritten as

$$\bar{\mathbf{R}} \, \bar{\mathbf{D}} = \mathbf{C} \tag{15}$$

for $NM \times M$ matrix $\bar{\mathbf{D}} = [\bar{\mathbf{d}}_1 \;\; \bar{\mathbf{d}}_2 \;\; \cdots \;\; \bar{\mathbf{d}}_M]$. It can therefore be shown that

$$\bar{\mathbf{w}}_\omega = \bar{\mathbf{D}} \mathbf{v}_\omega \,, \tag{16}$$

such that determination of $\bar{\mathbf{D}}$ provides the means to readily compute the joint-domain filter for arbitrary Doppler. In the next section we separate (15) into the constituent components

$$\bar{\mathbf{R}} \, \bar{\mathbf{d}}_m = \mathbf{c}_m \tag{17}$$

for $m = 1, \cdots, M$ to solve iteratively for each column of $\bar{\mathbf{D}}$.

III. Efficient Implementation #1: PCG

Rather than solving (17) through matrix inversion, via forward/backward substitution [15], or using an efficient block Toeplitz solver [16], an approximate solution can be iteratively refined through convex optimization techniques. It is known [17] that the solution to (17) is the same as the minimizer of the quadratic (and thus convex) cost function

$$f(\bar{\mathbf{d}}_m) = \bar{\mathbf{d}}_m^H \bar{\mathbf{R}} \, \bar{\mathbf{d}}_m - \bar{\mathbf{d}}_m^H \mathbf{c}_m - \mathbf{c}_m^H \bar{\mathbf{d}}_m \,, \tag{18}$$

which has the gradient

$$\nabla_{\bar{\mathbf{d}}_m^*} f(\bar{\mathbf{d}}_m) = \bar{\mathbf{R}} \, \bar{\mathbf{d}}_m - \mathbf{c}_m \tag{19}$$

that is identical to (17) and the Hessian

$$\nabla^2 f(\bar{\mathbf{d}}_m) = \bar{\mathbf{R}} \,. \tag{20}$$

Steepest descent would be simple to implement, yet it only has linear convergence. Newton/quasi-Newton methods are faster, but have complexity/memory requirements that are comparable to direct inversion. However, the preconditioned linear conjugate gradient (PCG) algorithm provides super-linear convergence to the optimum [17], with each iteration moving the optimal step in a direction conjugate to all previous directions. Since convergence would otherwise be dictated by the condition number of $\bar{\mathbf{R}}$, a positive-definite preconditioner matrix $\bar{\mathbf{M}}$ is introduced, providing a trade-space between solution quality and computational cost, with good performance achieved in just a few iterations. The block-Toeplitz structure also readily supports parallel processing.

For starting point $\bar{\mathbf{d}}_{m,i=0}$ and known $\bar{\mathbf{M}}^{-1}$, PCG [17] in this context is initialized as

$$\begin{aligned} &\mathbf{r}_0 = \mathbf{c}_{m,0} - \bar{\mathbf{R}}\,\bar{\mathbf{d}}_{m,0} \\ &\mathbf{u}_0 = \bar{\mathbf{M}}^{-1}\mathbf{r}_0 \\ &\mathbf{p}_0 = \mathbf{u}_0 \\ &i = 0 \end{aligned} \quad . \tag{21}$$

Given stopping threshold ϕ, then I iterations are performed via

$$\begin{aligned} &\text{while } i < I \text{ and } \|\mathbf{r}_i\| < \phi \\ &\quad \alpha_i = \frac{\mathbf{r}_i^H \mathbf{u}_i}{\mathbf{p}_i^H \bar{\mathbf{R}}\,\mathbf{p}_i} \\ &\quad \bar{\mathbf{d}}_{m,i+1} = \bar{\mathbf{d}}_{m,i} + \alpha_i \mathbf{p}_i \\ &\quad \mathbf{r}_{i+1} = \mathbf{r}_i - \alpha_i \bar{\mathbf{R}}\mathbf{p}_i \\ &\quad \mathbf{u}_{i+1} = \bar{\mathbf{M}}^{-1}\mathbf{r}_{i+1} \\ &\quad \beta_i = \frac{\mathbf{r}_{i+1}^H \mathbf{u}_{i+1}}{\mathbf{r}_i^H \mathbf{u}_i} \\ &\quad \mathbf{p}_{i+1} = \mathbf{u}_{i+1} + \beta_i \mathbf{p}_i \\ &\quad i = i+1 \\ &\text{end} \end{aligned} \quad , \tag{22}$$

with complexity dominated by matrix-vector multiplications.

However, using (10) and some arbitrary vector $\mathbf{f}$ we can decompose the following matrix-vector multiplication as

$$\begin{aligned} \bar{\mathbf{R}}\,\mathbf{f} &= (\bar{\mathbf{S}}_\Omega \bar{\mathbf{S}}_\Omega^H + \bar{\mathbf{R}}_{\text{nse}})\mathbf{f} \\ &= \bar{\mathbf{S}}_\Omega \bar{\mathbf{S}}_\Omega^H \mathbf{f} + \bar{\mathbf{R}}_{\text{nse}}\,\mathbf{f} \\ &= \bar{\mathbf{S}}_\Omega \mathbf{g} + \bar{\mathbf{R}}_{\text{nse}}\,\mathbf{f} \\ &= \mathbf{h} + \sigma_{\text{nse}}^2\,\mathbf{f} \ , \end{aligned} \tag{23}$$

where the last line imposes the assumption of white noise. Using Hadamard product $\odot$, the $(2N-1)K \times 1$ intermediate term can be efficiently computed via

$$\begin{aligned} \mathbf{g}_k &= \left[\sum_{m=1}^{M} \mathbf{A}^H \Big(e^{-j(m-1)\omega_k} (\mathbf{q})^* \odot (\mathbf{A}\tilde{\mathbf{s}}_m)^* \odot (\mathbf{A}\tilde{\mathbf{f}}_m) \Big) \right], \\ \mathbf{g} &= \begin{bmatrix} \mathbf{g}_1^T & \mathbf{g}_2^T & \cdots & \mathbf{g}_K^T \end{bmatrix}^T = \bar{\mathbf{S}}_\Omega^H \mathbf{f} \end{aligned} \tag{24}$$

where $\tilde{\mathbf{s}}_m$ and $\tilde{\mathbf{f}}_m$ are obtained from the mth N-length segment of each respective waveform that is subsequently padded with $N-1$ zeros (yielding $(2N-1)\times 1$ vectors), and the $\mathbf{A}$ and $\mathbf{A}^H$ matrices are the $(2N-1)\times(2N-1)$ discrete Fourier transform (DFT) and inverse DFT, respectively. The vector $\mathbf{q}$, with elements $q(n) = e^{-j2\pi nN/(2N-1)}$ for $n = 0, 1, \cdots, 2N-2$, is a complex sinusoid that equivalently shifts $\tilde{\mathbf{s}}_m$ by N samples via multiplication in the frequency domain.

The $\mathbf{h}$ term in (23) can then similarly be computed as

$$\begin{aligned} \mathbf{h}_m &= \tilde{\mathbf{I}}_{N\times(2N-1)} \left[\sum_{k=1}^{K} \mathbf{A}^H \Big(e^{j(m-1)\omega_k} (\mathbf{q}) \odot (\mathbf{A}\tilde{\mathbf{s}}_m) \odot (\mathbf{A}\mathbf{g}_k) \Big) \right] \\ \mathbf{h} &= \begin{bmatrix} \mathbf{h}_1^T & \mathbf{h}_2^T & \cdots & \mathbf{h}_M^T \end{bmatrix}^T = \bar{\mathbf{S}}_\Omega \mathbf{g} \ , \end{aligned} \tag{25}$$

where $\tilde{\mathbf{I}}_{N\times(2N-1)} = [\mathbf{I}_{N\times N}\ \mathbf{0}_{N\times(N-1)}]$ is a zero-appended identity matrix that truncates the last $(N-1)$ terms of a $(2N-1)\times 1$ vector. When (24) and (25) are performed sequentially using fast Fourier transforms (FFTs), and with each waveform FFT computed offline, the complexity of multiplication by $\bar{\mathbf{R}}$ reduces from $O(N^2M^2)$ to $O(2MK(2N-1)\log_2(2N-1))$.

The preconditioner $\bar{\mathbf{M}}$ can greatly accelerate convergence and can be used with (17) to initialize the starting point as

$$\bar{\mathbf{d}}_{m,0} = \bar{\mathbf{M}}^{-1}\mathbf{c}_m \, . \tag{26}$$

However, preconditioning incurs an additional computational cost, and thus $\bar{\mathbf{M}}$ needs to be selected such that the complexity of (21) and (22) do not significantly increase. In [18] a block-circulant approximation was shown to improve convergence and is efficiently solved with a per-iteration cost of at most $O(2M^2N\log_2(N))$ [19]. Alternatively, if the K Doppler frequencies are located about some center (like (11)), diagonal blocks of a block-circulant preconditioner can instead reduce the per-iteration complexity by a factor of M while still achieving faster convergence. These two implementations are denoted as PCG-1 (full block-circulant) and PCG-2 (diagonal block-circulant).

IV. Efficient Implementation #2: Proj-NIMPC

An alternative fast implementation can be realized by leveraging the projection-based approach of the extensive cancellation algorithm (ECA) [20], which was developed for passive bistatic radar. Combining ECA with the NIMPC model yields a projection form of NIMPC denoted as Proj-NIMPC. Rather than the piecewise model of (1) involving $N\times 1$ measurement snapshots, here a length-L range interval is addressed collectively. For the mth pulse, this $L\times 1$ received vector is modeled similar to (1) as

$$\dot{\mathbf{y}}_m = \sum_{\omega} \dot{\mathbf{S}}_{m,\omega}\,\dot{\mathbf{x}}_\omega + \dot{\mathbf{n}}_m \tag{27}$$

where this Toeplitz matrix is $L\times(L-N+1)$ and has the form

$$\dot{\mathbf{S}}_m = \begin{bmatrix} s_{m,1} & 0 & \cdots & \cdots & \cdots & 0 \\ \vdots & \ddots & & & & \vdots \\ s_{m,N} & & s_{m,1} & & & \vdots \\ 0 & \ddots & \vdots & \ddots & & \vdots \\ \vdots & & s_{m,N} & & \ddots & 0 \\ \vdots & & & \ddots & & s_{m,1} \\ \vdots & & & & \ddots & \vdots \\ 0 & \cdots & \cdots & \cdots & 0 & s_{m,N} \end{bmatrix} . \tag{28}$$

The joint-domain extension like (7) is now the $LM \times 1$ vector

$$\ddot{\mathbf{y}} = \sum_{\omega} \ddot{\mathbf{S}}_{\omega} \dot{\mathbf{x}}_{\omega} + \ddot{\mathbf{n}} , \tag{29}$$

where the $LM \times (L-N+1)$ concatenated matrix in (29) is

$$\ddot{\mathbf{S}}_{\omega} = [\dot{\mathbf{S}}_{1,\omega}^{T} \;\; \dot{\mathbf{S}}_{2,\omega}^{T} \;\; \cdots \;\; \dot{\mathbf{S}}_{M,\omega}^{T}]^{T} \tag{30}$$

in the same manner as (8), and likewise for the noise vector.

Similarly constructing the joint-domain clutter component

$$\ddot{\mathbf{S}}_{\Omega} = [\ddot{\mathbf{S}}_{\omega_1} \;\; \ddot{\mathbf{S}}_{\omega_2} \;\; \cdots \;\; \ddot{\mathbf{S}}_{\omega_K}] \tag{31}$$

like (11), this $LM \times (L-N+1)K$ matrix is used to form

$$\ddot{\mathbf{P}} = \ddot{\mathbf{I}} - \ddot{\mathbf{S}}_{\Omega} (\ddot{\mathbf{S}}_{\Omega}^{H} \ddot{\mathbf{S}}_{\Omega})^{-1} \ddot{\mathbf{S}}_{\Omega}^{H} , \tag{32}$$

which projects (29) onto the nullspace of the clutter via

$$\begin{aligned} \ddot{\mathbf{z}} = \ddot{\mathbf{P}}\, \ddot{\mathbf{y}} &= (\ddot{\mathbf{I}} - \ddot{\mathbf{S}}_{\Omega} (\ddot{\mathbf{S}}_{\Omega}^{H} \ddot{\mathbf{S}}_{\Omega})^{-1} \ddot{\mathbf{S}}_{\Omega}^{H})\, \ddot{\mathbf{y}} \\ &= \ddot{\mathbf{y}} - \ddot{\mathbf{S}}_{\Omega} (\ddot{\mathbf{S}}_{\Omega}^{H} \ddot{\mathbf{S}}_{\Omega})^{-1} \ddot{\mathbf{S}}_{\Omega}^{H}\, \ddot{\mathbf{y}} \\ &= \ddot{\mathbf{y}} - \ddot{\mathbf{S}}_{\Omega} (\ddot{\mathbf{S}}_{\Omega}^{H} \ddot{\mathbf{S}}_{\Omega})^{-1} \ddot{\mathbf{g}} \\ &= \ddot{\mathbf{y}} - \ddot{\mathbf{S}}_{\Omega}\, \ddot{\mathbf{a}} \\ &= \ddot{\mathbf{y}} - \ddot{\mathbf{h}} \end{aligned} \quad , \tag{33}$$

followed by subsequent standard range/Doppler processing of $\ddot{\mathbf{z}}$. Here $\ddot{\mathbf{g}}$ and $\ddot{\mathbf{h}}$ can be computed efficiently in the same manner as (24) and (25), respectively, albeit with the removal of $\mathbf{q}$ and the zero-padding increased to $(L-N)$. Solving for the intermediate term

$$\ddot{\mathbf{a}} = (\ddot{\mathbf{S}}_{\Omega}^{H} \ddot{\mathbf{S}}_{\Omega})^{-1} \ddot{\mathbf{g}} \tag{34}$$

can be achieved by rearranging (34) as

$$(\ddot{\mathbf{S}}_{\Omega}^{H} \ddot{\mathbf{S}}_{\Omega}) \ddot{\mathbf{a}} = \ddot{\mathbf{g}} , \tag{35}$$

which maps directly back into the framework of (13)-(17) to permit efficient computation via the PCG method outlined in (21) and (22). The computational distinction here is that PCG is now being applied to the $(L-N+1)K \times (L-N+1)K$ matrix in (35) and associated preconditioner, instead of the corresponding $NM \times NM$ matrices in the previous section.

V. Computational Complexity

Table I delineates the computational cost of the various NIMPC implementations. In addition to achieving substantially greater efficiency as discretized waveform dimensionality N and pulse number M grow, these fast methods also avoid the need to store the extremely large joint-domain matrices.

TABLE I. Complexity of NIMPC Implementations

NIMPC via [11]	$K(2N-1)\,(MN)^2 + (MN)^3 + D\,(MN)^2$
faster NIMPC via [14]	$K(2N-1)\,(MN)^2 + M(MN)^2 + (MN)^2$
CG	$2\,I\,M^2\,K\,(2N-1)\,\log_2(2N-1)$
PCG-1	$2IM^2K(2N-1)\log_2(2N-1) + 2IM^3N\log_2(N)$
PCG-2	$2IM^2K(2N-1)\log_2(2N-1) + 2IM^2N\log_2(N)$
Proj-NIMPC (direct (34))	$K^3(L-N+1)^2 + 2K^2M(L-N+1)\,L\log_2(L)$
Proj-NIMPC (w/ PCG-2)	$4(2I+1)\,KML\log_2(L)$

VI. Experimental Evaluation

Open-air measurements were collected to assess the different NIMPC joint-domain implementations on real MTI data. Here $M = 100$ nonrepeating PRO-FM waveforms [7] were generated using a Tektronix AWG70002A arbitrary waveform generator with $T = 4.5$ μs pulsewidth and 3-dB bandwidth of $B = 33.3$ MHz (so $BT = 150$), at a 3.55 GHz center frequency and 5 kHz pulse repetition frequency (PRF).

The various NIMPC implementations place a zero-centered Doppler null with $K = 10$ bins equally spaced over ±150 Hz, use $\sigma_{\text{nse}}^2 = 10^{-4}$, and discretize the waveforms to $N = 900$. The number of range window samples was $L = 5.9\times10^3$. When CG or PCG-2 were employed, $I = 10$ iterations were performed. The actual number of operations of each implementation for these particular parameters is listed in Table II. While the faster NIMPC approach of [14] does provide a factor of ~11 reduction relative to NIMPC [11] by exploiting Toeplitz structure, the PCG and projection/PCG implementations above achieve reductions of 4×10^4 and 3×10^5, respectively.

TABLE II. Complexities for Experimental Evaluation

NIMPC via [11]	9.5×10^{14}
faster NIMPC via [14]	8.7×10^{13}
CG NIMPC, $I = 10$	2.1×10^{10}
PCG-2 NIMPC, $I = 10$	2.3×10^{10}
Proj-NIMPC (direct (34))	1.9×10^{12}
Proj-NIMPC (PCG-2, $I = 10$)	3.1×10^{9}

Open-air measurements were collected from the roof of Nichols Hall on the University of Kansas campus, with collocated transmit/receive antennas aimed at a nearby traffic intersection to illuminate movers. Figure 1 illustrates the impact of transmitting a CPI of nonrepeating waveforms, where we clearly see the emergence of range sidelobe modulation (RSM) via the spreading of clutter across Doppler. Specifically, because the transmitter and receiver are operating concurrently due to the short range, the direct path leakage dominates the entire Doppler interval for the first 800 meters. The insert in Fig. 1 highlights the traffic intersection ~1 km away, where the RSM effect is still observed, albeit to a lesser degree since the responses are closer to the noise floor.

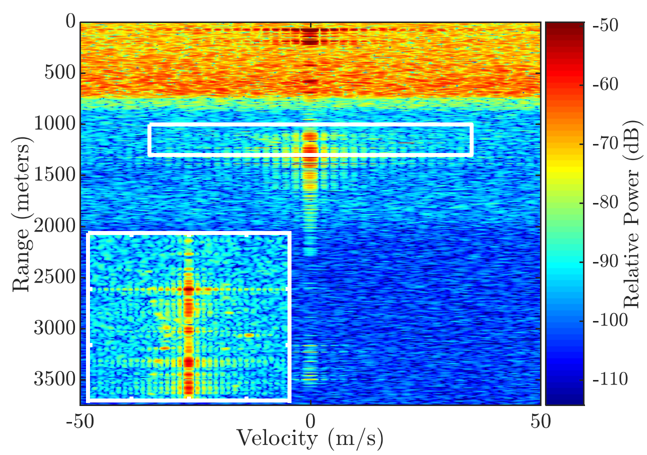

Fig. 1. Standard range-Doppler processing (no clutter cancellation) for open-air measurement using 100 pulsed PRO-FM waveforms

Of course, Fig. 1 employs no clutter cancellation. When it is applied (Fig. 2) via a simple Doppler-only projection, the direct path response and stationary clutter are suppressed. However, because the RSM effect spreads across Doppler, it persists after cancellation because the inherent coupling between range and slow-time Doppler are not being addressed. Further discussion can be found in [4] and references therein.

Figure 3 shows the improvement that can be achieved with joint-domain clutter cancellation via NIMPC [11]. Now the direct path spreading is almost completely mitigated (suppressed by about 35 dB) and the various movers (cars and trucks) in the intersection are clearly visible in the insert. Some Doppler sidelobe roll-off is still present, which is to be expected since NIMPC does not currently permit Doppler windowing. This computationally expensive, yet exact, result will serve as a performance benchmark for the efficient implementation results that follow.

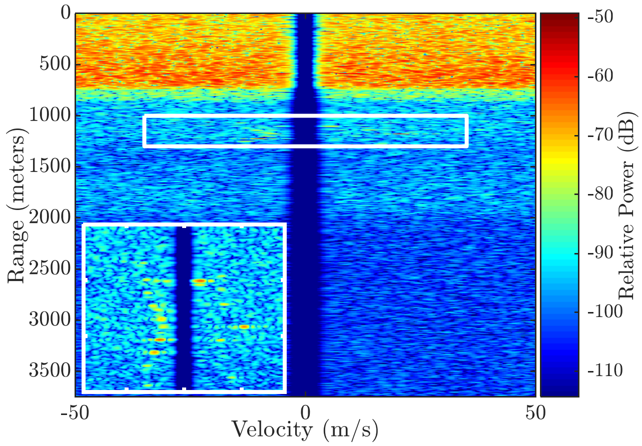

Fig. 2. Standard range-Doppler processing (with clutter cancellation) for open-air measurement using 100 pulsed PRO-FM waveforms

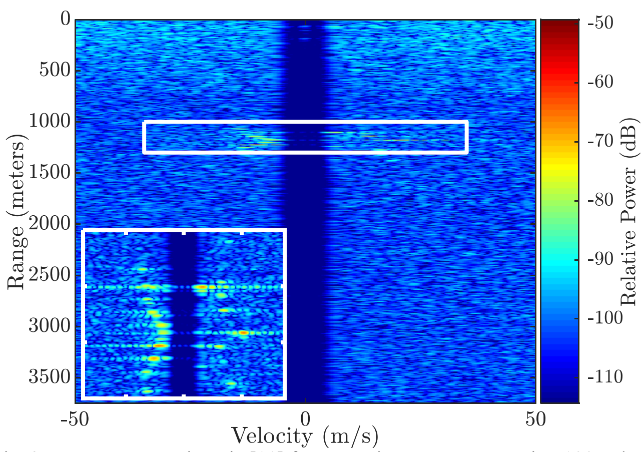

Fig. 3. NIMPC processing via [11] for open-air measurement using 100 pulsed PRO-FM waveforms

Figures 4 and 5 show the CG and PCG-2 implementations of NIMPC, each after 10 iterations. This value was selected to illustrate the trade-off between further refining the approximation with more iterations and controlling computational cost, which is nearly identical for these methods (per Table II). In the intersection region, both approaches yield results that are almost identical to the direct NIMPC result in Fig. 3. However, for the CG case in Fig. 4, which does not use preconditioning, residual RSM from the direct path is still noticeably visible because the approach takes longer to converge. In contrast, the PCG-2 case in Fig. 5 demonstrates there is a clear convergence advantage to preconditioning because the RSM, while not yet completely mitigated, is suppressed by an additional 8 dB by comparison.

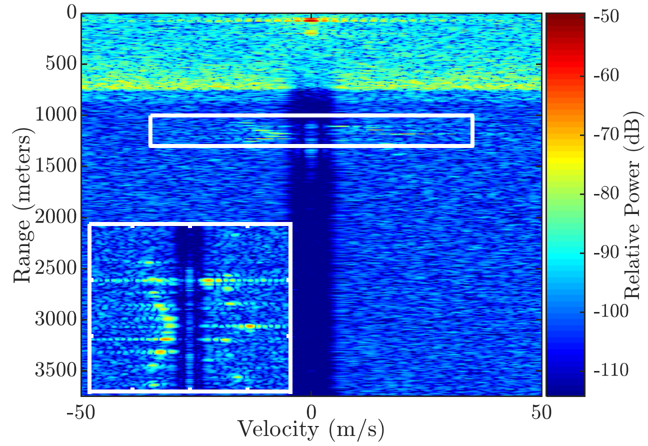

Fig. 4. NIMPC CG implementation ($I=10$) for open-air measurement using 100 pulsed PRO-FM waveforms

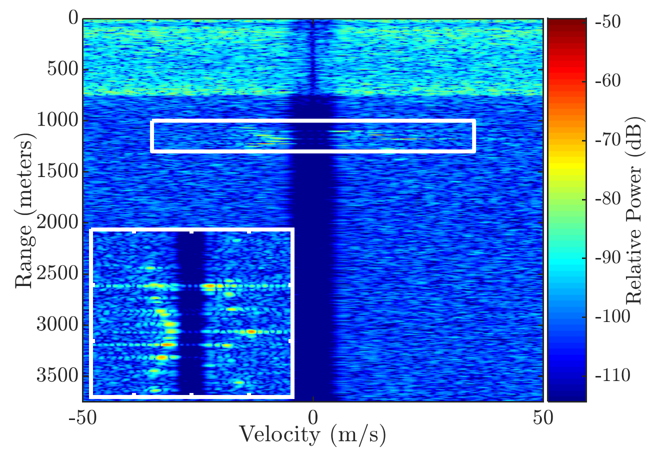

Fig. 5. NIMPC PCG-2 implementation ($I=10$) for open-air measurement using 100 pulsed PRO-FM waveforms

Figures 6 and 7 show the Proj-NIMPC implementations involving either direct inversion of the matrix in (34) or solving (35) via PCG-2, respectively, realize even better RSM suppression for the same 10 iterations. Indeed, the direct path response is now completely suppressed, and per Table II the latter result is achieved with the lowest computational cost.

As a final comparison using these open-air results, Fig. 8 depicts the convergence trends for the iterative NIMPC implementations, with original NIMPC [11] included for reference (lower dashed line) and standard uncoupled cancellation used to normalize (upper dashed line). The value plotted here is the average residual power over range interval [+10, +760] meters and across all Doppler, except for the clutter null region. While all four of these implementations do converge to the same level as full-dimension NIMPC performance, the Proj-NIMPC PCG-2 version clearly does so with the fewest iterations, thus underscoring its overall computational efficiency.

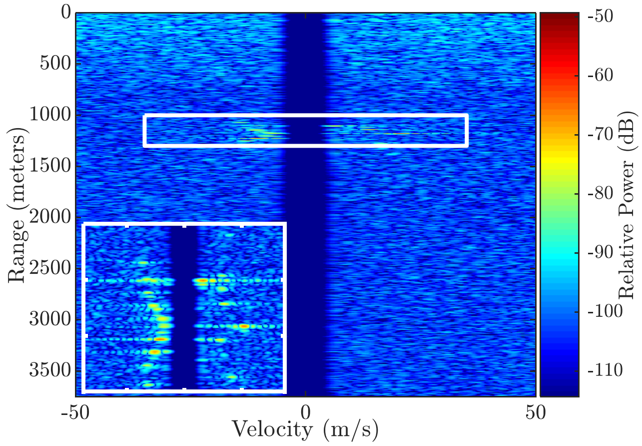

Fig. 6. Proj-NIMPC implementation (direct (34)) for open-air measurement using 100 pulsed PRO-FM waveforms

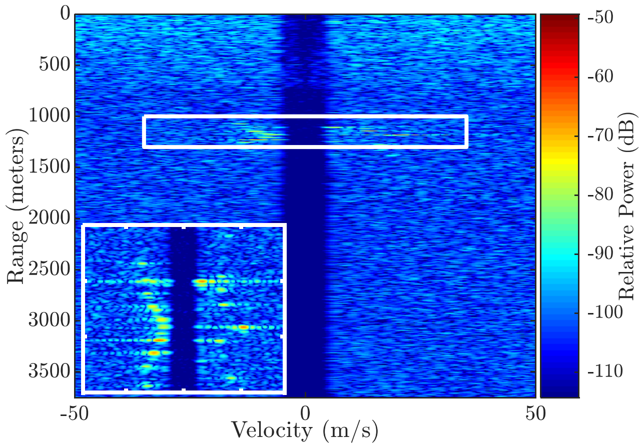

Fig. 7. Proj-NIMPC implementation (w/ PCG-2) for open-air measurement using 100 pulsed PRO-FM waveforms

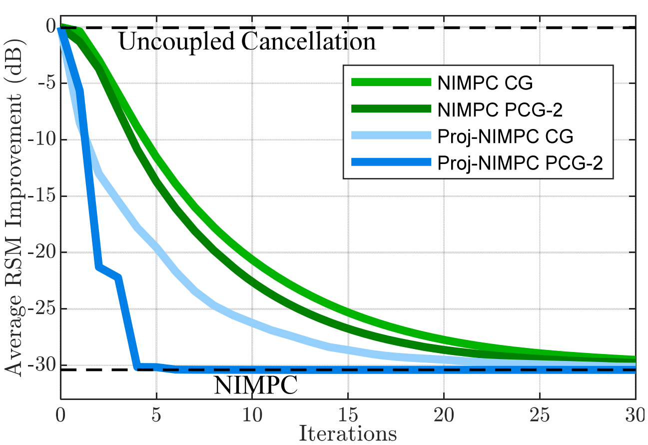

Fig. 8. Average direct path induced RSM versus iteration count

VII. Conclusions

Inspired by recent work [14] exploiting Toeplitz structure in the NIMPC method [11], which enables waveform-agile joint-domain clutter cancellation, CG [17-19] and ECA [20] schemes have been leveraged to reduce computational cost to a significant extent. Moreover, when these implementation approaches are combined, even greater efficiency improvement is realized, with experimentally demonstrated performance indistinguishable from that of original NIMPC.

References

[1] M.A. Richards, J.A. Scheer, W.A. Holm, eds., *Principles of Modern Radar: Basic Principles*, SciTech Publishing, 2010.

[2] S.D. Blunt, E.L. Mokole, "An overview of radar waveform diversity," *IEEE Aerospace & Electronic Systems Mag.*, vol. 31, no. 11, pp. 2-42, Nov. 2016.

[3] S.D. Blunt, M.R. Cook, J. Stiles, "Embedding information in radar emissions via waveform implementation," *Intl. Waveform Diversity & Design Conf.*, Niagara Falls, Canada, Aug. 2010.

[4] S.D. Blunt, J.K. Jakabosky, C.A. Mohr, P.M. McCormick, J.W. Owen, B. Ravenscroft, C. Sahin, G.D. Zook, C.C. Jones, J.G. Metcalf, T. Higgins, "Principles & applications of random FM radar waveform design," *IEEE Aerospace & Electronic Systems Mag.*, vol. 35, no. 10, pp. 20-28, Oct. 2020.

[5] B. Ravenscroft, J.W. Owen, J. Jakabosky, S.D. Blunt, A.F. Martone, K.D. Sherbondy, "Experimental demonstration and analysis of cognitive spectrum sensing and notching for radar," *IET Radar, Sonar & Navigation*, vol. 12, no. 12, pp. 1466–1475, Dec. 2018.

[6] A.C. O'Connor, "Frequency-domain synthesis of advanced pulse-compression filters," *IEEE Radar Conf.*, Boston, MA, Apr. 2019.

[7] J. Jakabosky, S.D. Blunt, B. Himed, "Spectral-shape optimized FM noise radar for pulse agility," *IEEE Radar Conf.*, Philadelphia, PA, May 2016.

[8] T. Yardibi, J. Li, P. Stoica, M. Xue, A.B. Baggeroer, "Source localization and sensing: a nonparametric iterative adaptive approach based on weighted least squares," *IEEE Trans. Aerospace & Electronic Systems*, vol. 46, no. 1, pp. 425-443, Jan. 2010.

[9] T. Higgins, S.D. Blunt, A.K. Shackelford, "Time-range adaptive processing for pulse agile radar," *Intl. Waveform Diversity & Design Conf.*, Niagara Falls, ON, Canada, Aug. 2010.

[10] B. Zhao, L. Kong, M. Yang, G. Cui, "Range-Doppler sidelobe and clutter suppression via time-range adaptive processing," *IEEE CIE Intl. Conf. on Radar*, Chengdu, China, Oct. 2011.

[11] T. Higgins, K. Gerlach, A.K. Shackelford, S.D. Blunt, "Aspects of non-identical multiple pulse compression," *IEEE Radar Conf.*, Kansas City, MO, May 2011.

[12] L.E. Brennan, L.S. Reed, "Theory of adaptive radar," *IEEE Trans. Aerospace & Electronic Systems*, vol. AES-9, no. 2, pp. 237-252, Mar. 1973.

[13] B. Ravenscroft, J.W. Owen, B.H. Kirk, S.D. Blunt, A.F. Martone, K.D. Sherbondy, R.M. Narayanan, "Experimental assessment of joint range-doppler processing to address clutter modulation from dynamic radar spectrum sharing," *IEEE Intl. Radar Conf.*, Washington, DC, May 2020.

[14] C. Sahin, J.G. Metcalf, B. Himed, "Reduced complexity maximum SINR receiver processing for transmit-encoded radar-embedded communications," *IEEE Radar Conf.*, Oklahoma City, OK, May 2018.

[15] D.S. Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, 2010.

[16] H. Akaike, "Block Toeplitz matrix inversion," *SIAM Journal of Applied Mathematics*, vol. 24, pp. 234–241, Mar. 1973.

[17] J. Nocedal, S.J. Wright, *Numerical Optimization: Second Edition*. Springer, 2006.

[18] F.-R. Lin, "Preconditioners for block Toeplitz systems based on circulant preconditioners" *Numerical Algorithms*, vol. 26, no. 4, pp. 365–379, April 2001.

[19] K.D. Czuprynski, J.B. Fahnline, S.M. Shontz, "Parallel boundary element solutions of block circulant linear systems for acoustic radiation problems with rotationally symmetric boundary surfaces," *ASME Noise Control & Acoustics Division Conf.*, Aug. 2012.

[20] F. Colone, D.W. O'Hagan, P. Lombardo, C.J. Baker, "A multistage processing algorithm for disturbance removal and target detection in passive bistatic radar," *IEEE Trans. Aerospace & Electronic Systems*, vol. 45, no. 2, pp. 698-722, Apr. 2009.