

Low-Complexity CRB-Based Beamforming Design for Downlink Monostatic ISAC Systems

Gabriela A. dos Santos, Caio G. de Figueredo, Walter C. Freitas Jr., F. Rodrigo P. Cavalcanti and Gábor Fodor, *Fellow, IEEE*

Abstract—This paper addresses joint transmit beamforming for a monostatic integrated sensing and communication base station with separate transmit and receive arrays. The goal is to design the precoder matrix that minimizes the Cramér–Rao bound for target direction estimation while satisfying per-user downlink signal-to-interference-noise ratio (SINR) requirements and individual per-target radar SINR floors. We propose a Split-Adam algorithm with an augmented Lagrangian that decouples sensing and communication beamformers, enforces downlink SINR constraints via closed-form iterative projection, and relies on analytic Wirtinger gradients of the Cramér–Rao bound that account for clutter coupling. Monte Carlo results confirm consistent gains over power-split benchmarks across a range of transmit powers and communication requirements.

Keywords—augmented Lagrangian, beamforming, Cramér–Rao bound, integrated sensing and communications.

I. INTRODUCTION

Integrated sensing and communication (ISAC) has emerged as a key design paradigm for sixth-generation (6G) wireless systems, allowing a single base station (BS) to share antennas, spectrum, and hardware between sensing and multi-user communication [1], [2]. Among the possible configurations, the monostatic setup, in which co-located antenna arrays transmit signals while simultaneously collecting echoes, is especially attractive for infrastructure deployments, requiring no dedicated sensing receiver while maximizing spectral efficiency [3], [4].

A widely adopted sensing metric in this context is the Cramér–Rao bound (CRB), which provides a fundamental lower bound on the variance of any unbiased target parameter estimate and is asymptotically tight for the maximum-likelihood estimate [5]. Minimising the CRB subject to communication quality-of-service (QoS) constraints directly captures the sensing–communication trade-off and produces beamformers that are optimal from an estimation-theoretic perspective [6].

Gabriela A. dos Santos, Caio G. de Figueredo, Walter C. Freitas Jr. and F. Rodrigo P. Cavalcanti are with the Postgraduate Program in Teleinformatics Engineering (PPGETI), Federal University of Ceará, and the Wireless Telecommunications Research Group (GTEL), Fortaleza-CE, Brazil, e-mail: {gabrielalbuquerque, caio, walter, rodrigo}@gtel.ufc.br.

Gábor Fodor is with Ericsson Research, 16480 Stockholm, Sweden, and also with the Division of Decision and Control, KTH Royal Institute of Technology, 11428 Stockholm, Sweden, e-mail: gabor.fodor@ericsson.com.

This work was supported in part by Ericsson Research, Sweden, and Ericsson Innovation Center, Brazil, under UFC.52 Technical Cooperation Contract Ericsson/UFC, in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. G. Fodor was supported by the Swedish Strategic Foundation Future Software Systems Project SAICOM, Grant No: FUS21-0004.

Liu et al. [6] established the seminal CRB-minimisation framework under per-user downlink signal-to-interference-plus-noise ratio (SINR) constraints and proved that semidefinite relaxation (SDR) is tight in certain instances, at the cost of $\mathcal{O}(N_t^6)$, where N_t is the number of transmit antennas, per-iteration complexity. For monostatic systems with simultaneous full-duplex communication, He et al. [3] proposed SDR and successive convex approximation (SCA) algorithms for joint transceiver and power allocation, still relying on interior-point solvers. More recently, Fang et al. [7], [8] characterised the optimal beamforming structure for CRB-based ISAC and developed SCA algorithms that avoid semidefinite lifting but still call a convex solver at each step. Hua et al. [9] studied the fundamental CRB–rate trade-off in multiple-input multiple-output ISAC, Zhao et al. [10] addressed robust CRB optimisation under channel uncertainty, and Liu et al. [11] addressed joint beamforming and reconfigurable intelligent surfaces design under CRB and signal-to-noise ratio constraints. These solver-based approaches are effective yet carry a computational footprint that scales unfavourably with the number of antennas.

This paper proposes a solver-free, gradient-based framework with the following contributions:

- Closed-form Wirtinger gradients of the CRB are derived, accounting for the coupling of the full beamformer matrix into the interference-plus-noise covariance through the clutter channel, a term absent from prior gradient derivations [7], [8].
- Individual per-target SINR constraints are included, enforced by an augmented Lagrangian with analytic gradient, allowing the optimiser to allocate resources across targets independently rather than through a single aggregate sensing metric.
- A Split-Adam strategy is adopted, maintaining separate Adam optimisers for the sensing and communication beam vectors and preventing downlink penalty forces from distorting the beamformer.
- Downlink SINR constraints are enforced by a closed-form iterative projection at every iteration, guaranteeing feasibility throughout without accumulating large multipliers for the communication side.
- Optimisation terminates via a relative CRB improvement criterion over a sliding window, with no fixed iteration count.

The remainder of this paper is organised as follows. Section II presents the system model and Section III the optimi-

sation problem. Section IV describes the proposed algorithm. Simulation results and conclusions are given in Sections V and VI, respectively.

Notation: Boldface lowercase and uppercase letters denote vectors and matrices. $(\cdot)^H$ is the conjugate transpose, $\text{tr}(\cdot)$ is the trace, $\|\cdot\|_F$ is the Frobenius norm, and $\mathbb{C}^{m \times n}$ is the set of complex $m \times n$ matrices.

II. SYSTEM MODEL

A monostatic ISAC BS equipped with N_t transmit and N_r receive antennas, both modelled as uniform linear arrays (ULAs), is considered. The BS simultaneously serves K single-antenna downlink users and senses M point targets in the presence of L passive clutter scatterers. Residual self-interference after analogue cancellation is modelled as $\mathbf{H}_{\text{SI}} \in \mathbb{C}^{N_r \times N_t}$, and is assumed to be perfectly known at the BS, as is standard in the monostatic ISAC literature [3].

The transmitted symbol vector satisfies $\mathbf{s} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$, so the transmit signal is $\mathbf{W}\mathbf{s}$, where $\mathbf{W} = [\mathbf{W}_s \ \mathbf{W}_c] \in \mathbb{C}^{N_t \times (M+K)}$ is the joint beamforming matrix, $\mathbf{W}_s \in \mathbb{C}^{N_t \times M}$ contains the sensing-dedicated beam vectors, and $\mathbf{W}_c \in \mathbb{C}^{N_t \times K}$, the information-bearing columns. The total transmit power is $\mathbb{E}[\|\mathbf{W}\mathbf{s}\|^2] = \|\mathbf{W}\|_F^2$, yielding the power budget $\|\mathbf{W}\|_F^2 \leq P_b$.

Each target m is modelled as a point scatterer with complex reflectivity β_0 at angle θ_m . The received signal at the N_r -element array is

$$\mathbf{y} = \sum_{m=1}^M \beta_0 \mathbf{a}(\theta_m, N_r) \mathbf{a}^\top(\theta_m, N_t) \mathbf{W}\mathbf{s} + \mathbf{B}\mathbf{W}\mathbf{s} + \mathbf{n}, \quad (1)$$

where $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I}_{N_r})$ is receiver noise and $\mathbf{B}$ is the aggregate interference channel defined below.

For a N -element ULA with half-wavelength spacing, the steering vector at angle θ is $\mathbf{a}(\theta, N) = \frac{1}{\sqrt{N}} [1, e^{j\pi \sin \theta}, \dots, e^{j(N-1)\pi \sin \theta}]^\top$, with derivative $\dot{\mathbf{a}}(\theta, N) = \partial \mathbf{a} / \partial \theta$. The aggregate interference channel at the receive array combines clutter and self-interference as

$$\mathbf{B} = \mathbf{H}_{\text{SI}} + \sum_{l=1}^L \beta_{c,l} \mathbf{a}(\gamma_l, N_r) \mathbf{a}^\top(\gamma_l, N_t), \quad (2)$$

where $\beta_{c,l}$ and γ_l are the complex reflectivity and angle of the l -th clutter patch. The interference-plus-noise covariance is then

$$\Psi(\mathbf{W}) = \sigma^2 \mathbf{I}_{N_r} + \mathbf{B}\mathbf{W}\mathbf{W}^H \mathbf{B}^H, \quad (3)$$

with σ^2 as the noise power. Note that Ψ depends on the full $\mathbf{W}$, including the communication columns, through $\mathbf{B}$.

The downlink SINR for user k is [6]

$$\text{SINR}_{c,k} = \frac{|\mathbf{g}_k^H \mathbf{w}_{c,k}|^2}{\sum_{j \neq k} |\mathbf{g}_k^H \mathbf{w}_{c,j}|^2 + \sigma^2}, \quad (4)$$

where $\mathbf{g}_k \in \mathbb{C}^{N_t}$ is the channel vector to user k and $\mathbf{w}_{c,k}$ is the k -th column of $\mathbf{W}_c$. With the minimum variance distortionless response (MVDR) receive beamformer $\mathbf{u}_m^* = \Psi^{-1} \mathbf{a}(\theta_m, N_r)$, which maximises the output SINR for target m , the sensing SINR for target m at angle θ_m is [3]

$$\text{SINR}_{s,m} = |\beta_0|^2 (\mathbf{a}_t^H \mathbf{W} \mathbf{W}^H \mathbf{a}_t^*) (\mathbf{a}_r^H \Psi^{-1} \mathbf{a}_r), \quad (5)$$

where $\mathbf{a}_t = \mathbf{a}(\theta_m, N_t)$, $\mathbf{a}_r = \mathbf{a}(\theta_m, N_r)$, and β_0 is the target reflectivity. The expression in (5) follows from applying the MVDR beamformer output to (1): the first factor measures the beamforming gain towards target m at the transmitter, while the second is the output SINR of the MVDR filter at the receiver.

III. PROBLEM FORMULATION

By the Slepian–Bangs formula [5], the Fisher information for the angle of target m , averaged over $\mathbf{s} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$, is

$$I_m(\mathbf{W}) = 2 \text{Re} \{ \text{tr}((\mathbf{D}_m \mathbf{W})^H \Psi^{-1} (\mathbf{D}_m \mathbf{W})) \}, \quad (6)$$

where $\mathbf{D}_m = \sqrt{|\beta_0|^2} (\dot{\mathbf{a}}_r \mathbf{a}_t^H + \mathbf{a}_r \dot{\mathbf{a}}_t^H)$. The per-target CRB is the inverse of the scalar Fisher information, $I_m^{-1}(\mathbf{W})$, so the total CRB summed over all targets is $\text{CRB}(\mathbf{W}) = \sum_{m=1}^M I_m^{-1}(\mathbf{W})$. The joint beamforming problem is

$$\underset{\mathbf{W}}{\text{minimize}} \quad \text{CRB}(\mathbf{W}) \quad (7)$$

$$\text{such that} \quad \text{SINR}_{c,k} \geq \tau_{c,k}, \quad k = 1, \dots, K, \quad (8)$$

$$\text{SINR}_{s,m} \geq \tau_{s,m}, \quad m = 1, \dots, M, \quad (9)$$

$$\|\mathbf{W}\|_F^2 \leq P_b. \quad (10)$$

Constraint (8) guarantees per-user communication QoS. Constraint (9) sets individual SINR floors per target [11], [14], ensuring that no target is completely starved of beamforming energy—a property that a single aggregate constraint cannot enforce. Problem (7)–(10) is non-convex due to the fractional coupled structure of $\text{CRB}(\mathbf{W})$ and $\text{SINR}_{s,m}$, and to the dependence of Ψ on all columns of $\mathbf{W}$.

IV. PROPOSED ALGORITHM

A. Wirtinger Gradients

The Wirtinger gradient of $\text{CRB}(\mathbf{W})$ with respect to $\mathbf{W}^*$ is derived by applying the chain rule in Wirtinger calculus [13] to (6) and accounting for the two paths through which $\mathbf{W}^*$ enters: directly via $\mathbf{W}^H$, and indirectly via $\Pi = \Psi^{-1}(\mathbf{W})$:

$$\nabla_{\mathbf{W}^*} \text{CRB} = \sum_{m=1}^M \frac{-1}{I_m^2} [2 \mathbf{D}_m^H \Pi \mathbf{D}_m \mathbf{W} - 2 \mathbf{B}^H \Pi \mathbf{D}_m \mathbf{W} \mathbf{W}^H \mathbf{D}_m^H \Pi \mathbf{B} \mathbf{W}], \quad (11)$$

1) *Proof of $\nabla_{\mathbf{W}^*} \text{CRB}$:* Since $\text{CRB}(\mathbf{W}) = \sum_{m=1}^M I_m^{-1}(\mathbf{W})$, the chain rule gives

$$\nabla_{\mathbf{W}^*} \text{CRB} = \sum_{m=1}^M \frac{-1}{I_m^2} \nabla_{\mathbf{W}^*} I_m, \quad (12)$$

so it suffices to compute $\nabla_{\mathbf{W}^*} I_m$. Since Ψ is Hermitian, I_m may be written as

$$I_m = 2 \text{Re} \{ \text{tr}(\mathbf{W}^H \mathbf{D}_m^H \Pi \mathbf{D}_m \mathbf{W}) \}. \quad (13)$$

Direct term. Treating Π as constant and using $\nabla_{\mathbf{W}^*} \text{tr}(\mathbf{W}^H \mathbf{A}) = \mathbf{A}$,

$$\nabla_{\mathbf{W}^*} I_m|_{\text{dir}} = 2 \mathbf{D}_m^H \Pi \mathbf{D}_m \mathbf{W}. \quad (14)$$

Indirect term. Differentiating the identity $\Psi \Pi = \mathbf{I}$ yields $d\Pi = -\Pi(d\Psi)\Pi$. The variation of $\Psi = \sigma^2 \mathbf{I} +$

$\mathbf{B}\mathbf{W}\mathbf{W}^H\mathbf{B}^H$ with respect to $\mathbf{W}^*$ (holding $\mathbf{W}$ fixed) is $d\Psi = \mathbf{B} d\mathbf{W} \mathbf{W}^H\mathbf{B}^H$, so

$$d\Pi = -\Pi\mathbf{B} d\mathbf{W} \mathbf{W}^H\mathbf{B}^H\Pi. \quad (15)$$

Substituting into I_m and applying the cyclic property of the trace to isolate $d\mathbf{W}$,

$$dI_m|_{\text{indir}} = -2 \operatorname{Re}\{\operatorname{tr}(\mathbf{W}^H\mathbf{D}_m^H\Pi\mathbf{D}_m\mathbf{W}\mathbf{W}^H\mathbf{B}^H\Pi\mathbf{B} d\mathbf{W})\}, \quad (16)$$

from which the corresponding Wirtinger gradient is

$$\nabla_{\mathbf{W}^*} I_m|_{\text{indir}} = -2\mathbf{B}^H\Pi\mathbf{D}_m\mathbf{W}\mathbf{W}^H\mathbf{D}_m^H\Pi\mathbf{B}\mathbf{W}. \quad (17)$$

Adding (14) and (17) into (12) yields (11). The indirect term, which arises from the dependence of Ψ on $\mathbf{W}$ through the clutter channel $\mathbf{B}$, is absent from derivations that assume white noise [7], [8].

2) *Gradient of the Augmented Lagrangian:* The sensing augmented Lagrangian term is $\mathcal{L}_m = \mu_m\nu_m + (\rho_s/2)\nu_m^2$, where $\nu_m = \max(0, \tau_{s,m} - \text{SINR}_{s,m})$ is the linear-scale violation, $\mu_m \geq 0$ is the dual variable, and $\rho_s > 0$ is the penalty parameter. Its gradient is

$$\nabla_{\mathbf{W}^*}\mathcal{L}_m = -(\mu_m + \rho_s\nu_m)\nabla_{\mathbf{W}^*}\text{SINR}_{s,m}, \quad (18)$$

so it suffices to derive $\nabla_{\mathbf{W}^*}\text{SINR}_{s,m}$.

Rewriting (5) as $\text{SINR}_{s,m} = |\beta_0|^2 f_1 f_2$, where $f_1 = \mathbf{a}_t^H \mathbf{Q} \mathbf{a}_t$ and $f_2 = \mathbf{a}_r^H \Pi \mathbf{a}_r$, the product rule gives

$$\nabla_{\mathbf{W}^*}\text{SINR}_{s,m} = |\beta_0|^2 [(\nabla_{\mathbf{W}^*} f_1) f_2 + f_1 (\nabla_{\mathbf{W}^*} f_2)]. \quad (19)$$

Direct term. Since $f_1 = \mathbf{a}_t^H \mathbf{W} \mathbf{W}^H \mathbf{a}_t$ depends on $\mathbf{W}^*$ only through $\mathbf{W}^H$,

$$\nabla_{\mathbf{W}^*} f_1 = \mathbf{a}_t \mathbf{a}_t^H \mathbf{W}. \quad (20)$$

Indirect term. $f_2 = \mathbf{a}_r^H \Pi \mathbf{a}_r$ depends on $\mathbf{W}^*$ only through $\Pi = \Psi^{-1}(\mathbf{W})$. Using $d\Pi = -\Pi\mathbf{B} d\mathbf{W} \mathbf{W}^H\mathbf{B}^H\Pi$ from the proof of (11),

$$df_2 = -\mathbf{a}_r^H \Pi \mathbf{B} d\mathbf{W} \mathbf{W}^H \mathbf{B}^H \Pi \mathbf{a}_r, \quad (21)$$

from which the Wirtinger gradient is

$$\nabla_{\mathbf{W}^*} f_2 = -\mathbf{B}^H (\Pi \mathbf{a}_r \mathbf{a}_r^H \Pi) \mathbf{B} \mathbf{W}. \quad (22)$$

Combining both terms yields

$$\begin{aligned} \nabla_{\mathbf{W}^*}\text{SINR}_{s,m} &= |\beta_0|^2 (\mathbf{a}_r^H \Pi \mathbf{a}_r) \mathbf{a}_t \mathbf{a}_t^H \mathbf{W} \\ &\quad - |\beta_0|^2 (\mathbf{a}_t^H \mathbf{Q} \mathbf{a}_t) \mathbf{B}^H (\Pi \mathbf{a}_r \mathbf{a}_r^H \Pi) \mathbf{B} \mathbf{W}. \end{aligned} \quad (23)$$

As in (11), the indirect term captures the feedback of the beamformer into Ψ .

B. Split-Adam with Closed-Form SINR_{DL} Projection

The key structural insight is that $\mathbf{W}_s$ drives CRB minimisation while $\mathbf{W}_c$ is shaped by the communication constraints. Coupling their gradients in a single Adam optimiser causes the downlink penalty to distort the sensing columns when multipliers grow large. We therefore maintain independent Adam first and second moment estimates for $\mathbf{W}_s$ and $\mathbf{W}_c$.

At each iteration, $\mathbf{W}_s$ is updated by the combined gradient

$$\mathbf{d}_s = \left[\nabla_{\mathbf{W}^*}\text{CRB} + \sum_m \nabla_{\mathbf{W}^*}\mathcal{L}_m \right]_{:,1:M}. \quad (24)$$

This gradient is passed to the Adam optimiser with learning rate $\eta > 0$ and standard moment hyperparameters [12]. No explicit gradient is applied to $\mathbf{W}_c$; instead, after each powerball projection $\|\mathbf{W}\|_F^2 \leq P_b$, the downlink SINR constraints are restored by rescaling each communication column. Imposing $\text{SINR}_{c,k} = \tau_{c,k}$ exactly and seeking a scalar α such that $\mathbf{w}_{c,k} \leftarrow \alpha \mathbf{w}_{c,k}$, substitution into (4) gives

$$\frac{\alpha^2 |\mathbf{g}_k^H \mathbf{w}_{c,k}|^2}{I_k + \sigma^2} = \tau_{c,k}, \quad (25)$$

where $I_k = \sum_{j \neq k} |\mathbf{g}_k^H \mathbf{w}_{c,j}|^2$. Solving for α and applying the rescaling only when the constraint is violated yields

$$\mathbf{w}_{c,k} \leftarrow \mathbf{w}_{c,k} \sqrt{\frac{\tau_{c,k}(I_k + \sigma^2)}{|\mathbf{g}_k^H \mathbf{w}_{c,k}|^2}}, \quad \text{if } \text{SINR}_{c,k} < \tau_{c,k}. \quad (26)$$

Up to N_{proj} such passes handle the coupling between users. This projection enforces feasibility at every step without accumulating large multipliers for the DL side. The sensing multipliers are updated every T_{upd} iterations as $\mu_m \leftarrow \text{clip}(\mu_m + \rho_s \nu_m, 0, \mu_{\text{max}})$, where $\mu_{\text{max}} > 0$ is an upper bound that prevents multiplier explosion during early iterations.

Algorithm 1 summarises the procedure. Letting $C_t = \text{CRB}(\mathbf{W})$ at iteration t , optimisation terminates when $|C_t - C_{t-p}|/C_{t-p} < \epsilon_{\text{rel}}$ for a sliding window of p iterations, with T_{max} as a safety bound. The initialisation places $\mathbf{W}_c$ at the zero-forcing (ZF) solution [6] scaled to achieve $\text{SINR}_{c,k} = \tau_{c,k}$, and $\mathbf{W}_s$ at matched-filter (MF) columns $\mathbf{a}(\theta_m, N_t)$, guaranteeing feasibility from the start. Additional seeds perturb $\mathbf{W}_s$ with 30% complex Gaussian noise, and the best feasible result is retained.

Algorithm 1 Split-Adam with Augmented Lagrangian

Require: $\mathbf{H}_{\text{SI}}, \mathbf{G} = [\mathbf{g}_1, \dots, \mathbf{g}_K], \{\theta_m\}, \{\gamma_l\}, P_b, \{\tau_{c,k}\}, \{\tau_{s,m}\}$

Ensure: $\mathbf{W}^*$

```

1: Initialise  $\mathbf{W}$  (ZF+MF),  $\mu_m = 0$ , Adam moments = 0
2: for  $t = 1, 2, \dots, T_{\text{max}}$  do
3:   Compute  $\Pi = \Psi^{-1}(\mathbf{W})$ ,  $\{I_m\}$ ,  $\text{CRB}(\mathbf{W})$ 
4:   Compute  $\mathbf{d}_s$ ; perform Adam step on  $\mathbf{W}_s$ 
5:   Project  $\mathbf{W}$  onto  $\|\mathbf{W}\|_F^2 \leq P_b$ 
6:   Enforce  $\text{SINR}_{c,k} \geq \tau_{c,k}$  via (26) ( $N_{\text{proj}}$  passes)
7:   if  $t \bmod T_{\text{upd}} = 0$  then
8:      $\mu_m \leftarrow \text{clip}(\mu_m + \rho_s \nu_m, 0, \mu_{\text{max}})$ ,  $\forall m$  with active  $\tau_{s,m}$ 
9:   end if
10:  if  $|C_t - C_{t-p}|/C_{t-p} < \epsilon_{\text{rel}}$  then
11:    break
12:  end if
13: end for
14: return best feasible  $\mathbf{W}$ 

```

The dominant per-iteration cost is Ψ^{-1} in $\mathcal{O}(N_r^3)$ and the gradient in (11) in $\mathcal{O}(N_t^2 N_r M)$. No inner convex program is solved and no auxiliary variables are introduced.




Fig. 1: Normalised transmit beampattern of the optimised $\mathbf{W}$ for a fixed realisation. Targets at 0° and 30° (red dash-dot), users at -40° and 60° (blue dash-dot), and clutter scatterers at -50° and 20° (grey dotted).

For $N_t = N_r = 8$ and $M = 2$, each iteration completes in a few milliseconds on a standard central processing unit (CPU), compared to tens of milliseconds for an SCA-SDR iteration [8], with convergence typically occurring in 250–450 iterations.

V. SIMULATION RESULTS

The system parameters are $N_t = N_r = 8$, $M = 2$ targets, $K = 2$ users, and $L = 2$ clutter scatterers. The transmit power budget is $P_b = 18$ dBW, the noise power is $\sigma^2 = -70$ dBm, the target reflectivity satisfies $|\beta_0|^2/\sigma^2 = -20$ dB, the clutter-to-noise ratio is 10 dB, and the residual self-interference level is $\alpha_{SI} = -110$ dB. The downlink path loss is $p_l = -103.6$ dB. All Monte Carlo results are averaged over $N_{MC} = 100$ independent channel realisations, with target, user, and clutter angles drawn uniformly from $[-60^\circ, 60^\circ]$ with a minimum angular separation of 5° . The algorithm is run with $N_{seeds} = 5$ multi-start initialisations, $\rho_s = 10$, $T_{upd} = 10$, $\mu_{max} = 10^3$, learning rate $\eta = 8 \times 10^{-3}$, $\epsilon_{rel} = 10^{-4}$, and $T_{max} = 600$.

Transmit beampattern. Fig. 1 shows the normalised transmit beampattern produced by the optimised $\mathbf{W}$ for a fixed channel realisation with targets at 0° and 30° , users at -40° and 60° , and clutter scatterers at -50° and 20° . The total beampattern (black) exhibits peaks at both target directions, driven primarily by $\mathbf{W}_s$ (red dashed), while $\mathbf{W}_c$ (blue dashed) concentrates energy towards the user directions. Notably, the total pattern places a relative null near the clutter direction at 20° , illustrating that the gradient term in (11) that accounts for $\Psi(\mathbf{W})$ actively suppresses interference at the receive array by shaping the transmit beamformer. The sensing and communication beamformers exhibit complementary spatial behaviour, confirming that the Split-Adam strategy successfully decouples their respective objectives without forcing the sensing columns to absorb the communication penalty.

CRB–communication trade-off. Fig. 2 shows the total CRB as a function of the downlink SINR requirement $\tau_{c,k}$, compared against three benchmarks: matched-filter (BM1: MF), zero-forcing combined with matched filter (BM2: ZF-MF), and an optimised power-split scheme (BM3) that allocates power between sensing and communication but does not




Fig. 2: Total CRB vs. downlink SINR requirement $\tau_{c,k}$, averaged over $N_{MC} = 100$ realisations. Annotations indicate the CRB gain (in dB) of the proposed method over BM3.




Fig. 3: Total CRB vs. transmit power budget P_b , averaged over $N_{MC} = 100$ realisations.

jointly optimise the beamformer shape. BM1 and BM2 exhibit a CRB that is nearly constant and independent of $\tau_{c,k}$, since neither adapts its beamformer shape to the communication constraint. The proposed method outperforms BM3 for $\tau_{c,k} \leq 12$ dB, with a gain of +1.1 dB at $\tau_{c,k} = 6$ dB that decreases monotonically as the communication requirement tightens. Beyond $\tau_{c,k} = 12$ dB, the CRB of the proposed method exceeds that of BM3, reaching -1.8 dB at $\tau_{c,k} = 18$ dB. This behaviour reflects the increasing cost of satisfying the per-user SINR constraints via gradient-based optimisation: as the feasible set shrinks, the projection onto the communication constraint consumes power that would otherwise be directed toward sensing, whereas BM3's fixed power-split structure incurs no such penalty. The right axis confirms that the proposed algorithm satisfies the SINR constraint accurately throughout the sweep, validating the closed-form projection in (26).

CRB versus transmit power. Fig. 3 illustrates the sensing–communication trade-off as a function of P_b . Three operating points are compared: the proposed method under $\tau_{c,k} = 12$ dB, the same optimiser with the downlink constraint relaxed to near zero (CRB only), and a sensing-only baseline with $K = 0$. At low power ($P_b = 10$ dBW), the proposed method incurs




Fig. 4: Per-target CRB vs. sensing SINR threshold $\tau_{s,1}$, with $\tau_{s,2}$ fixed at its nominal value (dashed line), averaged over $N_{MC} = 100$ realisations.

a penalty of nearly one order of magnitude in CRB relative to the unconstrained case, as a significant fraction of the power budget must be dedicated to satisfying the communication QoS. As P_b increases, the gap narrows and the three curves converge at $P_b = 22$ dBW, where the $\tau_{c,k}$ constraint is satisfied with a small fraction of the total power. The near-identical behaviour of the CRB-only and sensing-only curves across all power levels confirms that, when communication constraints are removed, the optimiser allocates resources in a manner equivalent to a dedicated sensing system.

Per-target CRB versus individual sensing SINR threshold. Fig. 4 demonstrates the effect of the individual per-target SINR constraints introduced in (9). The threshold $\tau_{s,1}$ is swept from -8.5 dB to -6.5 dB while $\tau_{s,2}$ remains fixed at -7.5 dB (nominal, indicated by the dashed vertical line). As $\tau_{s,1}$ increases, the constraint $\text{SINR}_{s,1} \geq \tau_{s,1}$ becomes harder to satisfy, causing the augmented Lagrangian multiplier μ_1 to accumulate larger values during optimisation and redirecting beamforming energy towards Target 1 at the expense of Target 2 under the fixed power budget P_b . This cross-target interference effect is the central motivation for individual per-target constraints: a single aggregate sensing metric would not prevent the optimiser from concentrating all sensing energy on one target at the expense of the other. Near the nominal threshold, both targets achieve comparable and jointly minimised CRBs, confirming that $\tau_{s,m} = -7.5$ dB is a well-calibrated operating point for this scenario.

VI. CONCLUSION

A solver-free beamforming framework for downlink monostatic ISAC systems was proposed, combining three key elements: closed-form Wirtinger gradients of the CRB incorporating clutter coupling through $\Psi(\mathbf{W})$, an Augmented Lagrangian with per-target dual variables to prevent target starvation, and a Split-Adam optimiser that separates sensing and communication updates to avoid gradient interference. Downlink SINR feasibility is enforced at every iteration through closed-form column rescaling, removing the need for large communication-side multipliers. The resulting per-iteration complexity is $\mathcal{O}(N_t^2 N_r M)$ with no inner convex program, achieving convergence in 250–450 iterations with millisecond-scale runtime on a standard CPU.

Monte Carlo simulations show consistent CRB improvements over power-split benchmarks across varying transmit powers and communication requirements, with gains up to 1.1 dB at low SINR thresholds. The optimised transmit beampattern simultaneously directs energy towards targets and users while suppressing clutter, an effect that emerges naturally from the clutter-aware gradient without explicit null-steering constraints. The per-target SINR formulation also provides effective control of sensing resource allocation between targets.

Several extensions remain open for future work. Robust formulations accounting for uncertainty in $\mathbf{H}_{S1}$ and clutter parameters γ_l, β_c, l would improve practical applicability. The point-scatterer model could be extended to fluctuating radar cross-sections, while replacing the line-of-sight ULA communication model with multipath fading would better reflect realistic cellular settings. On the algorithmic side, proving convergence of Split-Adam to a stationary point of the augmented Lagrangian remains an important theoretical challenge. Finally, improving multi-start and warm-start strategies for larger antenna arrays could help preserve the framework's low-complexity advantage as system dimensions scale.

REFERENCES

- [1] F. Liu, Y. Cui, C. Masouros, J. Xu, T. X. Han, Y. C. Eldar, and S. Buzzi, "Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 7, pp. 1728–1767, Jul. 2022.
- [2] J. A. Zhang, M. L. Rahman, K. Wu, X. Huang, Y. J. Guo, S. Chen, and J. Yuan, "Enabling joint communication and radar sensing in mobile networks—a survey," *IEEE Commun. Surveys Tuts.*, vol. 24, no. 1, pp. 306–345, 2022.
- [3] Z. He, W. Xu, H. Shen, D. W. K. Ng, Y. C. Eldar, and X. You, "Full-duplex communication for ISAC: Joint beamforming and power optimization," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 9, pp. 2920–2936, Sep. 2023.
- [4] S. Lu, F. Liu, and L. Hanzo, "The degrees-of-freedom in monostatic ISAC channels: NLoS exploitation vs. reduction," *IEEE Trans. Veh. Technol.*, vol. 72, no. 2, pp. 2643–2648, Feb. 2023.
- [5] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. Wiley, New York, 2002.
- [6] F. Liu, Y.-F. Liu, A. Li, C. Masouros, and Y. C. Eldar, "Cramér–Rao bound optimization for joint radar-communication beamforming," *IEEE Trans. Signal Process.*, vol. 70, pp. 240–253, 2022.
- [7] T. Fang, H. Mohammadi, H. Ruan, G. Zheng, and B. Ottersten, "Optimal ISAC beamforming structure and efficient algorithms for sum rate and CRLB balancing," *arXiv:2503.09489*, 2025.
- [8] T. Fang, N. T. Nguyen and M. Juntti, "Low-Complexity Cramér–Rao Lower Bound and Sum Rate Optimization in ISAC Systems," *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Hyderabad, India, 2025, pp. 1–5, doi: 10.1109/ICASSP49660.2025.10888351.
- [9] H. Hua, J. Xu, and T. X. Han, "MIMO integrated sensing and communication: CRB-rate tradeoff," *IEEE Trans. Wireless Commun.*, vol. 23, no. 3, pp. 2839–2854, Mar. 2024.
- [10] Z. Zhao, L. Zhang, R. Jiang, and D. W. K. Ng, "Joint beamforming scheme for ISAC systems via robust Cramér–Rao bound optimization," *IEEE Wireless Commun. Lett.*, vol. 13, no. 4, pp. 889–893, Apr. 2024.
- [11] R. Liu, M. Li, Q. Liu, and A. L. Swindlehurst, "SNR/CRB-constrained joint beamforming and reflection designs for RIS-ISAC systems," *IEEE Trans. Wireless Commun.*, vol. 23, no. 7, pp. 7456–7470, 2024.
- [12] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, San Diego, CA, 2015.
- [13] A. Hjørungnes and D. Gesbert, "Complex-valued matrix differentiation: Techniques and key results," *IEEE Trans. Signal Process.*, vol. 55, no. 8, pp. 4357–4362, Aug. 2007.
- [14] I. Siomina, A. Furuskär, G. Fodor, "A mathematical framework for statistical QoS and capacity studies in OFDM networks", *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, Tokyo, Japan, Sept. 2009.