electronics

MDPI

Article

Factorised Machine-Learning Inversion of 5G Coexistence Masks on Compact MIMO Apertures with Electrically Orthogonal Metamaterial Operators

Wyssem Fathallah [1,2], Dhaou Bouchouicha [3], Mongi Ben Moussa [4], Abdelilah Mansour [5] and Hedi Sakli [1,*]

[1] SYS'COM Laboratory LR99ES21, National Engineering School of Tunis, Tunis El Manar University, Tunis 1002, Tunisia; fathallahwyssem@gmail.com
[2] ISITCOM, University of Sousse, Sousse 4011, Tunisia
[3] Proceeds, Energies, Environment and Electrical Systems Laboratory, National Engineering School of Gabes, University of Gabes, Gabes 6029, Tunisia; bouchouicha.dhaou@gmail.com
[4] Departement of Physics, College of Sciences, Umm Al-Qura University, Makkah 21955, Saudi Arabia; msbenmoussa@uqu.edu.sa
[5] SERCOM Laboratory, Tunisia Polytechnic School, University of Carthage, La Marsa 2078, Tunisia; e_almansoury@hotmail.com
* Correspondence: saklihedi12@gmail.com

Abstract

Compact ultra-wideband (UWB) multiple-input multiple-output (MIMO) front-ends must stay matched over a wide continuum, reject licensed 5G interferers and remain uncorrelated at sub-wavelength spacing. This paper formulates coexistence as the inversion of a mask $M = (N(f), I(f), \text{rho-max}, \text{eta-min})$ and treats complementary split-ring resonators (CSRRs) and a split-ring resonator (SRR) fence as operators acting on electrically distinct regions. Sufficient conditions for a conditionally factorised inverse are derived (bounded interaction on the support of $N(f)$, bounded cross-sensitivity and a well-conditioned rejection Jacobian), together with a first-order bound on the inversion error. The rejection factor is implemented by inverse regressors with a forward-consistency check and an out-of-domain detector (random forest and k-nearest neighbours, R-squared = 0.9825 and 0.9921 on unit-cell data); non-uniqueness is handled by selecting, among forward-consistent candidates, the geometry with the largest fabrication margin. While ten target specifications are queried across the C-band and WLAN windows, only the two fabricated unit cells have undergone dedicated full-wave and experimental validation, whereas the remaining eight are compared against their nearest training-set neighbours, leaving multi-target generalisation across the specification space only partially validated. A two-element FR4 prototype designed as an experimental testbed for the union of a C-band and a WLAN mask, with 10 mm edge-to-edge spacing (0.07 wavelength at 2.1 GHz and 0.097 at the 2.9 GHz lower matched edge in full wave), is measured and scored against the mask. In measurement, frequency-domain coupling stays below −20 dB from 2.1 to about 14.8 GHz (full wave: 67 % of the band). Rejection is the weak coordinate: the measured reflection of the two-port board peaks at about −9.5 dB in both stop-bands, against −2.9 and −2.0 dB in full wave, so the −10 dB criterion is met over about 44 % of the combined mask, and the C-band mask is covered only from about 3.9 GHz. This partial compliance exposes the effects of uncompensated unit-cell-to-aperture transfer offsets, stop-band bandwidth mismatch and FR4 Q-factor damping, showing that the prototype demonstrates the diagnostic scoring pipeline rather than a fully compliant hardware realisation of coexistence-mask inversion. The four-configuration test, carried out in full wave at the fabricated design, shows an interaction residual (0.61 in magnitude norm) far above the rejection-margin threshold (0.019), so condition (C1) is not met and factorised inversion remains a strictly conditional formulation that is not experimentally validated by

Check for updates

Academic Editor: Yan Zhang

Received: 20 September 2026
Revised: 24 October 2026
Accepted: 30 October 2026
Published: 2 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

the current prototype. Moreover, because pattern-based correlation is evaluated only from simulated patterns at six spot frequencies (ECC at most 0.015) without measured pattern validation, and because group delay and fidelity factor have not been assessed, wideband performance claims are restricted to continuous-wave frequency-domain matching and isolation. The changes needed to close this gap, the tolerance budget of both cells and the validation that remains open are stated explicitly, and modelled results are separated from measured ones.

Keywords: coexistence mask; inverse design; MIMO antenna; metamaterial operators; fifth-generation interference; machine learning

1. Introduction

Fifth-generation (5G) terminals that occupy the unlicensed ultra-wideband (UWB) allocation 3.1 GHz to 10.6 GHz [1] illustrate a computational loop between a communication specification and compact hardware. Printed monopoles provide compact, low-cost radiators over that continuum [2–4], but the same spectrum overlaps licensed services—in particular, the C-band satellite downlink (3.7 GHz to 4.2 GHz) and WLAN (5.15 GHz to 5.35 GHz and 5.725 GHz to 5.825 GHz)—so band-notched UWB antennas are routinely required [5–7]. Complementary split-ring resonators (CSRRs) realise electrically small stop-bands [8–10]. Multiple-input multiple-output (MIMO) packing then raises mutual coupling, which is reduced by defected grounds, electromagnetic band-gap (EBG) lattices or metamaterial fences [11–14]. Machine-learning (ML) surrogates reduce the cost of the associated geometry search [15–18].

Three gaps remain at the intersection of those literatures, and they are computational rather than topological. First, dual-notch MIMO papers generally *compose* a rejection cell and a decoupling cell without stating the object being designed, so the specification survives only implicitly in a sequence of computer-aided design (CAD) steps. Second, inverse and reverse models are usually fitted at a single operating point [17,18], leaving open whether the learned map is a synthesiser or an interpolant of one layout. Third, notch control and isolation are applied sequentially without testing whether the response on the rejection band depends on the rejection geometry alone. If it does not, an inverse model trained on isolated cells is not a valid factor of the full-wave map, and the whole modular workflow rests on an untested assumption.

This work changes the object. A 5G coexistence *mask* $\mathcal{M}$ is inverted onto a compact MIMO aperture. Rejection and isolation act on electrically distinct regions—CSRRs in the patch and ground versus an SRR fence on the shared ground—and the relevant question is whether their effects on $\mathrm{supp}(N)$ are separable. Only when the sufficient conditions of Section 3.3 hold does the inverse factorise with a quantified error so that each factor can be learned independently. A dual-notch UWB MIMO prototype is retained as an experimental testbed aimed at one particular mask, C-band together with WLAN, rather than as the contribution itself, and is used to test whether those conditions and the target mask are actually satisfied on hardware. The C-band stop-band of the fabricated board covers only part of the 3.7 GHz to 4.2 GHz C-band mask. This shortfall is scored on the measured response (Section 6.3), and the changes that a retargeting must make are derived from it (Section 4.6). It is not claimed as a WiMAX notch. The X-band window 7.25 GHz to 8.4 GHz is deliberately kept inside the useful band as an operating window, not a stop-band.

The contributions are as follows:

1. Formulation of coexistence-mask inverse synthesis (CMIS): the design object is $\mathcal{M}$, the degrees of freedom are metamaterial operators on disjoint regions and separability is a residual to be measured rather than an identity assumed over the whole band. Sufficient conditions for a conditionally factorised inverse are derived, with a first-order bound on the inversion error and an explicit protocol, the four-configuration factorial, for testing them. The test is carried out in full wave at the fabricated design point, where condition (C1) is not met ($\|\Delta\|_{\mathcal{N}} = 0.61$ against $\varepsilon \leq 0.019$); accordingly, factorised inversion is presented strictly as a conditional theoretical formulation rather than an experimentally validated property.
2. A consistency-checked inverse map from a rejection specification onto CSRR geometry, trained on 5000 and 10,000 full-wave unit-cell samples, exercised on five targets per cell, equipped with an out-of-domain detector and an explicit treatment of non-uniqueness, and compared in cost with a direct optimiser. Dedicated full-wave and experimental validation is performed for the two fabricated cells, whereas the other eight targets are evaluated via nearest-neighbour database checks, so multi-target generalisation across the specification space is only partially validated.
3. Experimental frequency-domain characterisation of that aperture at 10 mm edge-to-edge spacing, with measured coupling below −20 dB from 2.1 GHz to 14.8 GHz (exceeded only near the upper band edge), a mask-compliance score computed on the measured curves that exposes the implications of partial compliance (44 % of M3 at −10 dB, leaving the lower C-band uncovered) for open-loop mask inversion and a tolerance budget for both rejection cells, while clearly identifying continuous-band/measured pattern-based ECC and time-domain pulse behaviour (group delay and fidelity factor) as open validations.
4. A first-order power-budget study, with an explicit SINR expression and a sensitivity analysis of its least-constrained input, showing that $N(f)$ and $I(f)$ solve two different coordinates of $\mathcal{M}$. This study is computed on modelled envelopes and is labelled as such throughout.

Section 2 positions the work against recent inverse-design and notched-MIMO papers. Section 3 states the mask, the operators, the separability hypothesis and its sufficient conditions. Section 4 details the inverse–forward pipeline. Section 5 reports the M3 hardware. Section 6 discusses the interaction residual, scores the mask compliance and illustrates the power budget. Section 7 compares the design with the literature. Section 8 discusses limitations, extensions to larger arrays and integrated sensing and communication (ISAC), and Section 9 concludes the paper.

2. Related Work and Positioning

Band-notched UWB antennas using CSRRs, slots or EBG cells are mature [5–7,10,19]. Compact UWB MIMO arrays with stubs, neutralisation lines, DGS or metamaterial fences are equally well documented [11–13,20–24]. The usual narrative is additive: etch a notch, then insert a decoupler. That narrative does not define an inverse problem, so it cannot support a computational claim of the form “the specification maps to the geometry”.

ML has entered the same literature along two axes. Forward surrogates predict *S*-parameters from geometry and replace the inner loop of a sweep [15,16,25,26]. Inverse and reverse models map a specification onto geometry [17,18]. Deep inverse designs embedding a forward-consistency loss [18] are closest in spirit to the present pipeline; they are demonstrated as single-layout reconstructions rather than as multi-target synthesisers tied to a coexistence mask and to a fabricated MIMO prototype, although in the present study dedicated full-wave verification is likewise completed only for the two fabricated target cells while the remaining eight queried targets are checked against nearest training samples.

The distinction pursued here is therefore methodological rather than topological. The fabricated antenna serves as an experimental testbed for a conditionally factorised inverse aimed at mask M3: the design object is $\mathcal{M}$, independent inversion of the operators is strictly conditional on satisfying the separability bounds (C1)–(C3) on $\mathrm{supp}(N)$ and a candidate geometry is accepted only if a companion forward surrogate reconstructs the query. A hardware comparison with recent notched MIMO designs is deferred to Section 7.

3. Coexistence-Mask Inverse Synthesis

3.1. Mask

The coexistence mask is the tuple

$$\mathcal{M} = \big(N(f),\ I(f),\ \rho_{\max},\ \eta_{\min},\ \mathcal{B}\big), \tag{1}$$

where $N(f) \in \{0,1\}$ marks frequencies that must be rejected ($|S_{11}|$ above -10 dB on the radiator), $I(f)$ is a minimum isolation (here, $I(f) = 20$ dB on the array band), $\rho_{\max}$ is an ECC ceiling (here, 0.5 as a diversity bound), $\eta_{\min}$ is a radiation-efficiency floor outside the stop-bands and $\mathcal{B}$ is the useful band. Three 5G-relevant masks are used throughout:

- M1: C-band downlink 3.7 GHz to 4.2 GHz;
- M2: WLAN 5.15 GHz to 5.35 GHz and 5.725 GHz to 5.825 GHz;
- M3: $N_{\mathrm{M1}} \cup N_{\mathrm{M2}}$, realised by the fabricated prototype.

The stop-bands quoted for the board in the original submission are 3.8 GHz to 4.2 GHz (400 MHz) and 4.9 GHz to 6.2 GHz (1.3 GHz). Reading the measured two-port S_{11} (Figure 1), these edges correspond to the points where $|S_{11}|$ crosses about -15 dB, not the -10 dB of the mask definition. Scored with the -10 dB criterion of (1), the realised rejection is narrower (Section 6.3): the fabricated board meets M3 only partially, and the interval 3.7 GHz to 3.8 GHz of M1 is left unrejected on the fabricated board (the full-wave model starts at 3.62 GHz). Section 4.6 derives what closing the gap requires.

The five coordinates of $\mathcal{M}$ do not play the same role in synthesis. $N(f)$ and $I(f)$ are *inputs* of the inverse: $N(f)$ is converted into the notch specification $\mathbf{x}$ of (12), and $I(f)$ drives the sizing of the fence. $\rho_{\max}$ and $\eta_{\min}$ are *acceptance constraints* checked a posteriori on the assembled aperture: a candidate that satisfies the forward-consistency test (8) but violates $\mathrm{ECC}(f) \le \rho_{\max}$ or $\eta(f) \ge \eta_{\min}$ on $\mathcal{B} \setminus \mathrm{supp}(N)$ is rejected, and the next feasible candidate of Section 4.3 is examined. This choice is deliberate: ECC and efficiency are aperture-level quantities that a unit-cell dataset cannot predict, so promoting them to inputs would require an aperture-level training set.

3.2. Operators, Separability and the Interaction Residual

Let $\theta = (\theta_N, \theta_I)$ collect the CSRR and SRR dimensions and let

$$F : (\theta_N, \theta_I) \longmapsto \big(S(f),\ \mathrm{ECC}(f),\ \eta(f)\big) \tag{2}$$

be the full-wave map of the aperture. Rejection and isolation are introduced by two geometry operators, E_N and E_I, which etch the CSRR cells and the SRR fence respectively on a common baseline layout Λ_0. Because etching acts on disjoint metallisation regions, E_N and E_I commute exactly: $E_N E_I \Lambda_0 = E_I E_N \Lambda_0$ is one and the same board. Operator ordering is therefore *not* the question, and an order-swap experiment would be vacuous.

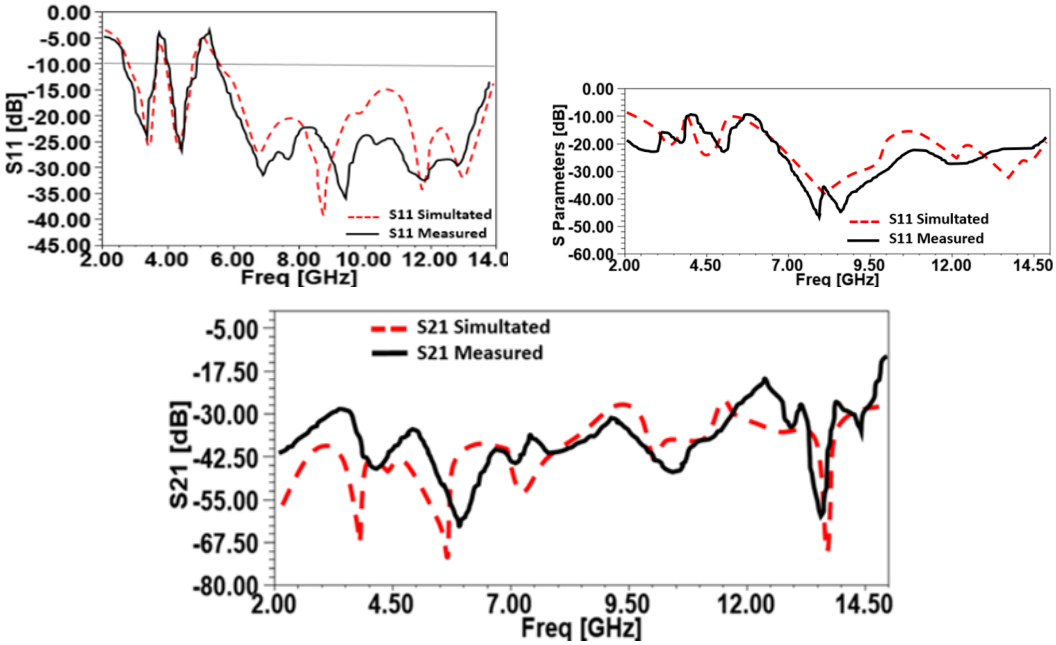


Figure 1. Matching and isolation (measured and full-wave simulated data). **Top left**: simulated and measured S_{11} of the dual-notched element. **Top right** and **bottom**: simulated and measured MIMO S_{11} and S_{21} with the SRR fence. Read on the curves, measured $|S_{21}|$ stays below -20 dB from 2.1 GHz to 14.8 GHz but only touches that level (-19.9 dB, within the reading error) near 12.4 GHz, and it exceeds it above about 14.8 GHz. The element-level notches are at 3.7 and 5.25 GHz.

The question that governs whether the inverse factorises is *separability*: does the response on the rejection band depend on θ_N alone? A first ingredient is that the mixed sensitivity vanishes,

$$\frac{\partial^2 F}{\partial\theta_N\,\partial\theta_I}\Big|_{\mathrm{supp}(N)} \approx 0, \tag{3}$$

whose finite-difference counterpart over the four corner layouts is the interaction residual

$$\Delta(f) = S^{N+I}(f) - S^{N}(f) - S^{I}(f) + S^{0}(f), \tag{4}$$

with 0 as the bare aperture, N the CSRR-only board, I the fence-only board and $N+I$ the complete board. Equation (4) is a 2×2 factorial contrast: it isolates the part of the response that cannot be attributed to either operator acting alone. Writing $\|\cdot\|_{\mathcal{N}}$ for a weighted L^∞ seminorm supported on $\mathrm{supp}(N)$, the additivity hypothesis is

$$\|\Delta\|_{\mathcal{N}} \le \varepsilon. \tag{5}$$

Condition (3), or its finite form (5), is *not* sufficient for an independent inversion. If the mixed derivative vanishes identically on a convex domain, the restriction of F to $\mathrm{supp}(N)$ is additive, $F|_{\mathcal{N}} = g(\theta_N) + h(\theta_I)$, but the rejection response still depends on θ_I through h. Solving for θ_N at a fence geometry different from the one used during training would then be biased by $h(\theta_I) - h(\theta_I^0)$. Two further conditions are needed. They are stated next.

3.3. Sufficient Conditions for a Factorised Inverse

Let $\Phi_N(\theta_N, \theta_I)$ denote the restriction of F to $\mathrm{supp}(N)$ (the reflection response on the rejection band), Φ_I the restriction to the isolation coordinate on $\mathcal{B}$ and (θ_N^0, θ_I^0) the baseline (no CSRR, no fence). The partial maps are $F_N(\theta_N) = \Phi_N(\theta_N, \theta_I^0)$ and $F_I(\theta_I) = \Phi_I(\theta_N^0, \theta_I)$. Assume, over the design domain $\Theta_N \times \Theta_I$, the following:

(C1) *Bounded interaction:* $\|\Delta\|_{\mathcal{N}} \le \varepsilon$ for every (θ_N, θ_I), i.e., (5) holds uniformly and not only at the fabricated point.

(C2) *Bounded cross-sensitivity:* the fence alone barely acts on the rejection band, $\|S^I - S^0\|_{\mathcal{N}} \le \delta_N$; symmetrically, the CSRRs alone barely act on the isolation coordinate, $\|S^N_{21} - S^0_{21}\|_{\mathcal{B}} \le \delta_I$.

(C3) *Well-conditioned factors:* the Jacobians $J_N = \partial F_N/\partial\theta_N$ and $J_I = \partial F_I/\partial\theta_I$ have smallest singular values bounded below by $\sigma_N > 0$ and $\sigma_I > 0$ on the domain.

Proposition. *Under (C1)–(C3), let $\theta^\star = (\theta^\star_N, \theta^\star_I)$ satisfy $\Phi_N(\theta^\star) = N$ and $\Phi_I(\theta^\star) = I$, and let $\hat\theta_N = F_N^{-1}(N)$, $\hat\theta_I = F_I^{-1}(I)$ be the factorised solutions. Then, to first order,*

$$\|\hat\theta_N - \theta^\star_N\| \le \frac{\varepsilon + \delta_N}{\sigma_N}, \qquad \|\hat\theta_I - \theta^\star_I\| \le \frac{\varepsilon_I + \delta_I}{\sigma_I}, \tag{6}$$

with ε_I the interaction residual (4) evaluated on the isolation coordinate.

Proof. By definition of (4), $\Phi_N(\theta_N, \theta_I) = F_N(\theta_N) + [\Phi_N(\theta^0_N, \theta_I) - \Phi_N(\theta^0_N, \theta^0_I)] + \Delta$. At $\theta^\star$, the bracket is bounded by δ_N (C2) and Δ by ε (C1); hence, $\|F_N(\theta^\star_N) - N\| \le \varepsilon + \delta_N$. Since $F_N(\hat\theta_N) = N$, the mean-value theorem and (C3) yield $\|\hat\theta_N - \theta^\star_N\| \le (\varepsilon + \delta_N)/\sigma_N$. The isolation factor follows using symmetry. □

When (6) is smaller than the fabrication tolerance, the factorisation

$$\theta^\star_N \approx F_N^{-1}(N), \qquad \theta^\star_I \approx F_I^{-1}(I) \tag{7}$$

is exact for practical purposes. Condition (C1) alone, which is what (3) expresses, controls only one of the two numerator terms. Only under conditions (C1)–(C3) can each factor be validly approximated and inverted by a regressor trained on the corresponding isolated structure without aperture-level recalibration; whenever (C1) or (C2) is violated, open-loop unit-cell inversion incurs an uncompensated transfer bias and cannot be interpreted as a validated factorisation. A candidate $\hat\theta = f_{\text{inv}}(\mathcal{M})$ is accepted only if a companion forward surrogate satisfies

$$\left\| f_{\text{fwd}}\left(f_{\text{inv}}(\mathcal{M})\right) - \mathcal{M} \right\|_W \le \tau, \tag{8}$$

with W placing a 120 MHz budget on the notch centre and a 5 dB budget on the notch depth (Section 4).

Experimental test. Conditions (C1) and (C2) are both evaluated with the four layouts 0, N, I and $N + I$ of (4): Δ yields ε, and the differences $S^I - S^0$ and $S^N_{21} - S^0_{21}$ yield δ_N and δ_I. The original submission evaluated only the main effect of the fence on the notched board,

$$\mu_I(f) = \left|S^{N+I}_{11}(f)\right| - \left|S^N_{11}(f)\right|, \tag{9}$$

in linear magnitude. Since $\mu_I = (|S^I| - |S^0|) + \Delta$ to first order, a small μ_I bounds the sum $\varepsilon + \delta_N$ at the fabricated point but cannot separate the two terms, and it says nothing about other values of θ_N. A conclusive test therefore requires full-wave solves of all four layouts at two or more values of θ_N, with (C3) evaluated by finite differences of the forward surrogate. That test is carried out at the fabricated design point in Section 6.2; a second value of θ_N remains open.

3.4. Choice of the Separability Threshold

The threshold ε is not a free tuning parameter. It is set by the two tolerances that already define the design. (i) *Magnitude:* a perturbation cannot be allowed to bring any rejected frequency back below the −10 dB criterion, so

$$\varepsilon_{\text{mag}} = \min_{f \in \text{supp}(N)} \left(|S^{N+I}_{11}(f)| - 10^{-10/20}\right), \tag{10}$$

i.e., the rejection margin of the realised board in linear magnitude. (ii) *Frequency:* the notch edges cannot move by more than the 120 MHz budget of (15), which is of the order of the shift produced by the FR4 permittivity tolerance (Section 5.4). Separability is declared when both criteria hold.

On the fabricated board, criterion (i) is already very tight. The measured two-port reflection peaks at −9.5 dB (Section 5), i.e., $|S_{11}| = 0.335$ against 0.316 at the criterion, so $\varepsilon_{\mathrm{mag}} \leq 0.019$ even over the board's own −10 dB intervals, and $\varepsilon_{\mathrm{mag}} < 0$ over the full mask. Any interaction, however small, can therefore move part of the stop-band below the criterion. The threshold is thus set by a rejection margin that the present board does not have, which is a second reason, beyond the M1 gap, why deeper notches are the first design change (Section 4.6). The full-wave response of the complete aperture does not provide a positive margin either: evaluated over the whole of M3, (10) yields approximately −0.16, so $\varepsilon_{\mathrm{mag}}$ is not defined over all of $\mathrm{supp}(N)$. The value $\varepsilon \leq 0.019$ obtained on the measured board is therefore kept as the reference threshold, and the verdict of Section 6.2 is provided at 0.5ε, ε and 2ε. Such a sensitivity check becomes fully meaningful only once $\varepsilon_{\mathrm{mag}}$ is positive.

3.5. Electrical Orthogonality

F_N is realised by CSRRs. The patch cell is excited by an electric field normal to the board; the ground cell sees the same orientation under the feed [8,9]. In the lumped picture of a CSRR-loaded line, the etched ring is a shunt tank (L_c, C_c) with

$$f_0 = \frac{1}{2\pi\sqrt{L_c C_c}}; \tag{11}$$

so f_0 is a nonlinear function of perimeter and slot width: a closed-form inversion of (11) is not accurate once the cell is loaded by a finite ground and a radiator, which is why an inverse surrogate is used. F_I is an SRR fence on the shared ground, operated above its magnetic resonance, where the loaded strip supports a negative-permeability window and opposes surface-wave coupling [14,27]. The two excitations—electric, on the radiator metallisation; magnetic, on the ground fence—are therefore distinct, which is the physical rationale for (C1) and (C2). However, this distinction in dominant excitation mode is only a qualitative physical motivation and does not guarantee that (C1) and (C2) hold quantitatively on a compact shared ground plane, where reactive near-field loading from the fence can still shift the CSRR resonances, as shown by the four-configuration test in Section 6.2. The two CSRR cells themselves sit in different layers, so F_N further splits into weakly interacting factors; that split is used as a first experimental hint, consistent with the dual-notch S_{11} overlay of Section 5.

4. Inverse Machine-Learning Map for the Rejection Operator

4.1. Learning Task

Conventional CSRR sizing is a nested sweep: n geometrical variables discretised over m values each cost m^n full-wave solves. The inverse task used here maps a specification onto geometry,

$$\mathbf{x} = [f_0,\ |S_{21}|_{\min},\ |S_{11}|_{\mathrm{res}}] \longmapsto \mathbf{y}, \tag{12}$$

with $\mathbf{y} = [d, s]$ for the circular cell (CSRR1) and $\mathbf{y} = [d', s', a']$ for the rectangular cell (CSRR2). Table 1 lists the variables and the dataset sizes. Each geometric coordinate is drawn independently and uniformly over the intervals of Table 1, whose upper bounds are set by the 4.4 mm cell footprint and whose lower bounds are set by the minimum etchable slot of the process. The neighbourhood of the fabricated cell is oversampled so that the retained design sits in a densely covered region. Every geometry reported in Tables 4 and 5 lies well inside this hull, and no extrapolation outside it is claimed.

Table 1. Variables of the inverse model, sampling intervals and size of the training sets. Inputs are the specification, and outputs are the geometry to be synthesised.

	Quantity	CSRR1	CSRR2
Inputs $\mathbf{x}$	notch frequency f_0 (GHz)	2.68–6.86	3.23–7.86
	notch depth $\lvert S_{21}\rvert_{\min}$ (dB)	−49.1 to −24.7	−36.8 to −19.7
	reflection at resonance $\lvert S_{11}\rvert_{\mathrm{res}}$ (dB)	−0.61 to −0.16	−0.77 to −0.20
Outputs $\mathbf{y}$	ring diameter d / side d' (mm)	2.60–4.90	1.70–3.40
	slot width s/s' (mm)	0.15–0.55	0.15–0.45
	gap a' (mm)	—	0.10–0.40
Number of full-wave samples		5000	10,000
Cell footprint $X \times Y$ (mm^2)		4.4 × 4.4	4.4 × 4.4

Geometric coordinates are sampled independently and uniformly over the stated intervals, with the neighbourhood of the fabricated cell oversampled. Input intervals are the resulting spans of the specification, not independently imposed bounds.

Datasets are generated by sweeping the microstrip-loaded unit cell in Ansys HFSS (finite-element frequency domain; FR4 with $\varepsilon_r = 4.4$, $\tan\delta = 0.02$; copper 0.035 mm) and recording S-parameters. Features and targets are standardised. Each dataset is split 80 %/20 % train/test; model selection uses five-fold cross-validation on the training subset only. Five regressors are compared: linear regression (LR), support-vector regression (SVR, RBF kernel), gradient boosting (GB), k-nearest neighbours (KNN) and random forest (RF). The searched grids are as follows: RF estimators $\in \{100, 200, 400\}$, max. depth $\in$ {None, 10, 20}; GB estimators $\in \{100, 300\}$, learning rate $\in \{0.05, 0.1\}$; KNN $k \in \{3, 5, 7, 9\}$ with uniform or distance weights; SVR $C \in \{1, 10, 100\}$, $\gamma \in$ {scale, 0.1}. LR is the closed-form baseline.

Performance is quantified by the coefficient of determination R^2 and the root-mean-square error:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left\|\hat{\mathbf{z}}_i - \mathbf{z}_i\right\|^2}, \tag{13}$$

where the squared error is averaged over the output coordinates of $\mathbf{y}$, as well as over the N samples. Since the two cells differ in the number of outputs and in the physical scale of those outputs, an absolute geometric RMSE is not comparable between them, and its numerical value depends on the sampling bounds of Table 1 rather than on the quality of the model alone. Tables 2 and 3 therefore report the dimensionless normalised error:

$$\mathrm{NRMSE} = \frac{\mathrm{RMSE}}{\sigma_{\mathbf{y}}}, \tag{14}$$

with $\sigma_{\mathbf{y}}$ being the standard deviation of the targets over the same split, so that a value of 1 corresponds to predicting the mean geometry and a value of 0 to exact recovery. Both metrics are evaluated on the held-out 20 % test split only.

A geometric error, normalised or not, is still not a filter specification: two cells can differ by a few micrometres and resonate hundreds of megahertz apart, or differ visibly and resonate identically. The regression scores are therefore used only to select candidate models. Acceptance is decided in the electromagnetic domain, by the forward-consistency budgets of (15) and by re-simulating the predicted cells.

4.2. Prediction Results and Forward Consistency

Table 2 shows that RF is retained for CSRR1 ($R^2 = 0.9825$, NRMSE= 0.132). The linear baseline collapses to $R^2 = 0.1525$, which confirms the strongly nonlinear map expected from (11) and rules out the possibility that the task is trivial. Table 3 shows that KNN is retained for CSRR2 ($R^2 = 0.9921$, NRMSE= 0.089), slightly ahead of RF ($R^2 = 0.9856$); the

denser 10,000-sample set favours local interpolation. Because KNN degrades outside the sampled hull, RF is kept as a consistency check, and both models must agree within fabrication tolerance before a geometry is accepted. The fabricated primary cells are $d = 3.8\,\text{mm}$, $s = 0.3\,\text{mm}$ (notch near 4 GHz, $|S_{21}| \approx -40\,\text{dB}$) and $d' = 2.4\,\text{mm}$, $s' = a' = 0.2\,\text{mm}$ (notch near 5.3 GHz, $|S_{21}| \approx -25\,\text{dB}$), on $X = Y = 4.4\,\text{mm}$ footprints. Figure 2 summarises the inverse–forward pipeline. Figures 3 and 4 compare full-wave and predicted S-parameters of those cells over the whole band, not at a single frequency.

Model-selection criteria beyond R^2. Four criteria guided the choice, in addition to the test scores. (i) *Robustness:* RF and KNN are the only two models with $R^2 > 0.96$ on the held-out split for both cells, and they lead the third model by at least 0.12 in R^2 (Tables 2 and 3). Their two predictions are also required to agree within the fabrication tolerance before a geometry is accepted, which protects against a model that happens to score well on one split. (ii) *Behaviour at the edge of the domain:* RF and KNN predictions are convex combinations of training targets, so they cannot return a geometry outside the sampled, fabricable intervals. SVR and GB can extrapolate to non-fabricable values, which is undesirable for an inverse map. (iii) *Prediction time:* on a laptop CPU (6th-generation Intel Core, single thread, scikit-learn), a single query takes about 10 ms for a 100-tree forest and 40–45 ms for a 400-tree forest (median; 95th percentile below 80 ms) and below 1 ms for KNN with $k = 9$. These timings were obtained with models of the same size and dimensions as ours trained on synthetic data, so they measure model cost and not data content. They are negligible compared with one full-wave solve. (iv) *Generalisation:* RF averages many decorrelated trees and is less sensitive to local sampling density, which suits the sparser 5000-sample CSRR1 set. KNN exploits the denser 10,000-sample CSRR2 set but degrades where the data are sparse, so it is never used without the RF cross-check and the out-of-domain detector of Section 4.4. Importantly, the held-out test metrics in Tables 2 and 3 measure statistical generalisation within the sampled probability distribution (which oversamples the neighbourhood of the fabricated cells) and do not, by themselves, establish electromagnetic synthesis generalisation across arbitrary new target specifications (Section 4.5).

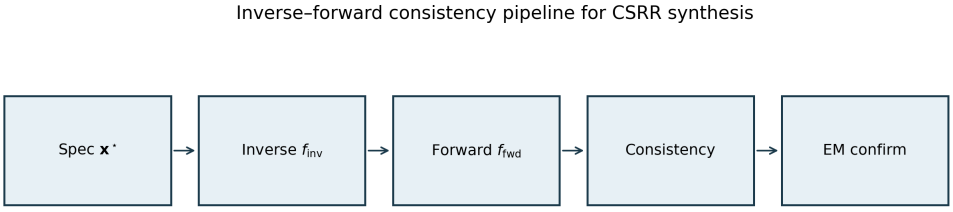


Figure 2. Inverse–forward consistency pipeline implementing F_N^{-1}. A candidate geometry is accepted only if the companion forward surrogate reconstructs the queried specification within the frequency and depth budgets of Section 4.

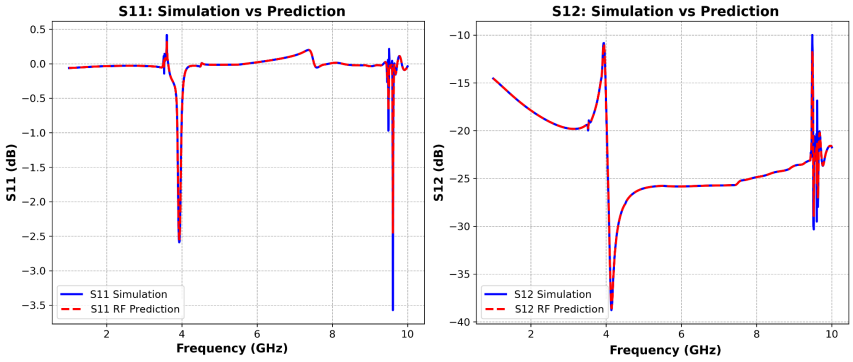


Figure 3. Full-wave versus predicted S-parameters of the retained CSRR1 cell (random forest).

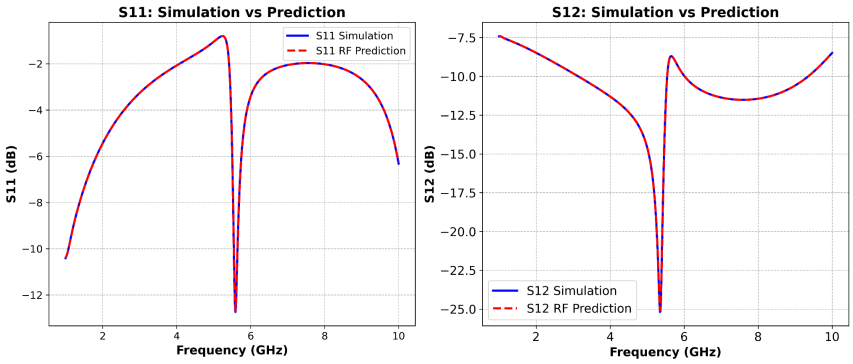


Figure 4. Full-wave versus predicted S-parameters of the retained CSRR2 cell (k-nearest neighbours, with random forest as consistency check).

Table 2. Regression performance for CSRR1 geometry prediction on the held-out test split. NRMSE is the target-normalised error of (14) and is dimensionless.

Algorithm	R^2	NRMSE
Gradient boosting	0.8435	0.396
SVR	0.6173	0.619
KNN	0.9634	0.191
LR	0.1525	0.921
Random forest	**0.9825**	**0.132**

Table 3. Regression performance for CSRR2 geometry prediction on the held-out test split. NRMSE is the target-normalised error of (14) and is dimensionless.

Algorithm	R^2	NRMSE
Gradient boosting	0.6835	0.563
SVR	0.4173	0.763
KNN	**0.9921**	**0.089**
LR	0.0128	0.994
Random forest	0.9856	0.120

An inverse prediction can be geometrically plausible yet electromagnetically inconsistent if several layouts map to nearby specifications. A companion forward surrogate $f_{\mathrm{fwd}} : \mathbf{y} \mapsto \mathbf{x}$ is therefore trained on the same samples. For a query $\mathbf{x}^{\star}$, the prediction $\hat{\mathbf{y}} = f_{\mathrm{inv}}(\mathbf{x}^{\star})$ is accepted only when both coordinates of the reconstruction error fall inside their budgets,

$$|\hat{f_0} - f_0^{\star}| \le 120\,\mathrm{MHz} \quad \text{and} \quad |\widehat{|S_{21}|}_{\min} - |S_{21}|^{\star}_{\min}| \le 5\,\mathrm{dB}, \tag{15}$$

with $(\hat{f_0}, \widehat{|S_{21}|}_{\min}) = f_{\mathrm{fwd}}(\hat{\mathbf{y}})$; otherwise, a local refinement is run before the electromagnetic confirmation (Figure 2). On a held-out CSRR1 test subset, the median frequency error after confirmation remains about 6 MHz, and the filter retains 99 % of the queries (Figure 5).

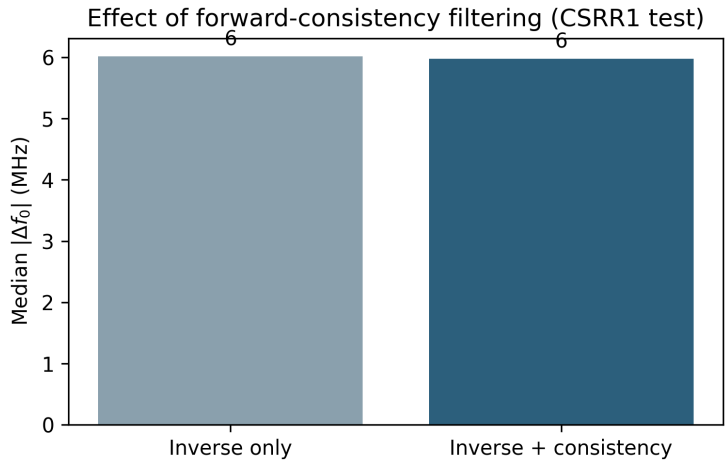


Figure 5. Median notch-frequency error on CSRR1 test queries with and without the forward-consistency filter.

4.3. Non-Uniqueness of the Inverse Problem

The map $\mathbf{y} \mapsto \mathbf{x}$ is many-to-one: using (11), an increase in ring perimeter (larger L_c) can be compensated by a wider slot (smaller C_c), so a continuum of geometries resonates at the same f_0. A high R^2 on the geometric coordinates therefore does not certify uniqueness or optimality. Worse, a regressor trained with a squared loss estimates the *conditional mean* of $\mathbf{y}$ given $\mathbf{x}$. When two disjoint branches of solutions exist, that mean can fall between them and correspond to no valid cell. The pipeline addresses this in three steps.

First, the object of the inversion is redefined as the *feasible set*

$$\Theta(\mathbf{x}^\star) = \left\{ \mathbf{y} \in \mathcal{Y} : \left\| f_{\mathrm{fwd}}(\mathbf{y}) - \mathbf{x}^\star \right\|_W \le \tau \right\}, \tag{16}$$

rather than a single point. Second, Θ is populated with three sources of candidates: the RF prediction, the KNN prediction and the k nearest training geometries in specification space, each screened by the forward surrogate. A conditional-mean prediction that falls between branches fails (15) and is discarded rather than fabricated. Third, one element of Θ is selected by an explicit secondary criterion: the largest fabrication margin, i.e., the candidate whose predicted f_0 moves least under a $\pm 50\,\mu\mathrm{m}$ perturbation of every dimension, computed with the forward surrogate; ties are broken by distance to the centre of the sampled hull. The size and spread of Θ indicate how strongly a specification constrains the geometry. The synthesised geometries of Tables 4 and 5 (Section 4.5) provide their order of magnitude. A linear fit of f_0 on (d, s) over the five CSRR1 geometries gives $\partial f_0/\partial d \approx -0.98\,\mathrm{GHz\,mm^{-1}}$ (rms residual 41 MHz), close to the estimate $-f_0/d \approx -1.05\,\mathrm{GHz\,mm^{-1}}$ that (11) gives for the ring perimeter, and $\partial f_0/\partial s \approx -3.1\,\mathrm{GHz\,mm^{-1}}$, which is less well determined with five points. A 120 MHz budget therefore admits a band of about ± 0.12 mm in d at fixed s, and the iso-f_0 direction trades about 0.3 mm of d for 0.1 mm of s. For CSRR2, s' and a' are almost collinear over the five geometries (correlation 0.97): only their combination is constrained by the specification, which is a direct sign of non-uniqueness. A systematic enumeration of Θ over the full datasets is left to future work. Multiple geometries satisfying the same specification are therefore not an inconvenience to hide but a resource. The secondary criterion turns them into robustness, and $\rho_{\max}$ and $\eta_{\min}$ (Section 3.1) can be used to choose among them at aperture level.

4.4. Domain of Validity and Out-of-Domain Detection

The inverse model is considered reliable inside the region of specification space covered by the training set, which is narrower than the per-coordinate spans of Table 1 because the

three inputs are correlated. A query $\mathbf{x}^\star$ is flagged as out-of-domain, and no geometry is returned without a full-wave search, if any of the following holds: (i) one coordinate lies outside the span of Table 1; (ii) the mean standardised distance from $\mathbf{x}^\star$ to its k nearest training specifications exceeds the 95th percentile of the same statistic computed on the training set (a non-parametric density test); (iii) RF and KNN predictions differ by more than the 50 µm fabrication tolerance. Near the edge of the domain, RF and KNN predictions are biased towards the interior because they average neighbouring training targets. This bias is expected to grow with distance to the boundary, and it is caught by the budget (15) before any geometry is fabricated. The three tests use only the training set and the two trained models, so they add no full-wave cost; the companion script implements them (Data Availability Statement).

4.5. Multi-Target Synthesis and Design Cost

To test whether the inverse model is a synthesiser rather than a post-fit of one layout, five rejection targets are issued for each cell across the C-band and WLAN windows (Tables 4 and 5). The resulting notch-frequency errors for these synthesis targets with respect to the nearest database samples are illustrated in Figure 6. The model returns a distinct geometry for each target, and the geometries move in the direction that (11) predicts, so the map is not degenerate.

The two kinds of verification in the tables must be distinguished. The fabricated primary cells (T2*) are verified by a dedicated full-wave solve at the predicted geometry, and then by measurement of the complete antenna. The other eight targets are compared with the nearest training sample (column "f_0 DB"). Database lookup detects a grossly wrong prediction, but not an error shared by the database, and it says nothing about points between grid samples. It is therefore a consistency check, not a measure of synthesis accuracy, and the eight non-fabricated rows should be read in that sense. Their specifications lie inside the spans of Table 1 and their geometries inside the sampled intervals. Independent full-wave solves of these eight geometries are the next verification step, and the predicted geometries are released for that purpose.

Implications for the generalisation capability of the inverse synthesis method. This distinction directly bounds what can be claimed regarding the generalisation capability of the proposed inverse synthesiser. Two levels of generalisation are involved: (i) *within-distribution statistical generalisation*, evaluated on the held-out 20 % test split ($R^2 = 0.9825$ and 0.9921), which confirms that the regressors interpolate accurately across the sampled unit-cell distribution; (ii) *independent multi-target electromagnetic generalisation*, which requires verifying via dedicated full-wave solves that newly synthesised geometries $\hat{\mathbf{y}} = f_{\text{inv}}(\mathbf{x}^\star)$ actually resonate at the requested target specifications $\mathbf{x}^\star$ across the C-band and WLAN domains. Because the training database oversamples the neighbourhood of the fabricated cells (T2*) and because the eight non-fabricated configurations (T1, T3, T4 and T5 for both CSRR1 and CSRR2) are presently checked only against their closest training-set samples rather than simulated in full wave, inter-sample interpolation errors or local surrogate bias away from T2* cannot be ruled out. Consequently, the multi-target generalisation capability of the inverse synthesis method remains only partially validated—specifically at the two T2* unit cells—until independent full-wave simulations of the remaining eight synthesised geometries are performed.

The design-cost argument needs the same care. Once the datasets exist, a new specification costs one confirmation solve instead of a factorial sweep (Figure 7), but the datasets themselves cost 5000 and 10,000 solves. A direct optimiser does not pay that entrance fee. The break-even point follows from a simple count. If a direct optimiser, for instance Bayesian optimisation with the full-wave solver in the loop, needs n_{opt} solves per specification, the

5000-sample CSRR1 database pays for itself after $5000/(n_{\mathrm{opt}} - 1)$ distinct specifications. Assuming n_{opt} between 20 and 60 for a two-variable cell, the break-even lies between about 85 and 260 specifications. For a single specification, or for a few, a direct optimiser is cheaper, and we state this explicitly. The case for the learning stage is consequently not total solve count but three properties an optimiser run does not leave behind: the marginal latency of a retarget drops from hours to seconds, which matters inside an interactive design loop; the same samples yield the forward surrogate that powers the consistency filter of (8); the database is reusable across masks, so the amortisation is a property of the laboratory rather than of one paper. Figure 7 should be read as marginal cost, which is what it plots, and not as total cost.

Table 4. Multi-target inverse synthesis for CSRR1 (circular cell). T2* is the fabricated C-band cell. "f_0 DB" is the nearest training sample (database check); for T2*, it coincides with the dedicated full-wave solve.

ID	f_0 Req. (GHz)	Depth (dB)	d/s (mm)	f_0 DB (GHz)
T1	3.80	−35	4.24/0.23	3.80
T2*	4.00	−40	3.80/0.30	4.00
T3	4.20	−30	4.07/0.17	4.19
T4	3.90	−25	4.41/0.17	3.90
T5	4.10	−45	3.57/0.37	4.11

Median $|f_0$ DB $- f_0$ req.$| = 4$ MHz. This measures consistency with the database, not synthesis accuracy.

Table 5. Multi-target inverse synthesis for CSRR2 (rectangular cell). T2* is the fabricated WLAN cell. Columns as in Table 4.

ID	f_0 Req. (GHz)	Depth (dB)	$d'/s'/a'$ (mm)	f_0 DB (GHz)
T1	5.00	−30	2.39/0.29/0.29	4.93
T2*	5.30	−25	2.40/0.20/0.20	5.30
T3	5.80	−30	2.08/0.28/0.29	5.81
T4	6.00	−20	2.19/0.18/0.11	5.97
T5	5.50	−35	2.05/0.38/0.38	5.48

Median $|f_0$ DB $- f_0$ req.$| = 20$ MHz. This measures consistency with the database, not synthesis accuracy.

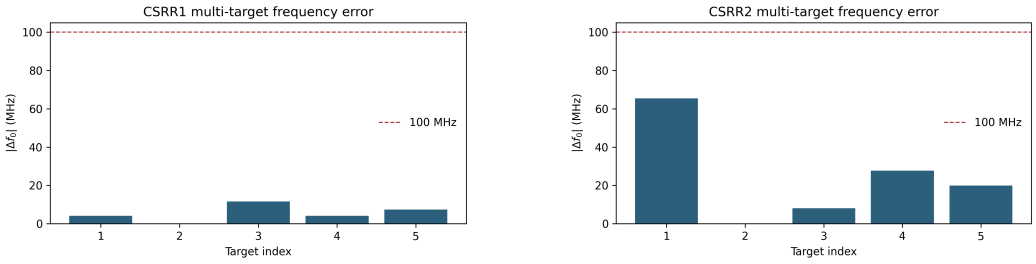


Figure 6. Notch-frequency error for the five CSRR1 (**left**) and CSRR2 (**right**) synthesis targets of Tables 4 and 5. Errors are computed with respect to the nearest training sample (database check).

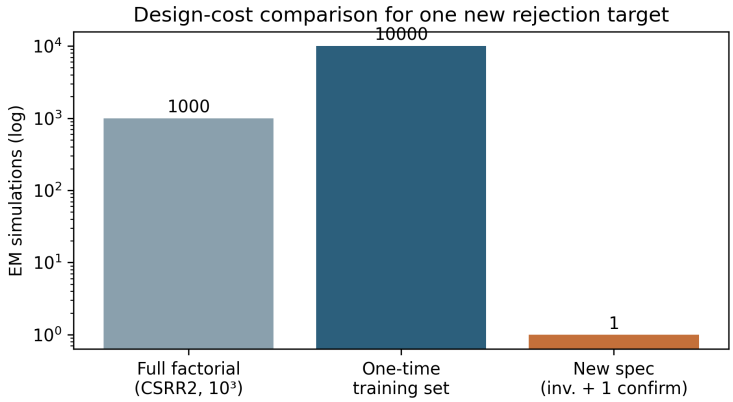


Figure 7. Marginal electromagnetic cost of one new rejection specification after the unit-cell datasets exist.

4.6. Retargeting a Complete Aperture: What Closing the M1 Gap Requires

The fabricated aperture rejects M1 over only about 3.91 GHz to 4.15 GHz at the −10 dB criterion, i.e., half of the interval, and leaves 3.7 GHz to 3.8 GHz unrejected in any reading (Table 10). This shortfall shows what a retarget of a complete aperture has to handle, including the transfer from the isolated cell to the integrated antenna. Two features of the problem matter. First, M1 requires a −10 dB rejection band of at least 500 MHz (12.7 % fractional), against about 240 MHz read on the measured aperture, so both the centre and the width must change. The rejection bandwidth is not an input of (12). It can be added as a fourth input, BW_{-10}, extracted from the S-parameters already stored for every training sample, so no new training solve is needed. Second, the notch of the loaded antenna does not coincide with the unit-cell resonance: the radiator and the finite ground shift it by a transfer offset $\delta_{\mathrm{tr}} = f_0^{\mathrm{ant}} - f_0^{\mathrm{cell}}$. For the fabricated cell ($f_0^{\mathrm{cell}} = 4.00$ GHz from Table 4), the measured notch peak is at 3.74 GHz on the isolated element ($\delta_{\mathrm{tr}} \approx -0.26$ GHz) and at 4.00 GHz on the two-port board ($\delta_{\mathrm{tr}} \approx 0$), with a reading uncertainty of about 0.03 GHz. For the rectangular cell ($f_0^{\mathrm{cell}} = 5.30$ GHz), the corresponding values are −0.05 and +0.40 GHz. The offset is therefore not a property of the cell: it depends on the structure around it, and it has to be calibrated on the board that is retargeted. The same argument applies to the width.

With the offset of the two-port board, the corresponding query is $\mathbf{x}^\star = [f_0 = 3.95\,\mathrm{GHz} - \delta_{\mathrm{tr}},\ |S_{21}|_{\min} \le -35\,\mathrm{dB},\ \mathrm{BW}_{-10} \ge 560\,\mathrm{MHz}]$, where 60 MHz of margin is kept on each side of M1. Answering it costs one inverse query. Validating the answer costs a full-wave solve of the complete two-port aperture with the re-synthesised cell, with the CSRR2 cell and the fence unchanged, and then a new fabrication. Neither step is part of this revision, and we make no claim that a retargeted aperture meets M1. The data already indicate where the difficulty lies. The required width is more than twice the measured one, and the measured notch of the two-port board rises above −10 dB by only 0.5 dB at its peak (Section 6.3), so the rejection must also become deeper. If no CSRR1 geometry in the sampled family provides that width, the feasible set (16) is empty. The model then signals that the mask calls for a different cell, such as a second resonator or a wider-band topology, instead of returning a geometry that cannot meet it. This is the expected behaviour of a mask-driven synthesis, and the retargeted design will also provide the second design point at which (C1) can be tested (Section 6.2).

5. Hardware Realisation of Mask M3

5.1. Baseline Element and Dual CSRR Rejection

The radiator is a circular patch on FR4 ($h = 1.6$ mm) of size 28 mm × 32 mm, fed by a 50 Ω microstrip line (Table 6). Bandwidth is enlarged by a partial ground, a V-slot

and corner chamfers until $S_{11} \leq -10$ dB from 3.15 GHz to 16 GHz. Simulated radiation efficiency exceeds 82 % over the lower and central band and falls at the upper edge because of the FR4 loss tangent, which is a substrate limit rather than a topology claim. A circular CSRR in the patch and a rectangular CSRR in the ground realise N_{M3} without enlarging the footprint (Figure 8a–c). Because the two cells occupy different metallisation layers, each notch can be moved with little shift of the other, which provides a qualitative indication of weak inter-cell coupling on the isolated element, without constituting an experimental validation of operator factorisation on the complete two-port aperture.

Table 6. Optimised dimensions of the baseline UWB element.

Parameter	Value (mm)	Parameter	Value (mm)
W	28	L_f	3.0
L	32	W_f	3.0
h	1.6	a	3.0
L_g	11.8	v	1.5
R	9.8	r	4.0

A dual-notched prototype was measured with an Agilent N5230A analyser after a two-port calibration. Simulated and measured S_{11} agree within the first notch of the isolated element (peaks at 3.77 and 3.74 GHz); the second notch of the element is shifted by +0.13 GHz, and by +0.09 and +0.4 GHz for the first and second notch of the two-port board. The shift of the element is ascribed to FR4 permittivity tolerance, etching and SMA parasitics (Figure 1; see Section 5.4, where the larger shift of the board is discussed). Matching with $S_{11} < -10$ dB is maintained over approximately 3 GHz to 15 GHz except in the stop-bands. The stop-bands are read here on the digitised measured curves of Figure 1, with a resolution of about 0.03 GHz and 0.25 dB. For the isolated element, $|S_{11}|$ exceeds −10 dB over 3.64 GHz to 3.97 GHz and 4.89 GHz to 5.48 GHz, with peaks of −4.3 and −3.9 dB at 3.74 and 5.25 GHz. For the two-port board, the notches are shallower and displaced: $|S_{11}|$ exceeds −10 dB only over 3.91 GHz to 4.15 GHz and 5.66 GHz to 5.91 GHz, with a peak of −9.5 dB at 4.0 and 5.7 GHz, whereas the full-wave aperture response peaks at −2.9 dB (3.88 GHz) and −2.0 dB (5.30 GHz) (Section 6.3). The measured reflection of the aperture is therefore a weak indicator of rejection: a notch that returns only about 11 % of the incident power at its deepest point cannot by itself be claimed to protect the receiver. The notch centres move by +0.26 and +0.45 GHz between the element and the aperture, which is the transfer offset used in Section 4.6. Outside the notches, the reflection coefficient of the element overlays the unloaded antenna, so the CSRRs behave as localised rejection elements rather than as parasitic radiators.

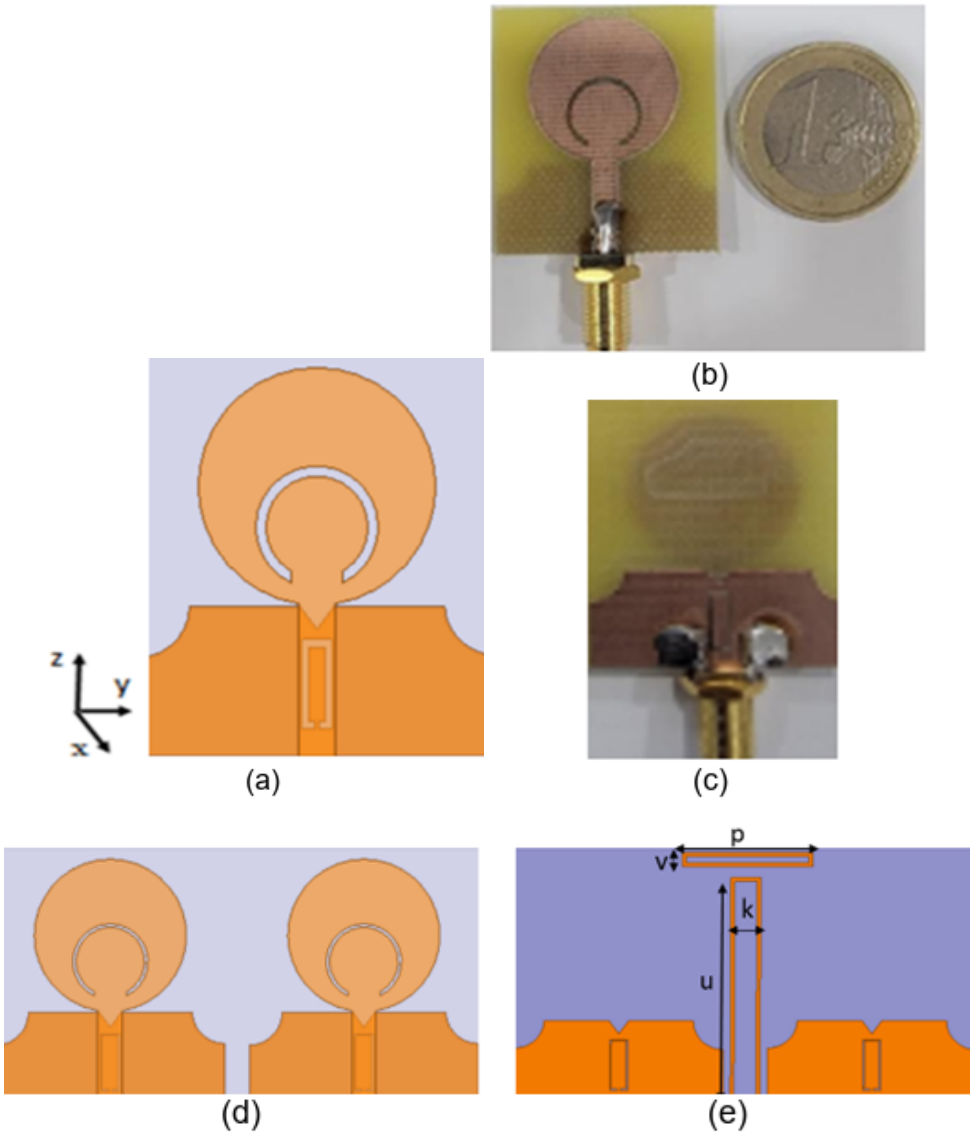


Figure 8. Hardware realisation of mask M3: (**a**) layout geometry of the single dual-CSRR element; (**b**) front view of the fabricated prototype with a 1 Euro coin for scale; (**c**) back view of the fabricated prototype showing the partial ground plane and rectangular CSRR; (**d**) front layout of the two-element MIMO array; (**e**) back layout of the two-element MIMO array showing the SRR isolation fence on the shared ground plane.

5.2. SRR Isolation Fence

Two dual-notched elements share a 60 mm board with 10 mm edge-to-edge spacing, without polarisation diversity, so decoupling is attributable to the fence (Figure 8d,e). In electrical terms, spacing is $0.07\lambda_0$ at the 2.1 GHz lower edge, $0.097\lambda_0$ at 2.9 GHz, the lower edge of the band over which the full-wave array is matched, about $0.17\lambda_0$ at 5 GHz and $0.33\lambda_0$ at 10 GHz. Whenever the compactness of the array is invoked in what follows, the 2.1 GHz value is quoted because that is where decoupling is hardest; it is the most favourable reading of a fixed physical gap and should be understood as such, and the 2.9 GHz value is the one supported by the full-wave matching. The MIMO board is electrically larger than the isolated element: the extended ground lowers the first resonance, and the matched array band starts near 2.1 GHz. All coupling values below refer to that array band. An SRR pair on the ground side ($r_{out} = 7$ mm, $r_{in} = 5$ mm, $p = v = 1$ mm, $k = u = 3$ mm) has a mean contour of about 37.7 mm and a half-wave magnetic resonance near 2.2 GHz, more than an octave below both CSRR stop-bands. The fence is therefore non-resonant

on supp(N), which is the geometric counterpart of (C2) in the isolated-resonator picture, even though finite-ground reactive loading in the complete two-port layout still produces a cross-sensitivity $\delta_N = 0.19$ and shifts the CSRR notch centres (Section 6.2).

Figure 1 also reports isolation. Without SRRs, S_{21} reaches approximately −11 dB at 2.5 GHz. With the fence, simulated coupling nulls reach about −70 dB at isolated frequencies (near 5.7 and 13.7 GHz); those nulls are narrow, process-sensitive and are *not* the figure of merit. The engineering metric is the band-wide upper bound. In full wave, the coupling is below −20 dB over 67 % of 2.1 GHz to 15 GHz (71 % above 2.9 GHz), with −13.6 dB at 2.5 GHz and a maximum of about −16 dB near 9 GHz (Figure 9). Below about 2.9 GHz, the simulated $|S_{11}|$ is above −10 dB, so the simulated array is not matched there. In measurement, it stays below −20 dB over about 98 % of the band, from 2.1 GHz to 14.8 GHz, with no usable margin near 12.4 GHz (-19.9 ± 0.4 dB), and rises to about −13 dB at 14.9 GHz. The isolation coordinate $I(f) = 20$ dB of the mask is therefore met over approximately 98 % of the array band (97 % to 98 % for a ±0.5 dB reading uncertainty), not over all of it (Table 10). The upper-edge excursion coincides with the frequency range where the electrical spacing exceeds $0.45\lambda_0$, and the fence is far above its design range. Measured SMA and cable parasitics are largest there too, but this is a likely cause and was not isolated experimentally. The measured isolation is better than the full-wave one at low frequency; both are reported, and the difference is not explained here. Surface currents with port 1 excited and port 2 terminated confirm that the ring pair traps the ground-plane current that would otherwise reach the second radiator (Figure 10).

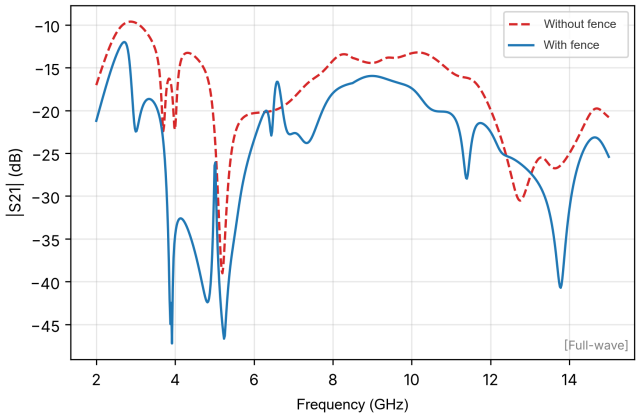


Figure 9. Full-wave. Isolation operator: $|S_{21}|$ of the two-port aperture without and with the SRR fence, to be compared with the mask $I(f) = 20$ dB. The measured isolation, which exceeds −20 dB above about 14.8 GHz, is shown in Figure 1.

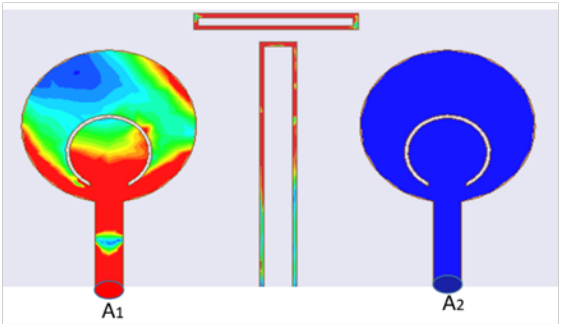


Figure 10. Surface-current distribution on the MIMO board with port 1 excited and port 2 terminated, illustrating the isolation operator.

5.3. ECC, Diversity Gain and Wideband Time-Domain Limitations

Two estimates of the envelope correlation coefficient must be distinguished, because they do not rest on the same assumptions [28]. The reference definition uses the embedded complex far-field patterns $\mathbf{F}_i(\theta, \phi)$ of port i, obtained with the other port terminated in 50 Ω:

$$\mathrm{ECC}_{\mathrm{FF}} = \frac{\left| \oint_{4\pi} \mathbf{F}_1(\theta,\phi) \cdot \mathbf{F}_2^*(\theta,\phi)\, d\Omega \right|^2}{\oint_{4\pi} |\mathbf{F}_1(\theta,\phi)|^2 d\Omega \oint_{4\pi} |\mathbf{F}_2(\theta,\phi)|^2 d\Omega}, \tag{17}$$

with $\mathbf{F}_i \cdot \mathbf{F}_j^* = E_{\theta,i}E_{\theta,j}^* + E_{\phi,i}E_{\phi,j}^*$ (isotropic, uniform multipath, equal polarisation power). The second estimate is the familiar *S*-parameter formula [29,30],

$$\mathrm{ECC}_S = \frac{\left| S_{11}^* S_{12} + S_{21}^* S_{22} \right|^2}{(1 - |S_{11}|^2 - |S_{21}|^2)(1 - |S_{22}|^2 - |S_{12}|^2)}, \tag{18}$$

and

$$\mathrm{DG} = 10\sqrt{1 - |\mathrm{ECC}|^2}. \tag{19}$$

Equation (18) equals (17) only for lossless antennas [28,29]. FR4 with $\tan\delta = 0.02$ dissipates a significant share of the accepted power, and Honma and Murata show that the *S*-parameter estimate already errs by more than 10 % once the radiation efficiency falls below 97 %. At low efficiency, it can carry essentially no information about the true correlation [28]. Contrary to the wording of the original submission, the error has no guaranteed sign, so (18) provides neither a lower bound nor a ceiling. It is reported here only for comparability with the literature, and it is not used as evidence of low correlation.

Computed with (18) on the full-wave *S*-parameters of the complete aperture, the ECC stays below about 0.005 above 2.9 GHz and peaks at about 0.25 near 2.7 GHz, where the simulated array is not matched (Figure 11). The pattern-based ECC (17) was evaluated from the full-wave embedded patterns of the complete aperture ($N + I$), with one port excited and the other terminated in 50 Ω. The patterns were sampled every 10° in θ (0°–180°) and ϕ (0°–350°) at 3, 3.9, 5.3, 8, 10 and 14 GHz and integrated with a $\sin\theta$ weight. Table 7 presents the following result: ECC ≤ 0.015 at all six frequencies, well below the 0.5 bound. The *S*-parameter estimate (18) is smaller than (17) by up to about three orders of magnitude inside the notches (3.9 and 5.3 GHz), which confirms that it is not a bound for this lossy FR4 board. The pattern-based values are therefore the ones to be used as correlation evidence, at the six sampled frequencies only. The in-notch values should be read with care, since the element is mismatched there ($|S_{11}|$ about −2 to −3 dB at the peaks). Two limitations of this ECC validation must be emphasised. First, evaluating (17) at only six discrete frequencies provides spot-frequency checks rather than a continuous broadband verification: because the aperture combines multiple high-Q resonant structures (two CSRRs per element and the SRR ground fence), rapid frequency variations in surface currents and far-field phase/polarisation can occur near the stop-band transitions and higher-order modes between the sampled frequencies. Second, no validation using 3D measured radiation patterns in an anechoic chamber is available, so the effects of fabrication tolerances, SMA connectors and feed-cable scattering on the far-field correlation remain unverified. The two ports are, in addition, coupled by less than −20 dB over about 98 % of the band (Section 5).

Table 7. Full-wave. Envelope correlation coefficient of the complete aperture ($N + I$): pattern-based value (17) from the embedded patterns (10° sampling in θ and ϕ) and S-parameter estimate (18) at the same frequencies.

f (GHz)	ECC (17), Patterns	ECC (18), S-Parameters
3.0	0.0060	1.2×10^{-4}
3.9	0.0066	2×10^{-5}
5.3	0.0152	1×10^{-5}
8.0	2×10^{-5}	4×10^{-5}
10.0	0.0097	3.9×10^{-3}
14.0	0.0021	$<10^{-5}$

Equation (19) then places the diversity gain close to 10 dB. That figure carries no information the ECC does not already contain—in the low-correlation regime, (19) is a saturating function and any ECC < 0.1 maps to essentially 10 dB—and it is quoted only for comparability with the tables of the literature. Complementary two-port metrics are bounded by the same S-matrix: where isolation is better than 20 dB, TARC follows the element matching outside the notches and capacity loss stays well under the 0.4 bit s^{-1} Hz^{-1} value usually taken as acceptable [12,30]. Radiation patterns of the dual-notched element remain essentially omnidirectional, with a peak gain near 5 dBi.

Time-domain behaviour and framing of wideband performance. Given the UWB nature of the front-end and the presence of three distinct resonant structures on the board (CSRR1, CSRR2 and the SRR fence), continuous-wave frequency-domain matching ($|S_{11}| < -10$ dB) and port decoupling ($|S_{21}| < -20$ dB) are necessary but not sufficient to establish distortionless wideband operation. In any band-notched UWB radiator, sharp magnitude rejection in the frequency domain is linked by causality to rapid phase transitions and group-delay peaks $\tau_g(f) = -\frac{1}{2\pi}\frac{\mathrm{d}\angle S_{21}(f)}{\mathrm{d}f}$ around the notch edges, as well as late-time ringing from energy stored in the resonant slots [10]. Neither the face-to-face group delay nor the time-domain pulse fidelity factor (FF) has been evaluated in simulation or measurement in the present study. Consequently, it remains to be verified that the phase dispersion induced by the CSRRs and the SRR fence remains localised within supp(N) and does not degrade pulse waveform integrity across the adjacent passband sub-bands (3.1 GHz to 3.7 GHz, 4.2 GHz to 5.15 GHz and 5.825 GHz to 10.6 GHz). All conclusions regarding the wideband performance of the prototype in this manuscript are therefore framed strictly in terms of frequency-domain S-parameter bandwidth and spot-frequency simulated pattern correlation, while continuous-band measured ECC and time-domain characterisation (group delay and fidelity factor) are explicitly identified as areas requiring further validation.

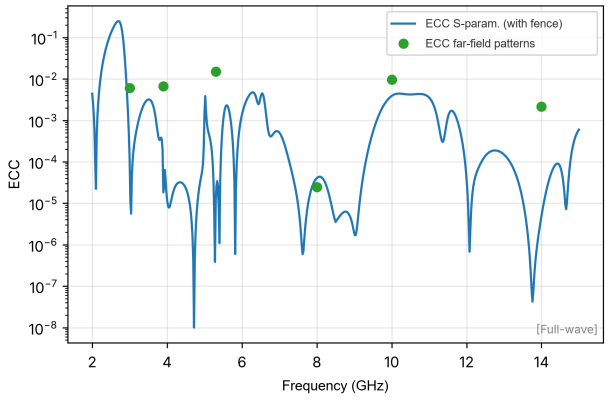


Figure 11. Full-wave. Envelope correlation coefficient: S-parameter estimate (18) with the SRR fence and pattern-based value (17) at six frequencies (Table 7).

5.4. Sensitivity to Fabrication and Substrate Tolerances

A first-order estimate sets the scale. For a slot etched in a ground plane, the CSRR capacitance scales with the effective permittivity $\varepsilon_{\text{eff}} \approx (\varepsilon_r + 1)/2$, so by (11),

$$\frac{\Delta f_0}{f_0} \approx -\frac{1}{2}\frac{\Delta \varepsilon_r}{\varepsilon_r + 1}. \tag{20}$$

For the usual FR4 tolerance $\varepsilon_r = 4.4 \pm 0.2$ this gives $\Delta f_0/f_0 \approx \mp 1.85\,\%$, i.e., about ± 74 MHz at 4 GHz and ± 98 MHz at 5.3 GHz. Both values lie inside the 120 MHz budget of (15). They explain the shift of the second notch of the isolated element between simulation and measurement (+0.13 GHz, peak at 5.12 and 5.25 GHz) only in part, and they do not explain the shift of the second notch of the two-port board (+0.4 GHz, peak at 5.3 and 5.7 GHz), which is four times larger than the permittivity estimate. Etching error, the alignment of the two CSRR layers and the SMA and cable parasitics must therefore be examined. The etching contribution can be estimated from the synthesised geometries of Tables 4 and 5, through the linear sensitivities of Section 4.3. For CSRR1, $\partial f_0/\partial d \approx -0.98\,\text{GHz mm}^{-1}$ and $\partial f_0/\partial s \approx -3.1\,\text{GHz mm}^{-1}$. For CSRR2, fitting on (d', s') yields $\partial f_0/\partial d' \approx -2.7\,\text{GHz mm}^{-1}$ and $\partial f_0/\partial s' \approx -4.4\,\text{GHz mm}^{-1}$ (rms residual 9 MHz). Each fit uses five points, and the slot coefficients are confounded with the other dimensions, so the resulting numbers are orders of magnitude rather than predictions. Table 8 combines them with (20).

Three conclusions follow. First, the 120 MHz budget of (15) is consistent with the fabrication process only if the etching error is held to about ± 20 µm for CSRR1, and even then it is slightly exceeded for CSRR2. The budget is therefore a synthesis tolerance, not a guaranteed fabrication tolerance, and a process with ± 50 µm etching would need a budget of about 180 MHz and 280 MHz for the two cells. Second, the CSRR2 notch is the more sensitive one, because its cell is smaller and its slots narrower. This matches the observation that the second notch moves more than the first between simulation and measurement. Third, the +0.4 GHz shift of the second notch of the board lies at the edge of the worst-case linear sum of etching and permittivity ($356 + 98 \approx 450$ MHz). It can be explained by fabrication alone only if all errors act in the same direction, which is unlikely. Layer alignment and connector parasitics therefore remain plausible contributors, and the shift is reported as an open fabrication question. A full-wave Monte Carlo analysis of the complete aperture would replace this first-order budget, and it is listed among the open validations in Section 8.

Table 8. First-order tolerance budget of the two notch frequencies. Etching: independent error on every CSRR dimension, combined as root sum of squares (RSS) and as worst-case linear sum. Permittivity: $\varepsilon_r = 4.4 \pm 0.2$ from (20).

Contribution to $\lvert\Delta f_0\rvert$ (MHz)	CSRR1 (4.0 GHz)	CSRR2 (5.3 GHz)
Permittivity ± 0.2	74	98
Etching $\pm$ 50 µm, RSS/worst case	161/203	258/356
Etching $\pm$ 20 µm, RSS/worst case	64/81	103/142
Total RSS, $\pm$ 50 µm etching	177	276
Total RSS, $\pm$ 20 µm etching	98	142
Budget of (15)	120	120

6. Interaction Residual, Mask Compliance and a Power-Budget Illustration

6.1. Provenance: Measured, Full-Wave and Modelled Data

This section combines three kinds of data, and every figure and table states which one it uses. **[Measured]**: VNA *S*-parameters of the fabricated board, digitised from Figure 1.

[Full-wave]: HFSS solves exported as Touchstone files, namely the element, and the four configurations of the two-port aperture (bare, CSRR-only, fence-only and complete) solved with identical settings (651 points, 2 GHz to 15 GHz, 20 MHz step). **[Modelled]**: *analytic S-parameter envelopes* that interpolate the measured and simulated corner values of the M3 campaign (notch edges, isolation bounds of about −11 dB without the fence versus −20 dB with it, ECC from (18)). The masks on the aperture response (Figure 12), the interaction residual (Figure 13), the isolation (Figure 9) and the ECC (Figure 11) are full-wave results. The mask compliance (Section 6.3) is scored on the measured curves, **digitised** from Figure 1 and on the full-wave aperture. The modelled envelopes are used only for the power-budget illustration (Section 6.4, Figure 14 and Table 11), where the four operator combinations A–D are needed on a common reference and with an assumed in-notch efficiency, because the bare and fence-only boards were solved in full wave (Section 6.2) but not fabricated. The modelled envelopes are not measurements, not full-wave outputs, and not a substitute for the Touchstone files of the same campaign. Two consequences must be kept in view when reading the modelled figures. First, the envelopes reproduce the corner values by construction, so agreement with those corners is not evidence. Second, the figures of merit derived from them should be read as *relative* statements about the roles of N and I, never as over-the-air throughput or as a certified compliance report. The generating script, the reconstructed traces and the sensitivity study of Section 6.4 are released so that a reader can substitute measured Touchstone files and recompute every number.

6.2. Four-Configuration Interaction Residual

The residual (4) was evaluated from four full-wave solutions of the same two-port aperture with identical settings (2 GHz to 15 GHz, 20 MHz step): the bare aperture (0), the CSRR-only aperture (N), the fence-only aperture (I) and the complete aperture ($N + I$); only $N + I$ was fabricated. Figure 12 shows the four reflection responses against the masks. Table 9 presents the quantities of Section 3.3 (complex S_{11} on $\mathrm{supp}(N)$ = M3, S_{21} on $\mathcal{B}$) and Figure 13 shows $|\Delta(f)|$.

Condition (C1) is not satisfied in the fabricated design. The residual is concentrated in the two notches: it peaks at 0.61 at 3.96 GHz and 0.48 at 5.20 GHz. It is dominated by a shift of the notch peaks when the fence is added (4.00 → 3.88 GHz and 5.16 → 5.30 GHz), which is of the order of the 120 MHz budget; a narrow notch that moves by this amount produces a large complex difference even if its depth is preserved. Condition (C2) holds only loosely: $\delta_N = 0.19$ and $\delta_I = 0.17$, both far above ε. Away from the mask, the residual is small near 9 GHz ($|\Delta| \leq 0.04$ between 8.5 and 9.5 GHz), where $|S_{11}|$ is already low. The conditioning σ_N (C3) and the bound (6) were not evaluated because they require the forward surrogate and a second design point. Consequently, the four-configuration analysis does not support factorised inversion at this design point; rather, it demonstrates that the proposed separability test is operational and falsifiable, confirming that all claims regarding factorised inversion must remain strictly conditional on (C1)–(C3) holding with a positive rejection margin.

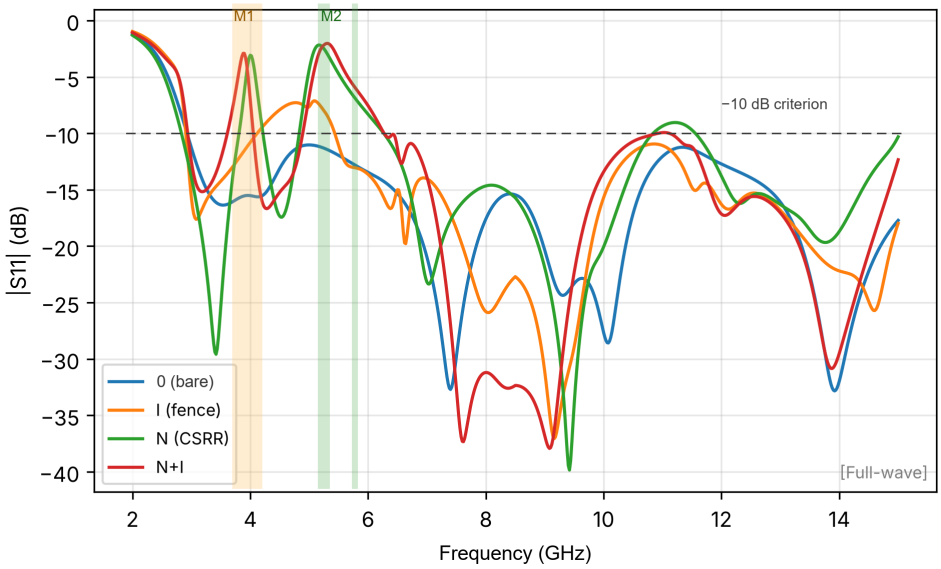


Figure 12. Full-wave. Reflection of the four configurations of the two-port aperture: bare (0), fence only (I), CSRR only (N) and complete $(N + I$, fabricated$)$, with the coexistence masks M1 (C-band) and M2 (WLAN) shaded and the -10 dB criterion dashed. M3 is their union.

Table 9. Full-wave. Four-configuration test of Section 3.3 at the fabricated design point. ε is the measured rejection margin of Section 3.4.

Quantity	Value	Remark
$\|\Delta\|_{\mathcal{N}} = \max_{\mathrm{supp}(N)} \|\Delta_{11}\|$	0.61	at 3.96 GHz
$\delta_N = \|S_{11}^{I} - S_{11}^{0}\|_{\mathcal{N}}$	0.19	at 5.16 GHz
$\delta_I = \|S_{21}^{N} - S_{21}^{0}\|_{\mathcal{B}}$	0.17	over 2.1 GHz to 15 GHz
ε (Section 3.4)	$\leq$0.019	rejection margin
Verdict at 0.5ε, ε, 2ε (0.0095, 0.019, 0.038)	(C1) not satisfied at all three levels	$\|\Delta\|_{\mathcal{N}} \approx 32\varepsilon$

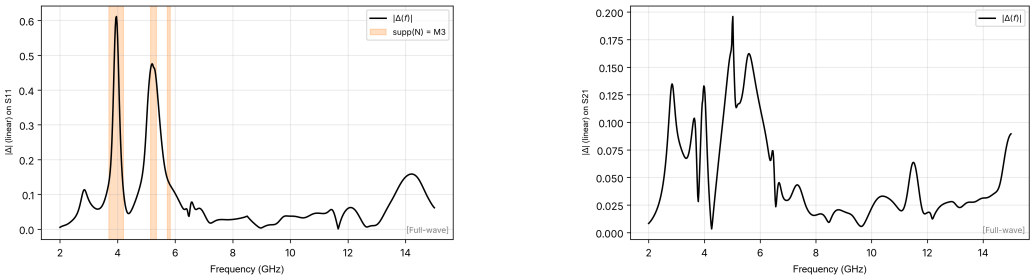


Figure 13. Full-wave. Interaction residual $|\Delta(f)|$ of Equation (4). Left: on S_{11}, shaded: $\mathrm{supp}(N) = \mathrm{M3}$. Right: on S_{21}.

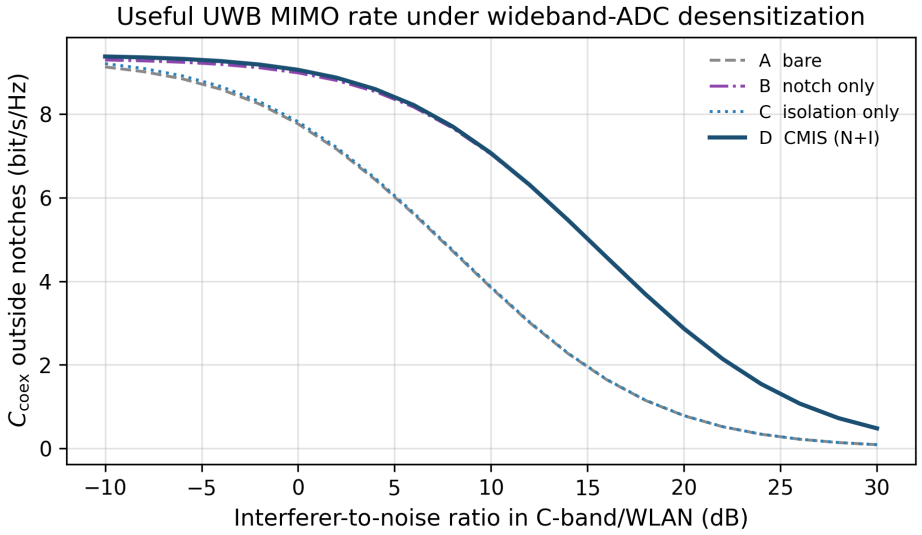


Figure 14. Modelled envelopes, not measured data. First-order power budget: useful MIMO rate (UWB minus notches) versus connector-referred interferer-to-noise ratio, $\mathrm{SNR}_0 = 15$ dB, $\eta_{\mathrm{notch}} = 0.12$. Cases: A, bare; B, notch only; C, isolation only; D, CMIS $(N + I)$. B and D overlap. Computed on the reconstructed envelopes of Section 6.1; see Table 11 for the sensitivity to η_{notch}.

6.3. Mask Compliance

Table 10 reports the fraction of each mask interval over which the aperture rejects, i.e., $|S_{11}| \geq -10$ dB as in the definition (1), and the fraction of the array band over which $|S_{21}| \leq -20$ dB, with the corresponding digitised compliance curves depicted in Figure 15. The measured values are read on the curves of Figure 1, digitised with a resolution of about 0.03 GHz and 0.25 dB; the ranges in brackets move the decision threshold by ±0.5 dB and bracket the reading error. A second column applies a −15 dB threshold, which reproduces the stop-band edges quoted in the original submission and lets the two readings be compared. The full-wave aperture $(N + I$, Touchstone export) is scored with the same criteria in its own column. The modelled-envelope values of the original submission are kept in a separate column. M3 is the mask the prototype was built for; M1 and M2 are restrictions of the same board and were not fabricated separately, so the table scores one design against three masks, not three designs.

The scoring is unambiguous: the fabricated board does *not* comply with M3 at the −10 dB criterion, and the modelled envelopes of the original submission overstated the compliance. The two-port board rejects 50 % of M1, 33 % of M2 and 44 % of M3. The first WLAN window, 5.15 GHz to 5.35 GHz, is not rejected at all at this criterion, because $|S_{11}|$ stays between about −12 and −14 dB there, whereas the second window is. The isolated element rejects 58 %, 67 % and 61 % of M1, M2 and M3, because its notches are deeper (Section 5). With the −15 dB threshold, the board reaches 75 %, 67 % and 72 %. The full-wave aperture response peaks at −2.9 dB (3.88 GHz) and −2.0 dB (5.30 GHz) and rejects 69 %, 100 % and 80 % of M1, M2 and M3 at the −10 dB criterion. The fabricated board is markedly shallower than its model (−9.5 dB). The FR4 permittivity and loss tangent, the etching accuracy of the narrow slots and the SMA transition are possible contributors; none was isolated, and the gap is reported as an open question. The realised lower notch begins at about 3.9 GHz on the board, whereas the C-band downlink mask opens at 3.7 GHz; 3.7 GHz to 3.8 GHz is left uncovered at either threshold. The board was synthesised for a 3.8 GHz to 4.2 GHz target and therefore under-covers M1 by construction.

Two limits of this scoring must be stated. First, $|S_{11}|$ is only a proxy for rejection. A lossy FR4 radiator can lose power in the CSRR without reflecting it, so a notch of −9.5 dB

does not by itself tell how much blocker power is refused. The quantity that matters to a receiver is the total efficiency, or the realised gain, inside the notch relative to the passband. It was not measured on the fabricated board, and it was not extracted from the full-wave model of the complete aperture. The value $\eta_{\text{notch}} = 0.12$ used in Section 6.4 is therefore an assumption, whose influence is swept in Table 11, and the in-notch efficiency is listed among the open validations (Section 8). Second, the numbers above are read from figures. They carry the reading error provided above and must be replaced by the output of `scripts/mask_compliance_measured.py` on the Touchstone files; only the order of magnitude is meant to be kept. The isolation coordinate is met over most, but not all, of the array band (Section 5). Section 4.6 states what closing the missing C-band interval requires. We report the shortfall because a framework whose selling point is mask inversion should be scored against the mask it did not meet, not only against the one it did.

Implications of partial compliance for the overall demonstration of coexistence-mask inversion. Satisfying the −10 dB rejection criterion over only about 44 % of M3 on the measured prototype—and leaving the lower portion of the C-band mask (3.70 GHz to 3.91 GHz) uncovered—has direct implications for what the present study demonstrates regarding the "coexistence-mask inversion" (CMIS) concept. Specifically, the fabricated two-port board cannot be claimed as a completed end-to-end experimental realisation of mask inversion; instead, it demonstrates the *diagnostic and falsifiable scoring pipeline* of CMIS, exposing three structural bottlenecks in the open-loop transfer from a unit-cell inverse model to a fabricated MIMO aperture:

1. *Point-wise versus interval-wise specification:* The inverse map (12) parameterises each CSRR by its resonance centre f_0 and peak attenuation $|S_{21}|_{\min}$ rather than by its −10 dB stop-band bandwidth BW_{-10} or band-edge frequencies. Even though CSRR1 places its measured two-port notch centre at 4.00 GHz, its 240 MHz rejection width at −10 dB is structurally too narrow to cover the 500 MHz C-band downlink allocation (3.7 GHz to 4.2 GHz), leaving the 3.7 GHz to 3.9 GHz segment unprotected. Inverting an interval mask therefore requires augmenting the feature vector with BW_{-10}, as formulated in Section 4.6.
2. *Uncompensated cell-to-aperture loading and operator interaction:* Because the CSRRs are synthesised on an isolated 50 Ω microstrip line and etched without aperture-level recalibration onto the finite MIMO ground alongside the SRR fence, reactive loading shifts the resonances by δ_{tr} (up to +0.40 GHz for CSRR2) and violates condition (C1) ($\|\Delta\|_{\mathcal{N}} = 0.61$). This displacement moves the upper notch away from the first WLAN sub-band (5.15 GHz to 5.35 GHz), where $|S_{11}|$ drops to −12 to −14 dB and fails the −10 dB threshold.
3. *Threshold sensitivity under FR4 Q-factor damping:* Whereas the full-wave simulated aperture achieves −2.9 and −2.0 dB peak reflections and satisfies the −10 dB criterion over 80 % of M3 (69 % of M1 and 100 % of M2), dielectric, copper and connector losses on the FR4 prototype damp the measured reflection peaks to −9.5 dB. When a notch peak hovers within 0.5 dB of the binary threshold, modest Q-factor damping halves the compliance score (from 80 % to 44 %). Consequently, a complete experimental demonstration of CMIS requires both closed-loop calibration of δ_{tr} and a stricter full-wave synthesis margin on $|S_{11}|_{\text{res}}$ to absorb substrate and fabrication losses.

Table 10. Mask compliance (fraction of the interval that meets the criterion). "Env.": modelled envelopes of the original submission, kept for comparison. "Meas.": two-port fabricated board, read on the measured curves of Figure 1 (digitised); the central value uses the stated threshold and the range moves it by ±0.5 dB. "Elem.": isolated dual-notched element, measured, −10 dB criterion. "FW": full-wave complete aperture ($N + I$), −10 dB criterion. M1 and M2 are restrictions of the M3 response, not separately fabricated boards.

Mask	Intervals (GHz)	Env. [Modelled]	Meas., −10 dB (Range)	Meas., −15 dB	Elem., −10 dB [Meas.]	FW, −10 dB [Full-Wave]
M1 C-band	3.70–4.20	0.84	0.50 (0.07–0.58)	0.75	0.58	0.69
M2 WLAN	5.15–5.35, 5.725–5.825	1.00	0.33 (0.22–0.33)	0.67	0.67	1.00
M3 union	M1∪M2	0.90	0.44 (0.12–0.49)	0.72	0.61	0.80
$\lvert S_{21}\rvert \leq -20$ dB	2.1–15	1.00	0.98 (0.97–0.98)	—	—	0.67

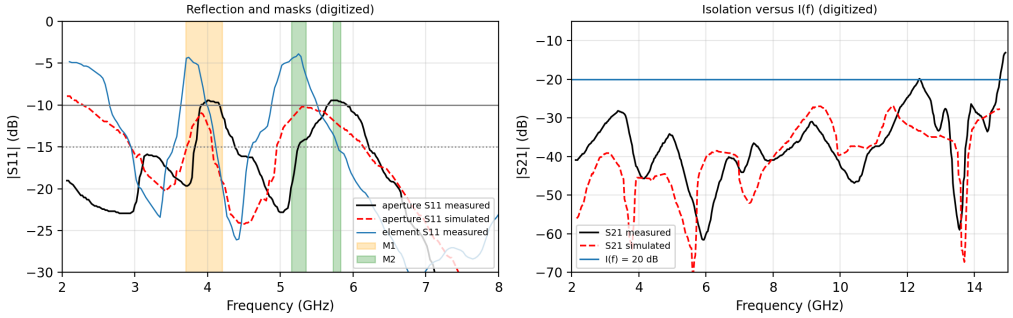


Figure 15. Measured, digitised from Figure 1. Left: reflection of the two-port board (measured and full-wave) and of the isolated element (measured) with the masks M1 and M2 shaded; the −10 dB criterion is the solid line and the −15 dB reading the dotted line. Right: measured and simulated $|S_{21}|$ against $I(f) = 20$ dB. Generated by `scripts/score_digitized.py`.

6.4. A First-Order Power Budget for Blocker Desensitisation [Modelled Envelopes]

The purpose of this subsection is to make explicit *why* a coexistence mask is worth inverting at the aperture, in a currency a system designer recognises. It is a first-order power-budget calculation, not a link-level simulation: there is no channel model, no fading realisation, no receiver nonlinearity model, no adjacent-channel leakage and the blocker is treated as an additive noise floor rather than as a compressing tone. The numbers below are illustrative and inherit every caveat of Section 6.1. They are not an experimental validation of the system.

Consider a blocker occupying N_{M3} in front of a wideband converter with no analogue preselector. Let

$$\alpha(f) = \left(1 - |S_{11}(f)|^2 - |S_{21}(f)|^2\right) \eta(f) \tag{21}$$

be the accepted, radiated power fraction. Outside the stop-bands $\eta = 0.82$, consistent with the simulated efficiency of the FR4 element; inside an etched stop-band, the nominal value is $\eta_{\mathrm{notch}} = 0.12$. The latter is the least constrained input of the model—it was not measured—so its influence is quantified explicitly in Table 11. The interferer-to-noise ratio seen by the converter is

$$\mathrm{INR}(f) = \mathrm{INR}_0 \cdot \mathbb{E}_N[\alpha(f)], \tag{22}$$

i.e., the connector-referred INR scaled by the mask-band average of α. With per-stream SNR $\mathrm{SNR}_0 = 15$ dB in the absence of mismatch and coupling, the instantaneous SINR is $\mathrm{SNR}(f)/(1 + \mathrm{INR}(f))$. For a 2×2 Kronecker correlation with coefficient $\rho = \sqrt{\mathrm{ECC}}$, the equivalent rate is

$$C = \log_2\!\left(1 + \mathrm{SINR}\,(1 + \rho)\right) + \log_2\!\left(1 + \mathrm{SINR}\,(1 - \rho)\right), \tag{23}$$

averaged over the useful band (FCC UWB minus the etched notches). Four combinations form a 2×2 factorial in the operators: A, neither N nor I; B, notch only; C, isolation only; D, both (CMIS). All four share the two-element baseline, and the depressed in-notch efficiency is applied only to the notched cases B and D, so that the comparison isolates the operators rather than a change of reference structure.

At $\mathrm{INR}_0 = 20$ dB, the useful-band rate is about 0.78 bit/s/Hz for A, 2.87 bit/s/Hz for B, 0.79 bit/s/Hz for C and 2.87 bit/s/Hz for D (Figure 14). Two readings matter. The increment D − A of about 2.09 bit/s/Hz is entirely a rejection gain: it is recovered by B alone, and C alone recovers essentially none of it. And D − B is null to within 0.01 bit/s/Hz, so on this metric the isolation operator contributes nothing. That is the expected outcome, not a defect: isolation is not a blocker-rejection mechanism. It is required by the other two coordinates of the mask, $I(f)$ and $\rho_{\max}$ (Figures 9 and 11), which this rate metric does not score. The two operators address two different specifications, and the theoretical rationale for composing them is that they act on distinct coordinates of $\mathcal{M}$ and, under the sufficient conditions (C1)–(C3), allow $\mathcal{M}$ to be inverted by parts, not that both happen to be metamaterials.

Because the rejection gain is driven by how much power the notch refuses to accept, it depends directly on η_{notch}. Table 11 sweeps that value over a range far wider than the plausible uncertainty. The gain degrades monotonically and gracefully, from 2.45 bit/s/Hz at $\eta_{\mathrm{notch}} = 0.05$ to 1.12 bit/s/Hz at $\eta_{\mathrm{notch}} = 0.45$. Even at 0.45—a stop-band that radiates almost half the accepted power, which would be a poor notch—the ordering of the four cases is unchanged and the benefit remains above one bit per second per hertz. The qualitative conclusion is therefore insensitive to the one input we could not measure; the specific value 2.09 is not, and should not be quoted without its sensitivity.

Table 11. Modelled envelopes, not measured data. Sensitivity of the power-budget gain to the in-notch radiation efficiency, at $\mathrm{INR}_0 = 20$ dB. Useful-band rates in bit/s/Hz. The nominal row is $\eta_{\mathrm{notch}} = 0.12$.

η_{notch}	A (Bare)	D (CMIS)	D − A
0.05	0.78	3.23	2.45
0.12	0.78	2.87	2.09
0.20	0.78	2.55	1.77
0.30	0.78	2.24	1.46
0.45	0.78	1.90	1.12

7. Comparison with Related Antennas

Table 12 compares the M3 aperture with representative UWB, MIMO and band-notched designs. The per-element footprint and the frequency-domain impedance bandwidth are competitive, and the upper edge extends into X-band. Isolation is comparable to EBG and parasitic schemes while being obtained at a 10 mm edge-to-edge gap, whereas similar levels are sometimes quoted at much larger separations, 44 mm in [20], for instance. A pattern-based ECC of at most 0.015 (simulated at six frequencies) sits comfortably under the usual 0.5 diversity bound but is explicitly *not* claimed as best-in-class against designs quoting ECC below 0.001 [12,21]; close packing is what penalises correlation here, and that is the constraint the design accepted. Most ECC values in Table 12, including the ones from the cited works, are S-parameter estimates and share the limitation discussed in Section 5.3. Moreover, because our pattern-based ECC is limited to six simulated frequencies without pattern measurements and because group delay and fidelity factor have not been assessed, the wideband comparison in Table 12 applies only to continuous-wave frequency-domain impedance and port-isolation parameters.

Two limits of Table 12 should be stated. Several entries are qualitative because the cited papers report figures rather than tabulated bounds, and a qualitative cell is not a comparison.

And the columns compare hardware, not methodology: the last row wins on none of them decisively—and on the fabricated two-port board the −10 dB rejection criterion covers only 44 % of M3 (Section 6.3). The claim of this work is on a different axis, namely that no referenced band-notched MIMO paper combines a multi-target unit-cell inverse CSRR pipeline (validated in full wave on the two fabricated cells and checked against the training database on eight additional targets), an explicit forward-consistency check, a conditional separability criterion with sufficient conditions and a computable four-configuration test and a fabricated dual-notched MIMO prototype scored quantitatively against an explicit coexistence mask. Table 12 establishes that the baseline frequency-domain hardware is competitive, which is a precondition for that claim rather than the claim itself.

Table 12. Comparison with related UWB/MIMO/band-notched antennas (values from the cited publications).

Ref.	Size (mm^2)	BW (GHz)	Notches (GHz)	Isol. (dB)	ECC	DG (dB)	Distinctive Feature
[10]	30 × 34	3–11	3.4–3.6, 5.1–5.9	n/a (1-el.)	n/a	n/a	Dual square CSRRs in patch
[12]	26 × 31	3.1–11	—	<−25	<0.001	>9.995	Ground stub + EBG
[20]	2-el. MIMO	3.1–13.5	—	≈−23	<0.09	—	Modified EBG decoupling
[19]	compact MIMO	UWB	quadruple notches	high	low	—	Mushroom EBG notches
[21]	54 × 54	3.1–11.8	triple notches	>20	<0.001	>9.99	Parasitic + SRR/EBG
[13]	planar MIMO	UWB	reconfigurable notch	high	low	—	Reconfigurable notch
[22]	40 × 20	2.5–12	3.3–3.8, 5–6.5	<−15	low	—	Self-isolated dual notch
[24]	compact MIMO	SWB	dual notches	high	low	—	Dual-notch SWB MIMO
[23]	compact 4-el.	UWB	triple notches	high	low	—	Close-packed triple-notch MIMO
This work	**28 × 32**/el.	**3.15–16/2.1–15**	**3.91–4.15, 5.66–5.91 (−10 dB)** †	**<−20 (2.1–14.8)**	**≤0.015 (sim. pattern, 6 freq.)**	**≈10**	**CMIS: conditional factorisation (4-layout test), consistent inverse ML**

Two-element board 60 mm wide; edge-to-edge spacing 10 mm, i.e., $0.07\lambda_0$ at the 2.1 GHz lower edge and $0.33\lambda_0$ at 10 GHz. † Measured two-port intervals at the −10 dB mask criterion (44 % of M3; at −15 dB: 3.8 GHz to 4.2 GHz and 4.9 GHz to 6.2 GHz). Isolation is the band-wide bound, not the simulated coupling nulls; measured isolation exceeds −20 dB above about 14.8 GHz. Entries marked "high" or "low" are qualitative in the cited source.

8. Discussion and Limitations

The revision closes several of the gaps identified in the original submission and leaves others open. Both are listed here, and the open items are collected at the end of this section so that each can be checked.

Separability and conditional status of factorised inversion. The criterion is now the set of sufficient conditions (C1)–(C3), with a computable four-configuration test and an explicit error bound (Sections 3.3 and 6.2). The test was performed at one design point, the fabricated one. (C1) is not satisfied ($\|\Delta\|_{\mathcal{N}} = 0.61$ against $\varepsilon \leq 0.019$), mainly because both notches shift by 0.12–0.14 GHz when the fence is added, and (C2) holds only loosely ($\delta_N = 0.19$, $\delta_I = 0.17$). All claims regarding "factorised inversion" in this paper are therefore strictly conditional on (C1)–(C3) and must not be interpreted as a general experimental validation of factorisation. On the contrary, the current prototype falsifies (C1) at the tested design point under the tight margin $\varepsilon \leq 0.019$, showing that open-loop composition of unit-cell CSRR designs with the SRR fence suffers from non-negligible reactive loading. A second design point and the conditioning σ_N remain to be evaluated; even then, the test is a finite-difference check at the tested points, not a proof over the design domain.

Rejection depth and implications for coexistence-mask inversion. The main weakness exposed by the revision is the depth and coverage of the notches on the two-port board. Read on the measured curves, the reflection peaks at −9.5 dB, the −10 dB criterion is met over about

44 % of M3 (leaving 3.70 GHz to 3.91 GHz in the lower C-band and 5.15 GHz to 5.35 GHz in the lower WLAN window uncovered), whereas the full-wave aperture response reaches −2.9 and −2.0 dB and meets it over 80 % of M3; the gap between the board and its model is not explained (Section 6.3). As detailed in Section 6.3, this partial compliance implies that the fabricated prototype demonstrates the diagnostic scoring procedure of CMIS rather than a compliant hardware realisation of coexistence-mask inversion. Achieving full mask compliance requires incorporating the −10 dB notch bandwidth BW_{-10} into the inverse query, calibrating the cell-to-aperture transfer offset δ_{tr} at the array level, and enforcing a deeper full-wave reflection target to overcome FR4 Q-factor damping. The mask definition uses $|S_{11}|$ as a proxy for rejection, but the quantity that protects a receiver is the in-notch efficiency, which has not been measured. Section 4.6 shows that closing M1 requires a stop-band both wider and deeper than the measured one. Whether a single CSRR1 ring can provide both is an open question, which the feasible set (16) answers before any fabrication.

Synthesis, generalisation capability and retargeting. Of the ten synthesised geometries of Tables 4 and 5, the two fabricated cells were verified by a dedicated full-wave solve and by measurement of the antenna. The other eight were checked against the nearest training sample only, which tests consistency with the database and not synthesis accuracy. Because the training set oversamples the neighbourhood of the fabricated T2* cells and independent electromagnetic simulations of the remaining eight configurations have not yet been conducted, the generalisation capability of the inverse synthesis method across the broader C-band and WLAN specification space remains only partially validated (Section 4.5). The M1 retarget of a complete aperture is derived (query, transfer offset, required width) but not solved in full wave. Only the rejection factor is learned; the SRR fence was sized parametrically, and applying (7) to $I(f)$ would make the pipeline uniform.

The system-level section is a modelling study for the power budget only, as stated in Section 6.1. The in-notch efficiency was assumed rather than measured; Table 11 shows the conclusion survives a wide sweep of that assumption, but the assumption remains.

ECC validation and time-domain behaviour. The *S*-parameter ECC is neither a bound nor a reliable value for a lossy FR4 board (Section 5.3). The pattern-based ECC was computed from full-wave embedded patterns at six frequencies only (10° sampling) and was not measured; given the presence of multiple high-*Q* resonators that can cause rapid pattern variations near the notch edges, this six-frequency simulation provides only spot-frequency evidence, and a densely sampled, ideally measured, pattern-based ECC is needed for a broadband diversity claim. The time-domain behaviour (group delay and fidelity factor) was not assessed. Group-delay excursions and late-time ringing are expected inside and near the edges of the stop-bands, since the CSRRs and the SRR fence are resonant, and whether the cells and the fence disturb the group delay and pulse fidelity outside them remains to be shown. All conclusions regarding wideband performance in this work are therefore restricted to continuous-wave frequency-domain impedance and isolation characteristics. FR4 loss limits efficiency at the upper band edge. The +0.4 GHz shift of the second notch of the board is not explained by the permittivity tolerance alone, and it lies at the edge of the first-order worst case of Table 8; it remains an open fabrication question. Measured isolation falls short of 20 dB above about 14.8 GHz, and it has no margin near 12.4 GHz. All measured values quoted from figures are digitised, with the resolution stated in Section 6.3.

Open validations. The following steps are not part of this paper. Each is defined precisely enough here to be carried out, and the post-processing scripts are released with the data (Data Availability Statement):

1. the four-layout residual (4) and the bound (6) at a second design point (the fabricated design is reported in Section 6.2), with σ_N from the forward surrogate;
2. independent full-wave solves of the eight non-fabricated geometries of Tables 4 and 5;

3. the full-waves solve, and then the fabrication, of the M1-retargeted aperture of Section 4.6;
4. the pattern-based ECC (17) over the whole band and from measured embedded patterns (six simulated frequencies are reported in Table 7);
5. the group delay and fidelity factor of a face-to-face link;
6. a full-wave Monte Carlo tolerance analysis replacing the first-order budget of Table 8;
7. the in-notch total efficiency, which replaces the assumed η_{notch} of Section 6.4;
8. a direct comparison with Bayesian optimisation on the same specifications, which replaces the break-even estimate of Section 4.5;
9. the re-scoring of Table 10 from the measured Touchstone files instead of digitised curves.

None of these steps can change the scoring of the fabricated board, which is measured. Items 1–3 and 6 decide whether the factorisation and the synthesis hold beyond the evidence given here.

Extension to more than two elements. For an M-element array, the isolation coordinate becomes the matrix condition $|S_{ij}(f)| \leq -I(f)$ for all $i \neq j$, and the correlation coordinate becomes the set of pairwise ECC_{ij} or, preferably, the correlation matrix computed from the M embedded patterns [28]. Rejection remains a per-element operator, so F_N is unchanged as long as the elements are identical. The fence, however, becomes a set of operators $E_{I,ij}$ placed between adjacent pairs, and the interaction residual (4) must be evaluated for every operator touching a given element. For a linear or 2×2 arrangement, each element is adjacent to at most two fences. In the worst case, this requires 2^{1+q} layouts per element, where q is the number of fences adjacent to it, rather than four. The conditions (C1)–(C3) carry over unchanged. Two new effects are expected. Non-adjacent coupling is not controlled by any fence, and the shared ground carries superposed currents from several ports. Both can only weaken (C2), so separability must be re-verified for each new layout. In a quad-element configuration [23], orthogonal element orientation would reduce the burden placed on the fences.

Integrated sensing and communication. The same front-end is a natural candidate for ISAC terminals, where a single aperture both communicates and senses. System-level ISAC frameworks such as semantic-driven multimodal sensing and communication [31] and beam-forming-based rate maximisation for multi-UAV ISAC networks [32] place requirements on the antenna that go beyond the present masks. First, a *sensing band* must be kept inside $\mathcal{B}$ and away from $\text{supp}(N)$. In UWB radar, this band is wide, and its group-delay dispersion, which this paper does not assess, becomes a design constraint rather than a check. Second, the omnidirectional search and directional tracking phases of [32] impose *radiation-pattern* requirements, such as coverage, beam steering and pattern diversity, that are not coordinates of the present $\mathcal{M}$. Third, the transmit–receive isolation needed for monostatic sensing is typically far stricter than the 20 dB communication value. Within CMIS, these requirements can be added as further coordinates of the mask: a sensing interval with a group-delay ceiling, a pattern-coverage constraint and a frequency-dependent $I(f)$. Whether the corresponding operators remain separable from the rejection operator would then be tested with the same factorial residual.

In summary, a strict distinction must be maintained between the theoretical formulation, the conditions for operator factorisation, and the partial experimental evidence provided by the present hardware: coexistence for a compact MIMO front-end is formulated as mask inversion; metamaterial operators on distinct regions make that inverse conditionally factorizable on $\text{supp}(N)$ *if and only if* (C1)–(C3) hold; while the consistency-checked regressor is validated on held-out unit-cell splits and two fabricated cells, complete experimental

validation of the factorisation, broad multi-target generalisation, full mask compliance and wideband pattern/time-domain performance remain to be established.

9. Conclusions

To keep the conclusions strictly aligned with the level of validation achieved, the findings of this study are summarised along three distinct axes: (i) the proposed theoretical framework, (ii) the conditions under which operator factorisation is valid and (iii) what has actually been demonstrated with the current prototype versus what remains to be validated.

(i) Proposed theoretical framework. Fifth-generation coexistence for a compact UWB MIMO front-end has been formulated as the computational inversion of a multi-coordinate mask $\mathcal{M} = (N(f), I(f), \rho_{\max}, \eta_{\min}, \mathcal{B})$, where $N(f)$ and $I(f)$ act as synthesis inputs while $\rho_{\max}$ and $\eta_{\min}$ serve as aperture-level acceptance constraints. For the rejection coordinate, an inverse machine-learning pipeline trained on 5000 (CSRR1) and 10,000 (CSRR2) full-wave unit-cell samples ($R^2 = 0.9825$ with random forest and $R^2 = 0.9921$ with k-nearest neighbours on held-out test splits) is coupled with a companion forward surrogate to filter out electromagnetically inconsistent predictions, a feasible-set selection rule to resolve non-uniqueness by maximising fabrication margin and a non-parametric out-of-domain detector. A break-even analysis shows that this database-driven pipeline is computationally advantageous when many specifications are queried across design cycles, whereas direct optimisation is cheaper for a single specification.

(ii) Conditions for valid factorisation. Treating CSRR rejection cells and an SRR decoupling fence as independent factors (F_N^{-1} and F_I^{-1}) is formulated strictly as a *conditional approach*. A vanishing mixed derivative alone is insufficient; independent inversion holds with a bounded first-order error $\|\hat{\theta}_N - \theta_N^\star\| \leq (\varepsilon + \delta_N)/\sigma_N$ *if and only if* three sufficient conditions are met over the design domain: uniformly bounded four-configuration interaction $\|\Delta\|_{\mathcal{N}} \leq \varepsilon$ on the support of the rejection mask (C1), bounded cross-sensitivity δ_N and δ_I (C2) and well-conditioned factor Jacobians $\sigma_N, \sigma_I > 0$ (C3), where the threshold ε is dictated by the realised rejection margin and frequency tolerance budget.

(iii) Demonstrated results and open validations. Experimental and full-wave evaluations on the 10 mm-spaced two-port FR4 prototype establish both the capabilities and the current boundaries of the implementation:

1. *Factorisation test:* Executing the four-configuration full-wave protocol at the fabricated design point yields $\|\Delta\|_{\mathcal{N}} = 0.61$ against the measured rejection-margin threshold $\varepsilon \leq 0.019$ (driven by 0.12–0.14 GHz notch shifts when the SRR fence is added) and $\delta_N = 0.19$. Condition (C1) is therefore *not* satisfied at the tested design point, so the current prototype demonstrates the falsifiability of the test rather than an experimental validation of factorised inversion.
2. *Multi-target synthesis and generalisation:* Although ten 5G rejection targets are queried across the C-band and WLAN windows, only the two fabricated unit cells (T2*) have undergone dedicated full-wave simulation and experimental verification, while the remaining eight configurations are checked only against their nearest training-set neighbours. Consequently, the multi-target generalisation capability of the inverse synthesis method is only partially validated.
3. *Coexistence-mask compliance:* In frequency-domain measurements, the prototype satisfies the isolation coordinate $|S_{21}| \leq -20$ dB over 98 % of 2.1 GHz to 15 GHz (2.1 GHz to 14.8 GHz), but satisfies the -10 dB rejection criterion over only 44 % of the combined mask M3 (80 % in full wave), peaking at -9.5 dB and leaving the lower C-band segment (3.7 GHz to 3.9 GHz) uncovered. This partial compliance shows that an end-to-end hardware demonstration of coexistence-mask inversion requires incorporating stop-band bandwidth (BW_{-10}) into the inverse query, calibrating the cell-to-aperture

transfer offset δ_{tr}, and enforcing a deeper full-wave reflection target to overcome FR4 Q-factor damping.

4. *Wideband diversity and time-domain behaviour:* Pattern-based correlation ($ECC_{FF} \leq 0.015$) has been evaluated only from simulated embedded patterns at six spot frequencies without 3D pattern measurements, and neither group delay nor time-domain fidelity factor has been evaluated across the resonant-loaded UWB aperture. All wideband performance conclusions are accordingly restricted to continuous-wave frequency-domain matching and isolation.

Completing the independent full-wave simulations of the eight non-fabricated unit-cell geometries, testing (C1)–(C3) at a second design point with deeper notches, solving and fabricating the bandwidth-retargeted C-band aperture, extending pattern-based ECC to continuous broadband measurements and characterising the time-domain group delay and fidelity factor constitute the next steps required to establish a complete experimental validation.

Author Contributions: Conceptualisation, W.F. and D.B.; methodology, A.M.; software, W.F.; validation, W.F., D.B. and H.S.; formal analysis, W.F. and A.M.; investigation, W.F.; resources, D.B., H.S. and M.B.M.; data curation, W.F.; writing—original draft preparation, W.F. and A.M.; writing—review and editing, D.B., M.B.M. and H.S.; visualisation, A.M.; supervision, H.S.; project administration, H.S.; funding acquisition, M.B.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research work was funded by Umm Al-Qura University, Saudi Arabia, under grant number: 26UQU4331171GSSR07.

Data Availability Statement: The unit-cell training sets, sampling intervals, selected hyperparameters, regression scripts, the digitised curves and the scripts for the open validations listed in Section 8, and the measured *S*-parameter files of the prototypes, together with the full-wave Touchstone files of the four aperture configurations and the embedded far-field patterns used in Table 7, will be deposited in a public repository upon acceptance; they are available from the corresponding author in the meantime.

Acknowledgments: The authors extend their appreciation to Umm Al-Qura University, Saudi Arabia, for funding this research work through grant number: 26UQU4331171GSSR07.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Federal Communications Commission. *Revision of Part 15 of the Commission's Rules Regarding Ultra-Wideband Transmission Systems*; Technical Report First Report and Order, ET Docket 98-153; FCC: Washington, DC, USA, 2002.
2. Fathallah, W.; Abdelhamid, C.; Baccouch, C.; Mohammad, A.; Jouili, K.; Sakli, H. UWB Printed MIMO Antennas for Satellite Sensing System (SRSS) Applications. *Int. J. Adv. Comput. Sci. Appl.* **2024**, *15*, 650–657. [CrossRef]
3. Mansour, A.; Baccouch, C.; Fathallah, W.; Sakli, H.; Ben Salah, B. UWB Antenna for V2V Communication Systems. In Proceedings of the 7th IEEE International Conference on Advanced Technologies, Signal and Image Processing (ATSIP), Sousse, Tunisia, 11–13 July 2024; pp. 1–4. [CrossRef]
4. Sharawi, M.S.; Ikram, M.; Shamim, A. A Two Concentric Slot Loop Based Connected Array MIMO Antenna System for 4G/5G Terminals. *IEEE Trans. Antennas Propag.* **2017**, *65*, 6679–6686. [CrossRef]
5. Choi, N.; Jung, C.; Byun, J.; Harackiewicz, F.J.; Park, M.J.; Chung, Y.S.; Kim, B.; Lee, B. Compact UWB Antenna with I-Shaped Band-Notch Parasitic Element for Laptop Applications. *IEEE Antennas Wirel. Propag. Lett.* **2009**, *8*, 580–582. [CrossRef]
6. Mukherjee, S.; Roy, A.; Mukherjee, A.; Kundu, S.; Sarkar, P.P.; Bhunia, S. Notch band characteristics improvement of a printed ultra wideband antenna by embedding frequency Selective Surface. *AEU-Int. J. Electron. Commun.* **2024**, *178*, 155276. [CrossRef]
7. Zadehparizi, F. Design of low-profile SWB monopole antenna with band-notched characteristics. *AEU-Int. J. Electron. Commun.* **2025**, *198*, 155865. [CrossRef]
8. Falcone, F.; Lopetegi, T.; Laso, M.A.G.; Baena, J.D.; Bonache, J.; Beruete, M.; Marques, R.; Martin, F.; Sorolla, M. Babinet principle applied to the design of metasurfaces and metamaterials. *Phys. Rev. Lett.* **2004**, *93*, 197401. [CrossRef] [PubMed]

9. Baena, J.D.; Bonache, J.; Martin, F.; Sillero, R.M.; Falcone, F.; Lopetegi, T.; Laso, M.A.G.; Garcia-Garcia, J.; Gil, I.; Portillo, M.F.; et al. Equivalent-circuit models for split-ring resonators and complementary split-ring resonators coupled to planar transmission lines. *IEEE Trans. Microw. Theory Tech.* **2005**, *53*, 1451–1461. [CrossRef]
10. Lai, H.; Lei, Z.Y.; Xie, Y.J.; Ning, G.L.; Yang, K. UWB antenna with dual band rejection for WLAN/WiMAX Bands using CSRRs. *Prog. Electromagn. Res. Lett.* **2011**, *26*, 69–78. [CrossRef]
11. Kumar, A.; Ansari, A.Q.; Kanaujia, B.K.; Kishor, J.; Matekovits, L. A review on different techniques of mutual coupling reduction between elements of any MIMO antenna. Part 1: DGSs and parasitic structures. *Radio Sci.* **2021**, *56*, e2020RS007122. [CrossRef]
12. Khan, A.; Bashir, S.; Ghafoor, S.; Qureshi, K.K. Mutual coupling reduction using ground stub and EBG in a compact wideband MIMO-antenna. *IEEE Access* **2021**, *9*, 40972–40979. [CrossRef]
13. Saleem, S.; Kumari, S.; Yadav, D.; Siddiqui, M.G.; Bhatnagar, D. A planar UWB-MIMO antenna with high isolation and reconfigurable single band-elimination characteristics. *AEU-Int. J. Electron. Commun.* **2023**, *170*, 154853. [CrossRef]
14. Sakli, H.; Abdelhamid, C.; Essid, C.; Sakli, N. Metamaterial-Based Antenna Performance Enhancement for MIMO System Applications. *IEEE Access* **2021**, *9*, 38546–38556. [CrossRef]
15. Wu, Q.; Chen, W.; Chen, Y.; Wang, H.; Hong, W. Machine-Learning-Assisted Optimization for Antenna Geometry Design. *IEEE Trans. Antennas Propag.* **2024**, *72*, 2083–2095. [CrossRef]
16. Sarkar, D.; Khan, T.; Talukdar, F.A.; Rengarajan, S.R. Machine Learning Assisted Hybrid Electromagnetic Modeling Framework and Its Applications to UWB MIMO Antennas. *IEEE Access* **2023**, *11*, 19645–19656. [CrossRef]
17. Sarkar, D.; Khan, T.; Talukdar, F.A.; Kishk, A.A. Prior-Knowledge Incorporated Forward and Reverse Neural Network Modeling for Estimating Performance Parameters of Band-Notched UWB MIMO Antennas. *IEEE Antennas Propag. Mag.* **2024**, *66*, 50–63. [CrossRef]
18. Adegboye, O.; Li, X.; Qian, L. AI-Driven Inverse Antenna Design in UWB Band: Embedding Forward Consistency Into Deep Learning Architectures. *IEEE Trans. Antennas Propag.* **2025**, *74*, 1356–1370. [CrossRef]
19. Modak, S.; Khan, T. A slotted UWB-MIMO antenna with quadruple band-notch characteristics using mushroom EBG structure. *AEU-Int. J. Electron. Commun.* **2021**, *134*, 153673. [CrossRef]
20. Fadehan, G.A.; Olasoji, Y.O.; Adedeji, K.B. Mutual Coupling Effect and Reduction Method with Modified Electromagnetic Band Gap in UWB MIMO Antenna. *Appl. Sci.* **2022**, *12*, 12358. [CrossRef]
21. Abbas, A.; Hussain, N.; Sufian, M.A.; Awan, W.A.; Jung, J.; Lee, S.M.; Kim, N. Highly selective multiple-notched UWB-MIMO antenna with low correlation using an innovative parasitic decoupling structure. *Eng. Sci. Technol. Int. J.* **2023**, *43*, 101440. [CrossRef]
22. Addepalli, T.; Nagaraju, V.S.; Kilaru, M.; Kumar, G.N.; Kamili, J.B.; Kumar, C.M. A Compact Self Isolated UWB-MIMO Antenna with WiMAX and WLAN Band Notched Characteristics for Portable Wireless Applications. *Wirel. Pers. Commun.* **2024**, *135*, 1363–1382. [CrossRef]
23. Jayant, S.; Srivastava, G.P. Close-Packed Quad-Element Triple-Band-Notched UWB MIMO Antenna With Upgrading Capability. *IEEE Trans. Antennas Propag.* **2023**, *71*, 353–360. [CrossRef]
24. Shekhawat, S.S.; Lodhi, D.; Singhal, S. Dual band notched superwideband MIMO antenna for 5G and 6G applications. *AEU-Int. J. Electron. Commun.* **2024**, *184*, 155419. [CrossRef]
25. Chen, W.; Wu, Q.; Yu, C.; Wang, H.; Hong, W. Multibranch Machine Learning-Assisted Optimization and Its Application to Antenna Design. *IEEE Trans. Antennas Propag.* **2022**, *70*, 4985–4996. [CrossRef]
26. Gao, P.; Wang, W.; Liu, Y.; Liu, Q.; Chen, Y.; Yan, X. Design of ultra-wideband antenna based on one-dimensional convolutional neural network. *Int. J. Microw. Wirel. Technol.* **2025**, *17*, 482–491. [CrossRef]
27. Pendry, J.B.; Holden, A.J.; Robbins, D.J.; Stewart, W.J. Magnetism from conductors and enhanced non-linear phenomena. *IEEE Trans. Microw. Theory Tech.* **1999**, *47*, 2075–2084. [CrossRef]
28. Honma, N.; Murata, K. Correlation in MIMO Antennas. *Electronics* **2020**, *9*, 651. [CrossRef]
29. Blanch, S.; Romeu, J.; Corbella, I. Exact representation of antenna system diversity performance from input parameter description. *Electron. Lett.* **2003**, *39*, 705–707. [CrossRef]
30. Sharawi, M.S. Printed multi-band MIMO antenna systems and their performance metrics. *IEEE Antennas Propag. Mag.* **2013**, *55*, 203–218. [CrossRef]
31. Peng, Y.; Xiang, L.; Yang, K.; Jiang, F.; Wang, K.; Wu, D.O. SIMAC: A Semantic-Driven Integrated Multimodal Sensing and Communication Framework. *IEEE J. Sel. Areas Commun.* **2026**, *44*, 673–688. [CrossRef]
32. Zhou, S.; Xiang, L.; Yang, K.; Wong, K.K.; Wu, D.O.; Chae, C.B. Beamforming-based Achievable Rate Maximization in ISAC System for Multi-UAV Networking. *arXiv* **2025**, arXiv:2507.21895. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.