

Joint RIS Phase and Power Allocation for a Hardware Impaired OTFS Waveform in Dense Vehicular Integrated Sensing and Communication

B P Vidyashree

PES College of Engineering

Punith Kumar MB

PES College of Engineering

Purushottama Lingadevaru

`purushothamatl@sit.ac.in`

Siddaganga Institute of Technology

Research Article

Keywords: Integrated sensing and communication, reconfigurable intelligent surface, OTFS, hardware impairment, vehicular networks, Cramér-Rao bound

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10287170/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Joint RIS Phase and Power Allocation for a Hardware Impaired OTFS Waveform in Dense Vehicular Integrated Sensing and Communication

Vidyashree B P^{1,*}

Punith Kumar M B¹

Purushottama Lingadevaru²

¹Department of Electronics and Communication Engineering, PES College of Engineering, Mandya, Karnataka, India

²Department of Electronics and Communication Engineering, Siddaganga Institute of Technology, Tumkur, Karnataka, India

* Corresponding author. E-mail: purushothamatl@sit.ac.in

All authors contributed equally to this work.

Abstract

Reconfigurable intelligent surfaces are increasingly proposed to support integrated sensing and communication in vehicular networks, and most of the resulting designs optimize the surface phase configuration and the sensing accuracy bound under the assumption of an ideal radio front end. Separately, a number of studies have quantified how oscillator phase noise, finite resolution data converters and power amplifier nonlinearity erode both link reliability and target estimation accuracy, almost always without a reconfigurable surface in the loop. This paper closes that gap for dense, smart city vehicular deployments by proposing RAHI-OTFS, a reconfigurable intelligent surface aware, hardware impaired orthogonal time frequency space waveform. The surface phase configuration, the delay-Doppler power allocation and the joint Cramér-Rao bound on delay and Doppler are optimized together through an alternating procedure, while the nonlinear power amplifier, the finite resolution analogue to digital converter and the oscillator phase noise are written directly into the transmit and receive signal model rather than bolted on afterwards. A complete derivation of the Fisher information for the impaired link is given, together with a low complexity solution algorithm. Simulation results for a representative roadside scenario show that the optimized configuration reduces the range root mean square error from 4.9 cm to 0.9 cm relative to a surface with random phase, and brings the bit error rate down from 0.28 to 0.16 under identical hardware impairments, while an idealized hardware baseline confirms that residual impairments still cost roughly five orders of magnitude in bit error rate relative to a perfect front end. These results indicate that surface optimization and hardware realism need to be treated as a single design problem rather than two separate ones.

Keywords: Integrated sensing and communication, reconfigurable intelligent surface, OTFS, hardware impairment, vehicular networks, Cramér-Rao bound

Principal Symbols


Symbol	Meaning
M, N_D	Number of delay bins, number of Doppler bins
$\Delta f, T, B, T_f$	Subcarrier spacing, symbol period, bandwidth, frame duration
N_{RIS}	Number of RIS elements
K, L	Number of communication vehicles, sensing targets
$\boldsymbol{\theta}$	RIS phase-shift vector ($ \theta_i = 1$)
$\mathbf{W}$	Delay-Doppler power allocation matrix
$G(\boldsymbol{\theta})$	Cascaded RIS power gain toward the sensing target
N_0	Thermal noise power
$N_{0,\text{eff}}$	Effective noise power including hardware impairment floors
τ, ν_D	Target delay (s), target Doppler shift (Hz)
ν_{PA}	Rapp PA smoothness parameter (dimensionless)
β	Oscillator 3 dB linewidth (Hz)
ρ	ADC quantization efficiency factor (dimensionless)
ρ_s	Sensing power fraction in the power split (dimensionless)
η	Communication-sensing trade-off weight ($\eta \in [0, 1]$)
κ	Multi-user interference leakage factor
$\mathbf{F}, \mathbf{C}$	Fisher information matrix, Cramér-Rao bound matrix
c, λ, f_c	Speed of light, carrier wavelength, carrier frequency

1 Introduction

Roads in a modern city are becoming a wireless environment in their own right. Roadside units talk to vehicles, vehicles talk to each other, and increasingly the same radio hardware is asked to also work out where the surrounding vehicles actually are. This convergence of sensing and communication onto one waveform and one piece of hardware, generally called integrated sensing and communication, is now treated as one of the defining capabilities expected of sixth generation networks [1–3]. Two characteristics of the vehicular setting make the design problem harder than the indoor or fixed wireless case that much of the early integrated sensing and communication literature was built around. First, vehicles move quickly enough that the channel changes within a single transmission frame, which pushes the choice of waveform toward something that represents the channel directly in the delay and Doppler domain rather than the time and frequency domain. Second, a city street is full of buildings, parked vehicles and other obstructions, so a clean line of sight between a roadside unit and every vehicle of interest cannot be assumed.

Reconfigurable intelligent surfaces address the second problem by giving a network operator a cheap, passive way to bend a signal around an obstruction [4, 5]. A surface placed on a building facade or a lamp post can reflect a transmitted signal toward a vehicle that has no direct path to the roadside unit, and because the surface element weights are software controlled, the same hardware can be reconfigured continuously as the geometry of the road changes. The first problem, channel variation within a frame, has a fairly well established answer in orthogonal time frequency space modulation, which places information symbols on a grid indexed by delay and Doppler rather than by subcarrier and time slot [6]. Because a doubly dispersive vehicular channel tends to look sparse and nearly static in the delay-Doppler domain, orthogonal time frequency space has repeatedly been shown to tolerate Doppler spread better than orthogonal frequency division multiplexing in high mobility settings [7–10], and a recent hybrid scheme published in this journal demonstrates the practical value of switching between the two waveforms as a function of estimated vehicle speed [11].

Put a reconfigurable surface and an orthogonal time frequency space waveform together, and a fairly large literature on jointly optimizing the surface phase configuration against a sensing accuracy bound and a communication rate has appeared over the last four years [12–16]. This line of work, reviewed in more detail in Section 2, has produced solid alternating optimization, semidefinite relaxation and learned unfolding algorithms for the joint design problem. What it has almost never done is ask what happens to that same design once the transmitter is built from real, inexpensive radio frequency components rather than an



Graphical Abstract. Overview of the RAHI-OTFS system: a hardware impaired roadside transmitter feeds a reconfigurable surface that serves communicating vehicles and illuminates a sensing target, with the resulting sensing bound and communication rate driving an alternating optimization loop that updates both the surface phase configuration and the delay-Doppler power allocation.

idealized linear front end. A second, largely separate literature has asked exactly that question for integrated sensing and communication systems without a reconfigurable surface, and has shown that oscillator phase noise, finite resolution analogue to digital conversion and power amplifier nonlinearity introduce error floors in both bit error rate and target range or velocity accuracy that no amount of additional transmit power can remove [17–19]. These two literatures rarely meet, and the omission is not a minor one: a surface phase configuration that is optimal under a noiseless, linear, infinite resolution transceiver model is not guaranteed to remain sensible once that transceiver is replaced with the kind of roadside unit a city would actually be willing to buy in bulk.

The principal novelty of this work is to treat reconfigurable surface optimization and transceiver hardware realism as one design problem for a dense vehicular deployment, rather than two design problems solved in isolation and never recombined. The contributions of the paper are as follows.

- (i) A waveform design problem, named RAHI-OTFS, is formulated in which the surface phase configuration, the delay-Doppler power allocation and the joint Cramér-Rao bound on target delay and Doppler are optimized together, rather than the surface being tuned against communication metrics alone with sensing treated as an afterthought.
- (ii) Three hardware impairments, namely power amplifier nonlinearity, finite resolution analogue to digital conversion and oscillator phase noise, are written directly into the transmit and receive signal chain that feeds both the communication and the sensing metric, so that the optimization is performed against the system the impairments actually produce rather than an idealized stand in for it.
- (iii) A full Fisher information and Cramér-Rao bound derivation is given for the delay-Doppler pair under the combined surface gain and hardware noise floor, together with a low complexity alternating

optimization algorithm that updates the surface phase and the power allocation in turn.

- (iv) The design is validated for a representative dense roadside scenario with four communicating vehicles and three sensing targets, and benchmarked against a random surface configuration and an idealized hardware configuration, isolating the separate contributions of surface optimization and hardware realism to the final performance.

The rest of the paper is organized as follows. Section 2 reviews prior work along the three threads identified above and states the gap this paper addresses. Section 3 sets out the geometry, channel and hardware impairment models. Section 4 presents the RAHI-OTFS problem formulation, the Cramér-Rao bound derivation and the optimization algorithm. Section 5 discusses the performance metrics used to evaluate the design. Section 6 presents and discusses the simulation results. Section 7 concludes the paper and outlines its limitations.

2 Related Work

2.1 Waveform design for integrated sensing and communication

Early integrated sensing and communication research largely treated orthogonal frequency division multiplexing as the default waveform, since its frequency domain structure is convenient for both data modulation and radar style range processing [20, 21]. Orthogonal time frequency space modulation, formulated by representing symbols in the delay-Doppler domain and mapping them to the time-frequency grid through a pair of transforms, was proposed as an alternative that keeps the effective channel close to static within a frame even when the physical channel is rapidly time varying [6]. Comparative studies have found that orthogonal time frequency space achieves a more favorable diversity and error rate trade-off than orthogonal frequency division multiplexing once the channel becomes sufficiently sparse and fast varying [7, 8], and coded performance analyses over genuinely high mobility channels back this up with explicit bit error rate curves [9]. More recent work has looked at how orthogonal time frequency space can coexist with deployed orthogonal frequency division multiplexing infrastructure rather than replacing it outright [10], at low complexity channel estimation for the delay-Doppler grid [22], and at generalizing the underlying transform to a wider family of delay-Doppler waveforms [23]. A waveform switching scheme published in this journal estimates vehicle speed and selects between orthogonal frequency division multiplexing and orthogonal time frequency space accordingly, which is a pragmatic acknowledgement that neither waveform dominates the other across the full range of vehicular speeds [11].

2.2 Reconfigurable intelligent surface assisted integrated sensing and communication

Once a transmitter is asked to both communicate and sense through a passive reconfigurable surface, the surface phase configuration becomes a shared resource that affects both functions simultaneously, and a substantial body of work has formulated this as a joint optimization problem. Wang et al. designed a dual-function waveform together with the surface phase shift under a constraint on the Cramér-Rao bound for target direction of arrival, solved through an alternating projection between the waveform and a discrete phase codebook [13], and a related formulation jointly optimizes the transmit waveform and the passive beamforming weights for a similar dual-functional radar-communication architecture [12]. An and co-authors extended this style of design to a joint transmit waveform and passive beamforming problem with an explicit radar power constraint [14], while Zhong et al. proposed a joint waveform and beamforming design aimed specifically at the multi-user interference and Cramér-Rao bound trade-off [15]. A more recent model-driven learning approach unfolds the alternating direction method of multipliers into a trainable network to solve the joint waveform and beamforming problem with lower online complexity [16].

Several papers have pushed this line of work toward the active or near active surface case, where the surface itself injects thermal noise that must be accounted for in the joint design. Zhu and co-authors formulate joint transceiver beamforming and reflection design for an active surface assisted system [24], and a companion paper from an overlapping author group derives a Cramér-Rao bound optimization specifically for the active surface noise contribution [25]. Lyu et al. minimize the Cramér-Rao bound for a millimeter wave surface assisted system [26], and Chen and co-authors design a joint beam training and sensing scheme

for a surface assisted architecture [27]. Tunable peak to average power ratio has also been folded into a surface assisted waveform design, showing that the multi-user interference and beampattern similarity trade-off can be traded against peak to average power ratio through a manifold optimization and alternating direction method of multipliers hybrid solver [28].

Within the specifically vehicular setting, Liu and co-authors set out the broad opportunities and open problems for combining reconfigurable surfaces with integrated sensing and communication [29], and the survey by Zhu, Mao and Kato catalogues how surfaces have been proposed across the sixth generation vehicular communication stack [30]. Meng and co-authors design a sensing-assisted communication scheme for surface aided vehicular networks [31], Chen and co-authors propose a surface based vehicle direction of arrival estimation method embedded in an integrated sensing and communication system [32], and Cheng and co-authors develop a robust beamforming design for surface aided sensing and communication aimed at intelligent transportation [33]. Closest to the present paper in spirit is a recent sensing-centric design for surface assisted vehicular integrated sensing and communication that combines alternating optimization with semidefinite relaxation under both a fairness-oriented and a sum-information objective [34]. Reconfigurable surfaces have also been examined more broadly as smart city infrastructure, with the survey by Kisseleff and co-authors and the original smart radio environment proposal by Wu and Zhang setting out the deployment case that the present paper’s roadside scenario draws on [4, 5], and a dedicated review of surface assisted vehicular network protocols and resource allocation rounds out this picture [35].

2.3 Hardware impaired waveform design

Every paper reviewed in the two subsections above assumes, explicitly or implicitly, a transceiver chain free of phase noise, quantization error and amplifier distortion. A separate body of work has examined what happens once that assumption is dropped, almost always without a reconfigurable surface present. Giroto de Oliveira and co-authors give a detailed closed form treatment of how oscillator phase noise distorts an orthogonal frequency division multiplexing radar image, quantifying peak power loss and sidelobe growth in both monostatic and bistatic configurations [17]. Aydogan and co-authors apply a hardware calibrated phase noise model taken from third generation partnership project specifications to a D-band orthogonal frequency division multiplexing system and report range root mean square error floors of four to five centimetres that persist regardless of signal to noise ratio [18]. Kim and co-authors survey near field, ultra-wideband integrated sensing and communication and identify phase noise, in-phase and quadrature imbalance and power amplifier nonlinearity together as the dominant hardware level threats to delay based sensing accuracy [36]. For orthogonal time frequency space specifically, Doosti-Aref and co-authors analyse bit error rate degradation under combined phase noise and in-phase and quadrature imbalance for a millimetre wave vehicular communication link, confirming that orthogonal time frequency space, while more tolerant of Doppler spread than orthogonal frequency division multiplexing, is not immune to front end non-idealities [19].

2.4 Summary and identified gap

Reading the three threads together exposes a fairly clean gap. Reconfigurable surface assisted integrated sensing and communication, including the vehicular specific designs in [31–34], optimizes the surface phase configuration against a sensing bound and a communication metric under an idealized, noiseless and linear transceiver. Hardware impaired waveform analysis, including [17–19], quantifies the resulting bit error rate and range or velocity error floors but does so without a reconfigurable surface anywhere in the signal path. Neither thread, to the knowledge of the present authors, performs the joint optimization in the first thread under the transceiver model used in the second. This matters for a practical reason rather than a purely academic one: a surface configuration computed under the assumption of a noiseless, infinite resolution radio is not automatically the right configuration to deploy once the same network is built from the inexpensive roadside hardware that makes city scale surface deployment economically realistic in the first place. The waveform design proposed in the next two sections is built specifically to close this gap for the dense vehicular case.

3 System Model

3.1 Deployment geometry

A roadside unit equipped with a dual function transmitter is fixed at the origin and serves a stretch of road through a reconfigurable surface mounted on a nearby building facade or lamp post, at position $\mathbf{p}_R \in \mathbb{R}^2$. The roadside unit has no usable direct path to most of the vehicles of interest, which is the point of placing the surface there in the first place, so the surface provides the dominant propagation path for both functions. The surface carries N_{RIS} passive reflecting elements arranged on a uniform rectangular grid and configured by a phase vector

$$\boldsymbol{\theta} = [\theta_1 \ \theta_2 \ \cdots \ \theta_{N_{\text{RIS}}}]^T, \quad |\theta_i| = 1, \ i = 1, \dots, N_{\text{RIS}}. \quad (1)$$

K communicating vehicles are positioned at $\mathbf{p}_{u,k} \in \mathbb{R}^2, k = 1, \dots, K$, travelling at speed $v_{u,k}$, and L further vehicles act as sensing targets at $\mathbf{p}_{t,\ell} \in \mathbb{R}^2, \ell = 1, \dots, L$, travelling at speed $v_{t,\ell}$, with $v_{t,\ell}$ allowed to be positive or negative depending on whether the target is approaching or receding. This arrangement, sketched in Fig. 1, is representative of a signalized intersection or a busy arterial road in a smart city deployment, where the roadside unit and surface are fixed infrastructure and the vehicle population changes continuously.




Figure 1: Roadside deployment geometry used in the simulation study, showing the fixed roadside unit and reconfigurable surface, four communicating vehicles travelling in two lanes, and three sensing target vehicles distributed across the same road segment.

3.2 Reconfigurable surface cascaded channel

The channel from the roadside unit to the surface and from the surface onward to a vehicle at position $\mathbf{p}$ is modelled through far field steering vectors over the uniform rectangular array, scaled by free space path loss and a directional antenna gain budget representative of a roadside unit transmit antenna and an on board unit receive antenna. Writing d for the element spacing, λ for the carrier wavelength and $\mathbf{n}_i$ for the two dimensional grid index of element i , the roadside unit to surface steering vector is

$$\mathbf{g}_1 = \sqrt{G_{\text{ant}}^2 \beta(d_1)} \left[e^{j2\pi \frac{d}{\lambda} \mathbf{n}_{i,1} \sin \alpha_1} \right]_{i=1}^{N_{\text{RIS}}}, \quad (2)$$

where $d_1 = \|\mathbf{p}_R\|$ is the roadside unit to surface distance, α_1 is the corresponding azimuth, $\beta(d) = (\lambda/4\pi \max(d, 1))^2$ is the free space path loss and G_{ant} folds in the roadside unit transmit and vehicle receive antenna gain split evenly across the two hops. The surface to vehicle steering vector $\mathbf{g}_2(\mathbf{p})$ for a node at position $\mathbf{p}$ takes the same form with $d_2 = \|\mathbf{p} - \mathbf{p}_{\text{RIS}}\|$ is the RIS-to-vehicle distance and α_2 is the corresponding

azimuth measured from the surface normal. The cascaded, surface dependent channel gain to that node is then

$$h(\mathbf{p}, \boldsymbol{\theta}) = \mathbf{g}_1^T \text{diag}(\boldsymbol{\theta}) \mathbf{g}_2(\mathbf{p}). \quad (3)$$

Equation (3) is the single quantity through which the surface configuration affects every communicating vehicle and every sensing target simultaneously, and is the reason the joint design in Section 4 treats $\boldsymbol{\theta}$ as a shared resource rather than optimizing it separately for sensing and for communication.

3.3 OTFS frame structure

The transmitted frame occupies M delay bins and N_D Doppler bins (the subscript D distinguishes this from the noise power N_0 introduced in Section 4.2), with subcarrier spacing Δf and symbol duration $T = 1/\Delta f$, giving a signal bandwidth $B = M\Delta f$ and a frame duration $T_f = NT$. Data symbols $x[m, n]$, $m = 0, \dots, M-1$, $n = 0, \dots, N_D-1$, drawn from a unit average power constellation and weighted by a power allocation $W[m, n] \geq 0$, are placed on the delay-Doppler grid. The time domain transmit signal is produced by an inverse symplectic finite Fourier transform followed by a Heisenberg transform; written with explicit, orthonormal scaling on every Fourier pair so that average signal power is preserved exactly regardless of the relative sizes of M and N_D , this is

$$X_{\text{tf}}[m, n] = \frac{1}{\sqrt{N}} \sum_{n'=0}^{N-1} \left(\frac{1}{\sqrt{M}} \sum_{m'=0}^{M-1} x[m', n'] e^{-j2\pi \frac{mm'}{M}} \right) e^{j2\pi \frac{nn'}{N}}, \quad (4)$$

$$x_t[q] = \frac{1}{\sqrt{M}} \sum_{m=0}^{M-1} X_{\text{tf}}[m, n_q] e^{j2\pi \frac{mm_q}{M}}, \quad q = m_q + n_q M, \quad (5)$$

where the column major index q maps the two dimensional time-frequency grid to the one dimensional time domain sample stream. An earlier, unscaled version of (4)-(5) in which the forward and inverse Fourier sums along the M length and N_D length axes are not individually normalized inflates the average transmit power by a factor of M/N_D relative to the symbol domain, an error that cancels exactly under noiseless round trip testing but corrupts the true signal to noise ratio once channel noise is injected between modulation and demodulation; the orthonormal scaling in (4)-(5) avoids this.

3.4 Hardware impairment models

Three transceiver non-idealities are written directly into the signal chain rather than treated as a separate, decoupled analysis. On the transmit side, the power amplifier is the dominant hardware bottleneck; on the receive side, oscillator phase noise and finite resolution analogue to digital conversion (ADC) are modelled in that order, consistent with the physical arrangement in a superheterodyne or zero-IF receiver where the downconverting oscillator precedes the ADC front end [17].

Power amplifier nonlinearity. The transmit signal x_t passes through a solid state power amplifier modelled by the Rapp characteristic

$$y[q] = \frac{x_t[q]}{\left(1 + \left(\frac{|x_t[q]|}{A_{\text{sat}}} \right)^{2\nu} \right)^{1/2\nu_p}}, \quad (6)$$

where ν_p controls the sharpness of the saturation knee and $A_{\text{sat}} = \sqrt{\mathbb{E}[|x_t|^2]} \cdot 10^{\text{IBO}/20}$ sets the saturation point relative to the chosen input back off IBO in decibels. This is the same Rapp formulation used in tunable peak to average power ratio waveform studies for integrated sensing and communication [28].

Receive-side ADC quantization. The receive ADC, modelled with B_{ADC} bits, is represented through the additive quantization noise model (AQNM). This model covers the ADC at the receiver; DAC non-idealities at the transmitter are not separately modelled and are assumed to be subsumed by the overall power allocation in (20), in which the quantized signal is written as a scaled, attenuated copy of the input plus an independent noise term,

$$z[q] = \sqrt{\rho} r[q] + e_q[q], \quad \mathbb{E}[|e_q|^2] = (1 - \rho) \mathbb{E}[|r|^2], \quad (7)$$

where $r[q]$ is the receiver input prior to quantization and $\rho = \rho(B_{\text{ADC}})$ is the quantization efficiency factor, taken from the standard low resolution lookup table for $B_{\text{ADC}} \leq 5$ and from $\rho \approx 1 - \frac{\pi\sqrt{3}}{2}2^{-2B_{\text{ADC}}}$ for higher resolutions.

Oscillator phase noise. The shared local oscillator is modelled as a discrete time Wiener process,

$$\phi[q] = \phi[q-1] + w[q], \quad w[q] \sim \mathcal{N}(0, 2\pi\beta T_s), \quad (8)$$

where β is the oscillator three decibel linewidth in hertz and T_s is the sample period, and the received signal is multiplied by $e^{j\phi[q]}$. The innovation variance in (8) scales linearly with linewidth and sample period, in line with the standard Wiener phase noise model used for oscillator characterization; a squared dependence on linewidth would imply an oscillator that drifts by tens of full rotations within a single frame; for the parameters used in Section 6 a more modest linewidth in the single digit hertz range is adopted, consistent with the moderate phase noise severities reported for sub-6 gigahertz and low millimetre wave front ends [17], as opposed to the considerably harsher D-band regime studied in [18].

No phase tracking or common phase error compensation is applied at the receiver in this paper, which is a deliberate simplification discussed further in Section 7 and consistent with several of the hardware impaired studies reviewed in Section 2 that leave compensation as a separate problem [17, 19].

4 Proposed RAHI-OTFS Waveform Design

The novelty argued for in Section 1 is made concrete in this section. The surface phase vector $\boldsymbol{\theta}$ and the delay-Doppler power allocation $\mathbf{W} = [W[m, n]]$ are optimized against a single objective that combines communication sum rate and sensing accuracy, while the hardware impairments of Section 3 sit inside every term of that objective rather than being checked afterward. This is what distinguishes RAHI-OTFS from the surface optimization literature reviewed in Section 2, where the same two variables are optimized under an implicit noiseless front end assumption.

4.1 Problem formulation

The communication side of the objective is the sum rate across the K vehicles served through the cascaded channel of (3). Writing P_k for the power allocated to vehicle k and κ for a residual multi-user interference leakage factor representing imperfect precoding,

$$\text{SINR}_k(\boldsymbol{\theta}) = \frac{P_k |h(\mathbf{p}_{u,k}, \boldsymbol{\theta})|^2}{N_{0,\text{eff}} + \kappa \sum_{j \neq k} P_j |h(\mathbf{p}_{u,j}, \boldsymbol{\theta})|^2}, \quad R_{\text{sum}}(\boldsymbol{\theta}, \mathbf{W}) = \sum_{k=1}^K \log_2(1 + \text{SINR}_k(\boldsymbol{\theta})), \quad (9)$$

where $N_{0,\text{eff}}$ is the thermal noise power inflated by the analogue to digital converter and phase noise floors derived in Section 4.2 below. The sensing side of the objective is the trace of the joint delay-Doppler Cramér-Rao bound matrix $\mathbf{C}(\boldsymbol{\theta}, \mathbf{W})$ for a representative target, derived in full in Section 4.2. The two are combined through a normalized weighted scalarization. Because R_{sum} has units of bits/s/Hz and $\text{tr}(\mathbf{C})$ has units of $\text{s}^2 + \text{Hz}^{-2}$, the two terms are divided by reference values R_0 and C_0 before being combined, rendering the objective dimensionless and making η interpretable as a true Pareto weight. In the simulation, R_0 is set to the sum rate achieved by the random-surface baseline and C_0 is set to the corresponding trace of the CRB, so that $\eta = 0.5$ corresponds to equal weighting of the fractional improvements in both metrics:

$$\max_{\boldsymbol{\theta}, \mathbf{W}} \quad \frac{\eta R_{\text{sum}}(\boldsymbol{\theta}, \mathbf{W})}{R_0} - \frac{(1 - \eta) \text{tr}(\mathbf{C}(\boldsymbol{\theta}, \mathbf{W}))}{C_0} \quad (10)$$

$$\text{subject to} \quad \sum_{m,n} W[m, n] \leq P_{\text{tot}}, \quad W[m, n] \geq 0 \quad \forall m, n, \quad (11)$$

$$|\theta_i| = 1, \quad i = 1, \dots, N_{\text{RIS}}, \quad (12)$$

$$\text{PAPR}(x_t) \leq \Gamma_{\text{PAPR}}, \quad (13)$$

where $\eta \in [0, 1]$ trades off communication rate against sensing accuracy and traces out a Pareto boundary as it is swept, and the PAPR constraint in (13) is enforced implicitly through the PA input back-off: setting IBO = 6 dB limits the output PAPR to approximately the saturation-level ratio $\Gamma_{\text{PAPR}} = 10^{\text{IBO}/10} \approx 4$ (i.e., 6 dB), consistent with the values reported in Table 1. When the surface is restricted to a b bit discrete phase codebook, (12) is replaced by $\theta_i \in \{e^{j2\pi\ell/2^b} : \ell = 0, \dots, 2^b - 1\}$.

4.2 Fisher information and Cramér-Rao bound under hardware impairment

For a representative target at $\mathbf{p}_{t,1}$ with true delay τ and Doppler shift ν_D , the echo observed across the full delay-Doppler grid follows the standard radar parameter model

$$y[m, n] = \sqrt{E_{\text{tot}} G(\boldsymbol{\theta})} x[m, n] e^{j2\pi(m\Delta f\tau - nT\nu_D)} + w[m, n], \quad (14)$$

where $G(\boldsymbol{\theta})$ denotes $|h(\mathbf{p}_{t,1}, \boldsymbol{\theta})|^2$, the cascaded surface power gain toward the target from (3), $E_{\text{tot}} = \sum_{m,n} W[m, n]$ is the total energy coherently integrated across the frame, and $w[m, n]$ is the effective receiver noise. The quantization model of (7) and the phase noise model of (8) each contribute an additive term to the effective noise floor,

$$N_{0,\text{eff}} = N_0 + \underbrace{(1 - \rho) \frac{E_{\text{tot}}}{MN_D} G(\boldsymbol{\theta})}_{\text{converter floor}} + \underbrace{2\pi\beta T_s \frac{E_{\text{tot}}}{MN_D} G(\boldsymbol{\theta})}_{\text{oscillator floor}}, \quad (15)$$

so that the effective signal to noise ratio entering the Fisher information is

$$\text{SNR}_{\text{eff}} = \frac{E_{\text{tot}} G(\boldsymbol{\theta})}{N_{0,\text{eff}}}. \quad (16)$$

Differentiating the log likelihood of (14) with respect to τ and ν_D and taking expectations over the noise gives the Fisher information matrix. The cross term $F_{\tau\nu_D}$ is proportional to $\sum_{m,n} m \cdot n$; for a centred rectangular grid with $m \in \{-(M-1)/2, \dots, (M-1)/2\}$ and $n \in \{-(N-1)/2, \dots, (N-1)/2\}$, this double sum factorises as $(\sum_m m)(\sum_n n) = 0 \cdot 0 = 0$ because each factor is a symmetric sum over a zero-mean index set. The Fisher information matrix is therefore diagonal [37]:

$$\mathbf{F} = 2 \text{SNR}_{\text{eff}} \begin{bmatrix} (2\pi\Delta f)^2 \sum_{m=-\frac{M-1}{2}}^{\frac{M-1}{2}} m^2 \cdot N_D & 0 \\ 0 & (2\pi T)^2 \sum_{n=-\frac{N_D-1}{2}}^{\frac{N_D-1}{2}} n^2 \cdot M \end{bmatrix}, \quad (17)$$

and the Cramér-Rao bound matrix used in (10) is its inverse,

$$\mathbf{C}(\boldsymbol{\theta}, \mathbf{W}) = \mathbf{F}^{-1} = \begin{bmatrix} C_{\tau\tau} & 0 \\ 0 & C_{\nu\nu} \end{bmatrix}. \quad (18)$$

The corresponding range and velocity root mean square error bounds reported in Section 6 follow from the usual delay to range and Doppler to velocity conversions,

$$\sigma_{\text{range}} = c\sqrt{C_{\tau\tau}}, \quad \sigma_{\text{vel}} = \frac{\lambda}{2}\sqrt{C_{\nu_D\nu_D}}. \quad (19)$$

Equations (15)-(18) are the mechanism through which hardware impairment enters the sensing side of the joint objective: increasing the ADC resolution or reducing the oscillator linewidth lowers $N_{0,\text{eff}}$, raises SNR_{eff} , and tightens both bounds in (19), while the surface configuration $\boldsymbol{\theta}$ enters through $G(\boldsymbol{\theta})$ in exactly the same place that it enters the communication sum rate in (9). This shared dependence on $\boldsymbol{\theta}$ is what makes the two halves of the objective genuinely coupled rather than additive.

4.3 Joint power allocation

The total power budget P_{tot} is split explicitly into a uniform communication sub-budget, which guarantees that every delay-Doppler bin carries usable communication power, and a sensing sub-budget concentrated around the target's expected delay-Doppler bin, which behaves as a dedicated probing pilot for the matched

filter implicit in (17). With $\rho_s \in [0, 1]$ the fraction of power reserved for sensing and (m_0, n_0) the grid bin nearest the target's true delay and Doppler,

$$W[m, n] = \underbrace{\frac{(1 - \rho_s)P_{\text{tot}}}{MN_D}}_{\text{communication floor}} + \underbrace{\rho_s P_{\text{tot}} \cdot \frac{\exp\left(-\frac{(m-m_0)^2}{2\sigma_m^2} - \frac{(n-n_0)^2}{2\sigma_n^2}\right)}{\sum_{m', n'} \exp\left(-\frac{(m'-m_0)^2}{2\sigma_m^2} - \frac{(n'-n_0)^2}{2\sigma_n^2}\right)}}_{\text{sensing boost}}. \quad (20)$$

The choice $\rho_s = 0.30$ represents a design parameter that can be tuned to shift performance between the two functions; the sensitivity of results to this choice is discussed in Section 6. An earlier design considered here fed a single blended communication and sensing gain grid into classical water filling. That approach was rejected because water filling, applied to a gain grid with even a modest peak to floor ratio, concentrates almost the entire power budget onto the single strongest bin and leaves the remaining bins at exactly zero power, which is appropriate for a dedicated radar pilot but starves every other communication symbol regardless of channel quality. The explicit split in (20) keeps the communication floor strictly positive by construction while still allowing a genuine, surface coupled sensing boost.

4.4 Alternating optimization algorithm

Problem (10)-(13) is non-convex in θ on account of the unit modulus constraint, so it is solved by alternating between the power allocation of (20), which has a closed form given θ , and a projected gradient update of θ given $\mathbf{W}$. Algorithm 1 summarizes the procedure. The gradient with respect to θ is estimated through a simultaneous perturbation stochastic approximation rather than an explicit analytic gradient, since the dependence of (17) on θ through $G(\theta)$ inside $N_{0,\text{eff}}$ makes a closed form derivative cumbersome without offering a proportionate benefit at the moderate surface sizes considered here. Discrete phase quantization, when the surface is restricted to a b bit codebook, is applied only once the continuous phase has converged, rather than at every outer iteration; quantizing every iteration was found in preliminary experiments to repeatedly disturb the gradient ascent direction and reduced the converged objective by roughly a factor of thirty relative to the quantize-once-at-convergence strategy used in Algorithm 1.

Algorithm 1 Alternating optimization for RAHI-OTFS

- 1: Initialize $\theta^{(0)}$ with random unit modulus phases, set iteration counter $t \leftarrow 0$
 - 2: **for** $t = 1$ to T_{max} **do**
 - 3: Compute $\mathbf{W}^{(t)}$ from (20) using $\theta^{(t-1)}$
 - 4: **for** $s = 1$ to S_{max} **do**
 - 5: Draw random perturbation direction $\mathbf{v} \in \{-1, +1\}^{N_{\text{RIS}}}$
 - 6: Estimate gradient of (10) along $\mathbf{v}$ by finite difference
 - 7: Update $\theta \leftarrow \theta + \mu_s \cdot (\text{gradient estimate})$, then project: $\theta_i \leftarrow \theta_i / |\theta_i|$
 - 8: **end for**
 - 9: $\theta^{(t)} \leftarrow \theta$
 - 10: **if** $t = T_{\text{max}}$ and discrete surface phases are required **then**
 - 11: Quantize each $\theta_i^{(t)}$ to the nearest of 2^b codebook phases
 - 12: Recompute $\mathbf{W}^{(t)}$ from (20) using the quantized $\theta^{(t)}$
 - 13: **end if**
 - 14: **end for**
 - 15: **return** $\theta^{(T_{\text{max}})}, \mathbf{W}^{(T_{\text{max}})}$
-

Each outer iteration costs $\mathcal{O}(K + L)$ channel evaluations for the power allocation step and $\mathcal{O}(N_{\text{RIS}} S_{\text{max}})$ perturbation evaluations for the phase update step, each itself $\mathcal{O}(K + L)$, giving an overall per-iteration cost of $\mathcal{O}(N_{\text{RIS}} S_{\text{max}}(K + L))$. The power allocation step is a closed form maximizer of its sub-problem given θ . The phase update uses SPSA, which provides a noisy but asymptotically unbiased gradient estimate; because the estimate is noisy, an exact non-decrease guarantee does not hold at every step in the strict sense applicable to exact gradient ascent. In practice, the step size $\mu_s = 0.05/\sqrt{s}$ satisfies the Robbins-Monro

conditions $\sum_s \mu_s = \infty$, $\sum_s \mu_s^2 < \infty$ sufficient for SPSA convergence in expectation [38], and Fig. 2 confirms that the objective rises monotonically for the scenario considered here. Global optimality of the joint non-convex problem is not claimed, consistent with the alternating optimization literature reviewed in Section 2.




Figure 2: Convergence of the joint objective in (10) across outer alternating optimization iterations for the scenario of Fig. 1, with $\eta = 0.5$.

Fig. 2 shows the joint objective rising smoothly across the twenty five outer iterations used throughout Section 6, without oscillation, which is the behaviour expected of the non-decreasing block coordinate argument above. Fig. 3 compares the resulting optimized phase configuration against an unoptimized, randomly drawn surface configuration of the same size, illustrating that the converged solution does not resemble a random draw and has visibly redistributed phase across the array.




Figure 3: Surface element phase shift after joint optimization compared against a randomly drawn configuration of the same array size, used as the unoptimized baseline throughout Section 6.

5 Performance Metrics

Four communication metrics and four sensing metrics are reported in Section 6, chosen to be mutually consistent with the objective in (10) and the constraints in (11)-(13) rather than reported simply because they

are conventional.

On the communication side, spectral efficiency and sum rate follow directly from (9) and are the quantities the alternating optimization in Algorithm 1 is actually maximizing, so reporting them is necessary to show the optimization is doing what it claims. Bit error rate is reported separately because spectral efficiency alone does not show the bit level cost of the impairments in Section 3, which manifest as an error floor rather than a smooth rate reduction once the ADC or the oscillator becomes the dominant impairment. Peak to average power ratio is reported because it is the quantity that sets the power amplifier operating point in (6) and therefore the quantity through which the waveform design connects to the hardware impairment model in the first place.

On the sensing side, the joint Cramér-Rao bound of (18) is reported because it is the quantity directly penalized inside (10), and the corresponding range and velocity root mean square errors of (19) convert that bound into physically interpretable units that map onto a following distance or a closing speed margin relevant to autonomous driving safety. Range resolution, $c/2B$, is reported alongside the root mean square error specifically to make a point that is often blurred in casual discussion of radar performance: resolution is a bandwidth limited quantity that does not depend on signal to noise ratio, whereas the Cramér-Rao bound based estimate can be considerably finer than one resolution cell once signal to noise ratio is high, and the results in Section 6 show exactly this separation. Detection probability is reported as the decision theoretic counterpart to the continuous estimation view taken by the Cramér-Rao bound, computed through a Neyman-Pearson detector at a false alarm probability of 10^{-6} , representative of an automotive radar constant false alarm rate operating point,

$$P_d = Q\left(Q^{-1}(P_{fa}) - \sqrt{\text{SNR}_{\text{eff}}}\right), \quad (21)$$

with SNR_{eff} taken from (16) and $Q(\cdot)$ the standard Gaussian tail function.

Two metrics that could plausibly have been included are deliberately left out. Outage probability is not treated as a quantity directly optimized by Algorithm 1, since the deterministic, geometry based channel model in (3) has no per-realization randomness of its own; outage is only a meaningful, non-degenerate probability once a stochastic small scale fading layer is reintroduced on top of the optimized deterministic configuration, so it is reported separately in Section 6 as a Monte Carlo validation rather than conflated with the metrics the optimizer actually controls. Fairness index is not reported because the current formulation allocates an equal power share to every communicating vehicle rather than a fairness constrained or proportional fair share, and reporting a fairness index without a corresponding fairness aware mechanism in the objective would suggest the design controls something it does not.

6 Results and Discussion

6.1 Simulation setup

The scenario follows Fig. 1: a 28 gigahertz carrier, an orthogonal time frequency space frame with $M = 64$ delay bins and $N_D = 32$ Doppler bins at 120 kilohertz subcarrier spacing, giving a 7.68 megahertz bandwidth, a 64 element reconfigurable surface arranged as an eight by eight array with three bit discrete phase resolution, four communicating vehicles travelling at up to 130 kilometres per hour relative closing speed, and three sensing target vehicles. Unless stated otherwise the total transmit power is 30 dBm, the ADC resolution is six bits ($B_{\text{ADC}} = 6$), the PA input back-off is six decibels ($\text{IBO} = 6$ dB, $\nu_p = 2$), the oscillator linewidth is five hertz ($\beta = 5$ Hz), the trade-off weight in (10) is $\eta = 0.5$, the sensing power fraction in (20) is $\rho_s = 0.30$, the MUI leakage factor is $\kappa = 0.15$, and the AO algorithm runs for $T_{\text{max}} = 25$ outer iterations with $S_{\text{max}} = 15$ SPSA inner steps per iteration with step size $\mu_s = 0.05/\sqrt{s}$. Three configurations are compared throughout: the proposed design with both surface optimization and full hardware impairment active; a baseline using a randomly drawn surface configuration with the same hardware impairment, isolating the value of surface optimization on its own; and an idealized hardware configuration that keeps the optimized surface but disables the PA, ADC and oscillator impairments, isolating the cost of hardware realism on its own.




Figure 4: Sum rate against range root mean square error for the proposed configuration, the random surface baseline, and the idealized hardware configuration, illustrating that surface optimization improves both quantities together rather than trading one against the other.

6.2 Joint communication and sensing trade-off

Table 1 reports every metric from Section 5 for the three configurations. The proposed design reaches a sum rate of 6.31 bits per second per hertz against 5.01 for the random surface baseline, and a range root mean square error of 0.87 centimetres against 4.87 centimetres for the same baseline, so the surface optimization improves communication and sensing together rather than at each other's expense, visible directly in Fig. 4. The bit error rate falls from 0.28 for the random-surface baseline to 0.16 for the proposed design under identical hardware impairment. Both values are well above a practical link-quality threshold (typically $< 10^{-3}$ for coded vehicular links); they characterise the hardware-impairment floor under the simplified channel model rather than representing operational communication quality. The idealized hardware configuration reaches a bit error rate at the simulation floor of 10^{-6} , confirming that the optimized surface phase configuration is capable of supporting a reliable link once front-end impairments are removed. The roughly five order of magnitude gap between the proposed design and its idealized hardware counterpart is the direct, quantified cost of the PA, ADC and oscillator impairments once the surface has already been optimized, and is the primary evidence in the paper for why the two design problems identified in Section 2 cannot safely be solved in isolation.

6.3 Sensitivity to trade-off parameters

The objective weight η and the sensing power fraction ρ_s are the two free design parameters in (10) and (20). To confirm that the choice $\eta = 0.5$, $\rho_s = 0.30$ is not unusually favourable, Table 1 results were re-evaluated at $\rho_s \in \{0.10, 0.20, 0.30, 0.40\}$. The sum rate changes by less than 0.4 bits/s/Hz and the range RMSE changes by less than 1.1 cm across this range, indicating that the joint performance is not strongly sensitive to the precise split for this geometry. Similarly, sweeping η from 0.3 to 0.7 shifts the balance between communication rate and sensing accuracy predictably without changing the qualitative ordering of the three configurations in Table 1. All subsequently reported results use $\eta = 0.5$ and $\rho_s = 0.30$ as nominal values.

6.4 Effect of analogue to digital converter resolution

Figs. 5 and 6 sweep the ADC resolution from two to ten bits with the surface configuration and power allocation held fixed at their optimized values. The range root mean square error in Fig. 6, computed directly from the closed form bound in (19), falls steeply between two and four bits and then flattens, consistent with $\rho(B_{\text{ADC}})$ in (7) approaching unity once the ADC noise floor term in (15) becomes negligible relative

Table 1: Communication and sensing performance for the proposed design against the random surface and idealized hardware baselines.

Metric	Proposed	Random surface	Idealized hardware
Spectral efficiency (bits/s/Hz)	1.58	1.25	1.58
Sum rate (bits/s/Hz)	6.31	5.01	6.31
SINR (dB)	21.3	11.9	21.3
Bit error rate	1.6×10^{-1}	2.8×10^{-1}	1.0×10^{-6}
Peak to average power ratio (dB)	6.5	6.0	11.7
Energy efficiency (Mbit/J)	48.5	38.5	48.5
Range RMSE (cm)	0.87	4.87	0.87
Velocity RMSE (cm/s)	0.45	2.50	0.45
Range resolution (m)	19.5	19.5	19.5
Sensing SNR (dB)	31.7	16.8	31.7
Detection probability ($P_{fa} = 10^{-6}$)	1.00	0.98	1.00




Figure 5: Bit error rate against ADC resolution for the proposed design, estimated by Monte Carlo simulation of the full time domain chain. The broad downward trend is visible but, unlike the closed form bound in Fig. 6, is not perfectly smooth, as discussed in the text.




Figure 6: Range root mean square error against ADC resolution, following the converter noise floor term in (15).

to the oscillator phase noise term, which does not depend on B_{ADC} and therefore sets a floor that additional converter resolution cannot remove. The bit error rate in Fig. 5, estimated by Monte Carlo simulation of the full time domain chain, shows the same broad downward tendency but with visible run to run variation superimposed on it; this is expected, since at the oscillator linewidth used here phase noise is already the larger of the two impairments at every resolution tested, so the converter resolution has a real but secondary effect on bit error rate that a finite number of simulated frames does not average out as cleanly as the closed form bound in Fig. 6. Beyond roughly six to seven bits neither curve improves much further, an observation directly useful for specifying roadside hardware cost against expected performance.

6.5 Effect of surface size




Figure 7: Sum rate and range root mean square error against the number of surface elements, each point obtained by re-running the alternating optimization for that array size.

Fig. 7 repeats the alternating optimization for surface sizes between sixteen and one hundred and forty four elements. Sum rate rises and range root mean square error falls monotonically with array size, as expected since a larger surface offers more degrees of freedom for the phase configuration to direct energy toward both the communicating vehicles and the sensing targets simultaneously, though the marginal benefit per additional element is visibly smaller beyond roughly one hundred elements for this geometry, which is a relevant data point when weighing the manufacturing cost of a larger surface against the achievable gain for a given road segment.

6.6 Effect of transmit power

Fig. 8 sweeps transmit power from minus ten to thirty dBm. The idealized hardware curve falls without bound as expected of an unimpaired link, while the proposed and random surface curves both flatten into an error floor at higher power, the signature of an impairment limited rather than noise limited regime. Fig. 9 adds an auxiliary Rician small scale fading layer with an eight decibel Rician factor on top of the deterministic channel purely for this outage evaluation, since, as noted in Section 5, the deterministic model used everywhere else in the paper has no per-realization randomness for outage to be computed against. The proposed configuration's outage probability begins falling from full outage once the transmit power passes roughly twenty dBm and reaches zero by thirty dBm, while the random surface baseline follows the same general shape delayed by approximately five decibels, only beginning its descent past twenty five dBm. The two curves therefore tell a consistent story with the bit error rate result above: surface optimization buys a several decibel head start in the power budget needed for reliable communication to the non line of sight vehicles in this geometry, rather than being the difference between an always reliable and a never reliable link.




Figure 8: Bit error rate against transmit power budget for the proposed design, the random surface baseline and the idealized hardware configuration.




Figure 9: Communication outage probability, evaluated under an auxiliary Rician small scale fading layer as discussed in Section 5, against transmit power budget for the proposed design and the random surface baseline.

6.7 Peak to average power ratio and detection probability




Figure 10: Left: peak to average power ratio for the three configurations, showing the reduction produced by the power amplifier model of (6) once it is active. Right: target detection probability at a false alarm probability of 10^{-6} .

The left panel of Fig. 10 shows that the idealized hardware configuration, which has no power amplifier compression, exhibits a higher peak to average power ratio than either hardware impaired configuration, confirming that the Rapp model of (6) is doing what it is meant to do at the chosen input back off. The right panel shows detection probability close to unity for the proposed and idealized hardware configurations and a modest reduction for the random surface baseline, a smaller relative effect than the corresponding bit error rate gap because the coherent integration in (16) accumulates energy across the entire frame for the detection statistic, whereas the bit error rate is governed by the considerably lower per-symbol energy.

6.8 Delay-Doppler ambiguity

Representative joint delay-Doppler ambiguity mainlobe




Figure 11: Representative joint delay-Doppler ambiguity mainlobe, included as a qualitative illustration of the mainlobe shape implied by the diagonal FIM structure in (17). The axis scales are indicative only and do not correspond to the exact numerical waveform parameters of the simulation.

Fig. 11 is included as a schematic rather than a literal computation, to give a visual sense of the joint delay-Doppler mainlobe shape implied by the Fisher information structure in (17); the figure caption is deliberately explicit about this distinction.

6.9 Comparison against published designs

It is worth noting that the "random surface" baseline used in Table 1 is a worst-case reference rather than a competitive one; a more informative baseline would be a surface optimized for communication rate alone, without any sensing constraint. Such a configuration would achieve higher communication metrics than the random-phase case and thus narrow the gap relative to the proposed joint design. Adding this baseline is left as future work because it requires a second AO pass with a purely communication objective, which lies outside the scope of the present impairment-focused study.

Table 2 places RAHI-OTFS against five recent designs from Section 2, chosen because each is a joint surface and waveform or beamforming design for integrated sensing and communication. None of the five models more than one hardware impairment, and none targets the dense vehicular setting with an explicit joint delay-Doppler bound; this table is the direct, tabulated counterpart of the gap argued for narratively in Section 2.

Table 2: Comparison of RAHI-OTFS against recent reconfigurable surface assisted integrated sensing and communication designs.

Reference	Waveform	Surface joint optimization	Hardware impairment	Application
Wang et al. [13]	Constant modulus	Phase vs. CRB (DoA)	None	Generic dual function radar
Zhong et al. [15]	Dual function	Phase vs. interference, CRB	None	Generic ISAC
Zhu et al. [25]	Active surface	Phase vs. CRB	Surface noise only	Generic ISAC
Bazzi and Chafii [28]	Tunable PAPR	Phase vs. PAPR, interference	PAPR considered, no PA/ADC/PN model	Generic ISAC
Yang et al. [34]	Beamformed	Phase vs. SNR, sum information	None	Vehicular ISAC
This work (RAHI-OTFS)	OTFS	Phase, power vs. rate, joint delay-Doppler CRB	PA, ADC, phase noise	Dense vehicular ISAC

7 Conclusion and Future Work

This paper set out to answer a fairly specific question: does optimizing a reconfigurable surface for integrated sensing and communication still make sense once the surrounding hardware is allowed to be the kind of imperfect, inexpensive radio a city would actually deploy at scale, rather than the noiseless front end assumed in most of the joint surface and waveform literature. The RAHI-OTFS design proposed here answers that question by folding the surface phase configuration, the delay-Doppler power allocation and a hardware aware Cramér-Rao bound into one alternating optimization problem, with the power amplifier, converter and oscillator impairments written into the signal model rather than checked against it afterward. For the dense roadside scenario simulated in Section 6, surface optimization improved sum rate and range accuracy together, hardware impairment cost roughly five orders of magnitude in bit error rate even after the surface was optimized as well as it could be, and the two effects were shown to be large enough, and different enough in character, that neither can be safely ignored when specifying a real deployment.

The paper has real limitations that should be stated plainly. First, the hardware impairment models, while standard in the literature reviewed in Section 2, are simplified relative to a measured front end: no phase tracking or common-phase-error compensation is applied, the PA model carries no amplitude-to-phase (AM/PM) distortion term, and the ADC model is the Gaussian AQNM approximation rather than a bit-accurate simulation. Second, the channel model uses far-field free-space path loss without multi-path scattering; the bit error rate values reported (0.16 and 0.28 for the hardware-impaired configurations) are therefore best interpreted as characterising the impairment floor of this simplified model rather than as predictions of operational link quality in a dense urban environment. Third, no DAC non-ideality is modelled on the transmit side, and the interaction between PA compression and OTFS waveform shaping is captured only through the Rapp AM/AM characteristic without an AM/PM term. The alternating optimization in Algorithm 1 is not guaranteed to reach a global optimum of the non-convex joint problem, and the simultaneous perturbation gradient estimate, while computationally convenient at the surface sizes considered here, would need revisiting at much larger array sizes. The vehicular geometry simulated is representative rather than drawn from measured road traffic data. Future work should address phase noise compensation within the same joint framework, extend the channel model to a genuinely multipath, ray traced urban environment, and examine whether a learned or unfolded solver of the kind used in [16] can close the gap to a better local optimum at lower online computational cost than Algorithm 1.

Declarations

Funding

This research received no external funding.

Conflict of Interest

The authors declare that they have no conflict of interest.

Author Contributions

All authors contributed equally to this work.

Data Availability

The simulation code and data generated for this study are available from the corresponding author on reasonable request.

References

- [1] Y. Cui, H. Ding, L. Zhao, and J. An. Integrated sensing and communication: a network level perspective. *IEEE Wireless Communications*, 31(1):103–109, 2024.
- [2] A. Kaushik et al. Toward integrated sensing and communications for 6G: key enabling technologies, standardization, and challenges. *IEEE Communications Standards Magazine*, 8(2):52–59, 2024.
- [3] J. A. Zhang, M. L. Rahman, K. Wu, X. Huang, Y. J. Guo, S. Chen, and J. Yuan. Enabling joint communication and radar sensing in mobile networks: A survey. *IEEE Communications Surveys & Tutorials*, 24(1):306–345, 2022.
- [4] Q. Wu and R. Zhang. Towards smart and reconfigurable environment: intelligent reflecting surface aided wireless network. *IEEE Communications Magazine*, 58(1):106–112, 2019.
- [5] S. Kisseleff, W. A. Martins, H. Al-Hraishawi, S. Chatzinotas, and B. Ottersten. Reconfigurable intelligent surfaces for smart cities: research challenges and opportunities. *IEEE Open Journal of the Communications Society*, 1:1781–1797, 2020.
- [6] S. K. Mohammed, R. Hadani, A. Chockalingam, and R. Calderbank. OTFS: a mathematical foundation for communication and radar sensing in the delay-doppler domain. *IEEE BITS the Information Theory Magazine*, 2(2):36–55, 2022.
- [7] L. Gaudio, G. Colavolpe, and G. Caire. OTFS vs. OFDM in the presence of sparsity: a fair comparison. *IEEE Transactions on Wireless Communications*, 21(6):4410–4423, 2022.
- [8] P. Raviteja, Y. Hong, E. Viterbo, and E. Biglieri. Effective diversity of OTFS modulation. *IEEE Wireless Communications Letters*, 9(2):249–253, 2020.
- [9] S. Li, J. Yuan, W. Yuan, Z. Wei, B. Bai, and D. W. K. Ng. Performance analysis of coded OTFS systems over high-mobility channels. *IEEE Transactions on Wireless Communications*, 20(9):6033–6048, 2021.
- [10] A. Shafie, J. Yuan, P. Fitzpatrick, T. Sakurai, and Y. Fang. On the coexistence of OTFS modulation with OFDM-based communication systems. *IEEE Transactions on Communications*, 72(11):6822–6838, 2024.
- [11] Amina Darghouthi, Abdelhakim Khelifi, Belgacem Chibani, and Hmaied Shaiek. OTFS vs OFDM waveforms switching use based on mobile speed estimation. *Wireless Personal Communications*, 139:1499–1513, 2024.
- [12] X. Wang, Z. Fei, Z. Zheng, et al. Joint waveform design and passive beamforming for RIS-assisted dual-functional radar-communication system. *IEEE Transactions on Vehicular Technology*, 70:5131–5136, 2021.

- [13] X. Wang, Z. Fei, J. Huang, and H. Yu. Joint waveform and discrete phase shift design for RIS-assisted integrated sensing and communication system under cramér-rao bound constraint. *IEEE Transactions on Vehicular Technology*, 71:1004–1009, 2022.
- [14] D. An, J. Hu, and C. Huang. Joint design of transmit waveform and passive beamforming for RIS-assisted ISAC system. *Signal Processing*, 204:108854, 2023.
- [15] K. Zhong, J. Hu, C. Pan, M. Deng, and J. Fang. Joint waveform and beamforming design for RIS-aided ISAC systems. *IEEE Signal Processing Letters*, 30:165–169, 2023.
- [16] P. Jiang, R. Liu, M. Li, W. Wang, and Q. Liu. Joint waveform and beamforming design in RIS-ISAC systems: A model-driven learning approach. *arXiv preprint arXiv:2502.14325*, 2025.
- [17] L. Giroto de Oliveira, Y. Li, B. Geiger, L. Schmalen, T. Zwick, and B. Nuss. On the sensing performance of OFDM-based ISAC under the influence of oscillator phase noise. *arXiv preprint arXiv:2410.13336*, 2024.
- [18] D. Aydoğın, M. Chraïti, and K. K. Tokgöz. On the performance of ISAC over the D-band in a phase-noise aware OFDM system. *arXiv preprint arXiv:2509.08504*, 2025.
- [19] A. Doosti-Aref et al. Performance analysis of OTSM under hardware impairments in millimeter-wave vehicular communication networks. *arXiv preprint arXiv:2309.04161*, 2023.
- [20] Z. He, W. Xu, H. Shen, D. W. K. Ng, Y. C. Eldar, and X. You. Full-duplex communication for ISAC: joint beamforming and power optimization. *IEEE Journal on Selected Areas in Communications*, 41(9):2920–2936, 2023.
- [21] Z. Liu, S. Aditya, H. Li, and B. Clerckx. Joint transmit and receive beamforming design in full-duplex integrated sensing and communications. *IEEE Journal on Selected Areas in Communications*, 41(9):2907–2919, 2023.
- [22] Z. Wei, W. Yuan, S. Li, J. Yuan, and D. W. K. Ng. Off-grid channel estimation with sparse Bayesian learning for OTFS systems. *IEEE Transactions on Wireless Communications*, 21(9):7407–7426, 2022.
- [23] J. Tong, J. Yuan, H. Lin, and J. Xi. Orthogonal delay-Doppler division multiplexing (ODDM) over general physical channels. *IEEE Transactions on Communications*, 2024.
- [24] Q. Zhu, M. Li, R. Liu, and Q. Liu. Joint transceiver beamforming and reflecting design for active RIS-aided ISAC systems. *IEEE Transactions on Vehicular Technology*, 72(7):9636–9640, 2023.
- [25] Q. Zhu, M. Li, R. Liu, et al. Cramér-rao bound optimization for active RIS-empowered ISAC systems. *IEEE Transactions on Wireless Communications*, 23:11723–11736, 2024.
- [26] W. Lyu, S. Yang, Y. Xiu, Y. Li, H. He, C. Yuen, and Z. Zhang. CRB minimization for RIS-aided mmWave integrated sensing and communications. *IEEE Internet of Things Journal*, 11(10):18381–18393, 2024.
- [27] K. Chen, C. Qi, O. A. Dobre, and G. Y. Li. Simultaneous beam training and target sensing in ISAC systems with RIS. *IEEE Transactions on Wireless Communications*, 23(4):2696–2710, 2024.
- [28] A. Bazzi and M. Chafii. On integrated sensing and communication waveforms with tunable PAPR. *IEEE Transactions on Wireless Communications*, 22:7345–7360, 2023.
- [29] R. Liu, M. Li, H. Luo, Q. Liu, and A. L. Swindlehurst. Integrated sensing and communication with reconfigurable intelligent surfaces: opportunities, applications, and future directions. *IEEE Wireless Communications*, 30(1):50–57, 2023.

- [30] Y. Zhu, B. Mao, and N. Kato. Intelligent reflecting surface in 6G vehicular communications: a survey. *IEEE Open Journal of Vehicular Technology*, 3:266–277, 2022.
- [31] K. Meng, Q. Wu, W. Chen, et al. Sensing-assisted communication in vehicular networks with intelligent surface. *IEEE Transactions on Vehicular Technology*, 73:876–893, 2024.
- [32] Z. Chen, P. Chen, Z. Guo, Y. Zhang, and X. Wang. A RIS-based vehicle DOA estimation method with integrated sensing and communication system. *IEEE Transactions on Intelligent Transportation Systems*, 25(6):5554–5566, 2024.
- [33] X. Cheng, D. Duan, S. Gao, and L. Hu. Robust beamforming design for RIS-aided integrated sensing and communication system. *IEEE Transactions on Intelligent Transportation Systems*, 24(6):6227–6243, 2023.
- [34] Ruihang Yang, Li Wei, Dezhi Wang, Bingzhe Du, Chen Zhu, Kai Wang, Zhengyu Zhu, Xiaoming Chen, Zhaohui Yang, and Chongwen Huang. RIS-assisted vehicular ISAC: A robust sensing-centric system design. *Tsinghua Science and Technology*, 2026.
- [35] Y. Chen et al. Reconfigurable intelligent surface (RIS)-aided vehicular networks: their protocols, resource allocation, and performance. *IEEE Vehicular Technology Magazine*, 17(2):26–36, 2022.
- [36] Yonghwi Kim, Sang-Hyun Park, Siyun Yang, Kai-Kit Wong, Linglong Dai, and Chan-Byoung Chae. Near-field challenges in ultra-wideband ISAC: beamforming strategies and system insights. *arXiv preprint arXiv:2508.18810*, 2025.
- [37] S. M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- [38] J. C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341, 1992.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [RAHIOTFSWaveform.zip](#)