OTFS - Predictability in the Delay-Doppler Domain and its Value to Communication and Radar Sensing

Saif Khan Mohammed[1], Ronny Hadani[2], Ananthanarayanan Chockalingam[3], and Robert Calderbank[4], *Fellow, IEEE*
[1]Department of Electrical Engineering, Indian Institute of Technology Delhi, India
[2]Department of Mathematics, University of Texas at Austin, USA
[3] Department of Electrical and Communication Engineering, Indian Institute of Science Bangalore, India
[4]Department of Electrical and Computer Engineering, Duke University, USA

In our first paper [2] we explained why the Zak-OTFS input-output (I/O) relation is predictable and non-fading when the delay and Doppler periods are greater than the effective channel delay and Doppler spreads, a condition which we refer to as the crystallization condition. We argued that a communication system should operate within the crystalline regime.

In this paper, we provide an explicit formula for reconstructing the Zak-OTFS I/O relation from a finite number of received pilot symbols in the delay-Doppler (DD) domain. This formula makes it possible to study predictability of the Zak-OTFS I/O relation for a sampled system that operates under finite duration and bandwidth constraints. We analyze reconstruction accuracy for different choices of the delay and Doppler periods, and of the pulse shaping filter. Reconstruction accuracy is high when the crystallization condition is satisfied, implying that it is possible to learn directly the I/O relation without needing to estimate the underlying channel. This opens up the possibility of a model-free mode of operation, which is especially useful when a traditional model-dependent mode of operation (reliant on channel estimation) is out of reach (for example, when the channel comprises of unresolvable paths, or exhibits a continuous delay-Doppler profile such as in presence of acceleration). Our study clarifies the fundamental origins of predictability by revealing how non-predictability appears as a consequence of aliasing in the DD domain. This perspective leads to a canonical decomposition of the effective DD channel as a sum of predictable and non-predictable components, which we refer to as the crystalline decomposition. Vanishing of the non-predictable component of the channel is equivalent to satisfying the crystallization condition.

Finally, we measure the benefits of predictability in terms of bit error rate (BER) performance. We consider two cases. In the first, we measure performance given perfect knowledge of the I/O relation. We show that performance is optimal when the crystallization condition holds, that performance approaches that of TDM when the Doppler period vanishes, and approaches that of FDM when the delay period vanishes. In the second, we measure performance given imperfect knowledge of the I/O relation, as is the case when it is not possible to learn the underlying channel. We show that model-free operation is successful when the crystallization condition holds, and that performance is only slightly worse than performance given perfect knowledge of the I/O relation.

We also compare the performance of Zak-OTFS with that of a well-studied conventional multi-carrier approximation to Zak-OTFS, which we refer to as MC-OTFS. We show that the I/O relation of MC-OTFS is predictable to a lesser degree than that of Zak-OTFS, and as a result the performance of MC-OTFS is inferior as the Doppler spread increases.

***Index Terms*—OTFS, Delay-Doppler domain, channel predictability, bit error performance, radar sensing.**

I. Introduction

6G presents an opportunity to reflect on the fundamentals of wireless communication, as it becomes more and more difficult to estimate channels, and we encounter Doppler spreads measured in KHz (e.g., 1.3 KHz Doppler at 28 GHz carrier and 50 km/hr speed, and 2.3 KHz Doppler at 5 GHz carrier and 500 km/hr speed) [1]. It is even an opportunity to question the standard model-dependent approach to wireless communication that requires channel estimation.

It is common knowledge that a time-domain (TD) pulse is an ideal waveform for pure delay channels as it is possible to separate reflections according to their range, and, similarly, a frequency domain (FD) pulse is an ideal waveform for pure Doppler channels as it is possible to separate reflections according to their velocity. In part II of this tutorial paper, we explore the proposition that a pulse in the delay-Doppler (DD) domain is an ideal waveform for doubly spread channels comprising of reflections of various ranges and velocities. In Part I [2], we explained that a pulse in the DD domain is a quasi-periodic localized function, and that when viewed in the time domain, is realized as a pulse train modulated by a tone, (hence the name *pulsone*).

In Part I, we described a modulation scheme referred to as Zak-OTFS, which uses the inverse Zak transform [3], [4] to convert information symbols mounted on DD pulses to the time domain for transmission. We emphasized that the Zak-OTFS input-output (I/O) relation is predictable and non-fading when the delay and Doppler periods are greater than the effective channel delay and Doppler spreads, a condition we call the *crystallization condition*. We argued that to achieve robust performance, a communication system should operate within this *crystalline regime*.

In Part I, we described the predictability and non-fading attributes of Zak-OTFS in the context of continuous time and infinite bandwidth. Here in Part II, we study predictability in the context of a sampled communication system with finite duration and bandwidth constraints, and we present a discrete DD domain system model that enables bit error rate (BER) performance evaluation through simulation. We measure predictability through an explicit formula for reconstructing the I/O relation from a finite number of received pilot samples

in the DD domain. The reconstruction accuracy depends on the choice of the delay-Doppler periods and the pulse shaping filters, and accuracy is high when the crystallization condition is satisfied. In the crystalline regime, it is possible to learn the I/O relation without needing to estimate the underlying channel. This opens up the possibility of model-free operation, which can enable communication when traditional model-dependent modes requiring channel estimation are out of reach (for example when the channel comprises of non-resolvable paths, or admits a continuous delay-Doppler profile, as in the presence of acceleration). We now highlight our main contributions.

Origins of non-predictability: We show that non-predictability and fading result from aliasing in the DD domain, and we describe how aliasing occurs when the channel delay spread is greater than the delay period, or the channel Doppler spread is greater than the Doppler period. Fundamental understanding of aliasing leads to the *crystalline decomposition*, which is a canonical decomposition of the effective channel response filter into a predictable component and a non-predictable component. The crystallization condition holds if and only if the non-predictable component vanishes.

Benefits of predictability: Given the I/O response at one DD domain point in a frame, it is possible to predict the I/O response at all other points in the frame. Predictability implies that the received power profile is flat (no fading), that we have engineered a two-dimensional Gaussian channel. We illustrate the practical benefits of predictability by evaluating BER performance as a function of the received SNR. We first suppose that channel estimation is perfect in order to understand the impact of fading. We show that performance is superior in the crystalline regime, that performance approaches that of TDM as the Doppler period shrinks, and that performance approaches that of FDM as the delay period shrinks. We then evaluate performance when we do not have the fine delay-Doppler resolution necessary for accurate channel estimation, and as a consequence, model-dependent approaches fail. We show that in the crystalline regime, model-free operation is successful, and that performance is only slightly worse than performance with perfect knowledge of the effective channel. We also describe how better transmit and receive filters serve to extend the range of reliable operation.

Optimality of Zak-OTFS: Over the past few years several variants of OTFS have been reported in literature [14]. A multicarrier approximation to Zak-OTFS, which we refer to as MC-OTFS, has been the focus of most research attention so far [5], [6]. We show that the I/O relation of MC-OTFS is less predictable than that of Zak-OTFS. As the Doppler spread increases, the BER performance of MC-OTFS is inferior to that of Zak-OTFS. Some recent works on OTFS have started focusing on Zak transform based approach [15], [16], [17], [18]. However, none of these works investigate the subtle aspect of predictability of the I/O relation in the DD domain. In particular, none emphasizes the important fact that only Zak-OTFS where the information bits are encoded as a discrete quasi-periodic function and filtering is applied through twisted convolution admits a predictable I/O relation if the crystallization condition holds. We feel that this assertion is an important theoretical contribution of this paper.

Radar applications: We derive the radar ambiguity function for the Zak pulsone and demonstrate that unambiguous delay-Doppler estimation is possible in the crystalline regime, when the delay period is greater than the delay spread of the radar scene, and the Doppler period is greater than the Doppler spread. We highlight the similarity between the structure of the carrier waveform proposed by Woodward in his influential text [7] (a train of narrow TD Gaussian pulses modulated with a broad Gaussian envelope), and the Zak pulsone carrier waveform (a train of narrow impulses modulated by a sinusoid). Note that in radar applications, a signal is modulated onto a carrier waveform (in our case, the Zak pulsone), and that the ambiguity function of the concatenated system depends on both signal and carrier (see [8] and [9]).

The paper is organized as follows. Section II explains how unpredictability in communications results from aliasing in the DD domain, which occurs when one of the channel spreads is greater than the corresponding pulsone period. This geometric perspective leads to a canonical decomposition of the effective channel response filter into a predictable component and a nonpredictable component (the crystalline decomposition). Section III addresses finite bandwidth and duration constraints, expressing the I/O relation for TDM, FDM, and Zak-OTFS in matrix-vector form. Section IV uses BER simulations when the channel is perfectly known, to illustrate that Zak-OTFS is non-fading in the crystalline regime. Section V uses BER simulations when the channel is not known, to illustrate that Zak-OTFS is predictable, and that model-free operation is possible in the crystalline regime. Section VI uses BER simulations to illustrate that Zak-OTFS is more predictable than MC-OTFS (a widely studied multicarrier approximation) making the case that model-free operation is more possible. Section VII derives the radar ambiguity function of the Zak-OTFS carrier waveform, illustrating that unambiguous delay-Doppler estimation is possible in the crystalline regime. Section VIII presents conclusions.

II. The Zak-OTFS Input-Output (I/O) Relation and Predictability in the Crystalline Regime

A. The Zak-OTFS I/O Relation for a Sampled System

The Zak-OTFS I/O relation (in the absence of AWGN) is presented in Section VI-C of Part I [2] (see also Table III in Part I). In Zak-OTFS, the frame duration is roughly an integer multiple N of the delay period τ_p, and the bandwidth is roughly an integer multiple M of the Doppler period $\nu_p = 1/\tau_p$. The information symbols are arranged as a 2-D array $x[k,l]$, $k = 0, 1, \cdots, M-1$, $l = 0, 1, \cdots, N-1$ and are encoded as a discrete DD domain information signal which is given by

$$x_{\text{dd}}[k+nM, l+mN] \triangleq x[k,l]\, e^{j2\pi n \frac{l}{N}}, \quad m, n \in \mathbb{Z}. \tag{1}$$

Being discrete DD domain signal means that $x_{\text{dd}}[k,l]$ is a quasi-periodic function on the information grid with period M along the delay axis and period N along the Doppler axis, i.e.,

$$x_{\text{dd}}[k+nM, l+mN] = x_{\text{dd}}[k,l]\, e^{j2\pi n \frac{l}{N}}, \tag{2}$$

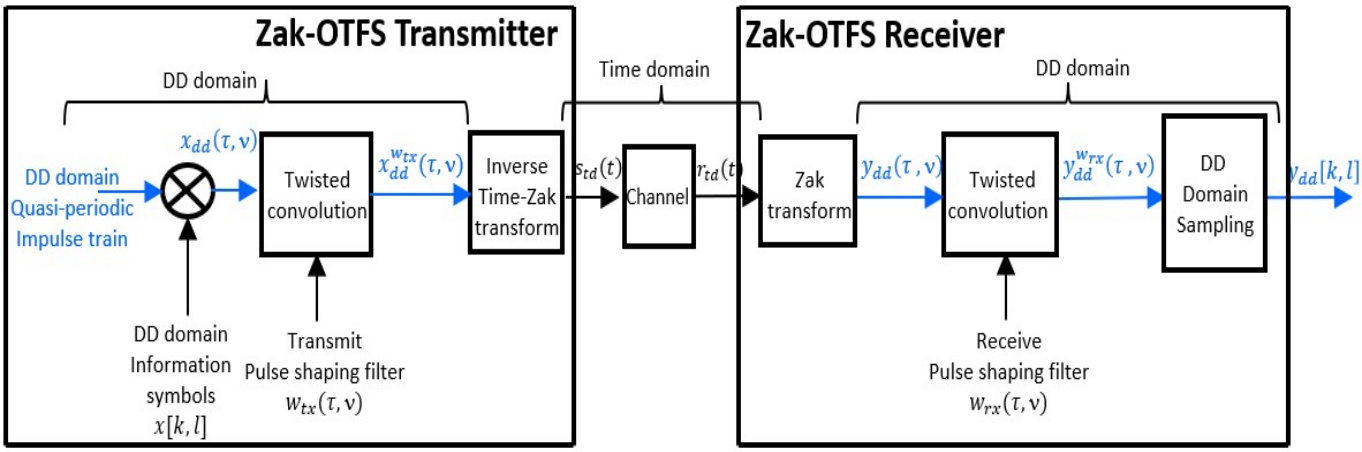

Fig. 1: Signal processing in Zak-OTFS.

for any $n, m \in \mathbb{Z}$. The discrete DD domain signal is then lifted to a continuous DD domain signal by means of

$$x_{\text{dd}}(\tau,\nu) = \sum_{k,l\in\mathbb{Z}} x_{\text{dd}}[k,l]\,\delta\left(\tau - k\frac{\tau_p}{M}\right)\delta\left(\nu - l\frac{\nu_p}{N}\right). \quad (3)$$

The complete transceiver signal processing is depicted in Fig. 1 (see also Table-III in Part I). As shown in Fig. 1, the DD domain transmit signal $x_{\text{dd}}^{w_{tx}}(\tau,\nu)$ is the twisted convolution of $x_{\text{dd}}(\tau,\nu)$ with a transmit pulse $w_{tx}(\tau,\nu)$, i.e., $x_{\text{dd}}^{w_{tx}}(\tau,\nu) = w_{tx}(\tau,\nu) *_\sigma x_{\text{dd}}(\tau,\nu)$, where $*_\sigma$ denotes the twisted convolution operation. We then apply the inverse Zak transform to obtain the TD realization $s_{\text{td}}(t)$. At the receiver, the received TD signal $r_{\text{td}}(t)$ is converted to its DD domain representation $y_{\text{dd}}(\tau,\nu)$ via the Zak transform. The channel acts on $x_{\text{dd}}^{w_{tx}}(\tau,\nu)$ by twisted convolution, so that $y_{\text{dd}}(\tau,\nu) = h(\tau,\nu) *_\sigma x_{\text{dd}}^{w_{tx}}(\tau,\nu)$. After twisted convolution of $y_{\text{dd}}(\tau,\nu)$ with a receive DD pulse $w_{rx}(\tau,\nu)$ we obtain

$$\begin{aligned} y_{\text{dd}}^{w_{rx}}(\tau,\nu) &= w_{rx}(\tau,\nu) *_\sigma y_{\text{dd}}(\tau,\nu) \\ &= w_{rx}(\tau,\nu) *_\sigma \left(h(\tau,\nu) *_\sigma x_{\text{dd}}^{w_{tx}}(\tau,\nu)\right) \\ &= w_{rx}(\tau,\nu) *_\sigma \left(h(\tau,\nu) *_\sigma \left[w_{tx}(\tau,\nu) *_\sigma x_{\text{dd}}(\tau,\nu)\right]\right) \\ &= \underbrace{\left(w_{rx}(\tau,\nu) *_\sigma h(\tau,\nu) *_\sigma w_{tx}(\tau,\nu)\right)}_{h_{\text{dd}}(\tau,\nu)} *_\sigma x_{\text{dd}}(\tau,\nu), \end{aligned} \quad (4)$$

where the last step follows from the associativity of the twisted convolution operation. Equation (4) gives the Zak-OTFS I/O relation in the continuous DD domain. Simply put, the I/O relation states that the output $y_{\text{dd}}^{w_{rx}}(\tau,\nu)$ is the twisted convolution of the input $x_{\text{dd}}(\tau,\nu)$ with the effective continuous DD channel filter $h_{\text{dd}}(\tau,\nu)$. Finally, we sample this continuous output signal along the information grid consisting of integer multiples of τ_p/M along the delay axis and integer multiples of ν_p/N along the Doppler axis, to obtain a discrete DD domain output signal

$$\begin{aligned} y_{\text{dd}}[k,l] &= y_{\text{dd}}^{w_{rx}}\left(\tau = k\frac{\tau_p}{M}, \nu = l\frac{\nu_p}{N}\right) \\ &= \sum_{k',l'\in\mathbb{Z}} h_{\text{dd}}[k-k', l-l']\, x_{\text{dd}}[k',l']\, e^{j2\pi\frac{(l-l')}{N}\frac{k'}{M}}, \end{aligned} \quad (5)$$

where $h_{\text{dd}}[k,l]$ is the discrete effective channel filter, given by sampling the continuous effective channel filter, i.e.,

$$h_{\text{dd}}[k,l] \triangleq h_{\text{dd}}(\tau,\nu)\Big|_{\left(\tau=\frac{k\tau_p}{M},\,\nu=\frac{l\nu_p}{N}\right)}. \quad (6)$$

Equation (5) constitutes the canonical form of the Zak-OTFS I/O relation. It expresses the output signal as a discrete twisted convolution[1] of the discrete effective channel filter and the input signal, i.e.,

$$\begin{aligned} y_{\text{dd}}[k,l] &= \sum_{k',l'\in\mathbb{Z}} h_{\text{dd}}[k',l']\, x_{\text{dd}}[k-k', l-l']\, e^{j2\pi\frac{(k-k')}{M}\frac{l'}{N}} \\ &= h_{\text{dd}}[k,l] *_\sigma x_{\text{dd}}[k,l]. \end{aligned} \quad (7)$$

We remark that the R.H.S of (5) is a weighted double sum of terms, where the (k',l')-th term is given by

$$h_{\text{dd}}[k,l] *_\sigma \left[\delta[k-k']\,\delta[l-l']\right],\ k,l\in\mathbb{Z}, \quad (8)$$

multiplied by the weight $x_{\text{dd}}[k',l']$.

B. Predictability of the Zak-OTFS I/O relation in the Crystalline Regime

In this section, we show that the Zak-OTFS I/O relation is predictable in the crystalline regime. Specifically, we show that when the crystallization condition holds, the channel response to a *green* pilot located at $(k^{(g)}, l^{(g)})$ can be accurately estimated from the channel response to a *blue* pilot located at $(k^{(b)}, l^{(b)})$.

By definition, the blue pilot is a discrete DD domain impulse signal, given by

$$x^{(b)}[k,l] = \sum_{m,n\in\mathbb{Z}} e^{j2\pi\frac{nl^{(b)}}{N}}\delta[k-(k^{(b)}+nM)]\delta[l-(l^{(b)}+mN)]. \quad (9)$$

[1] $u[k,l] *_\sigma v[k,l] = \sum_{k',l'\in\mathbb{Z}} u[k',l']v[k-k',l-l']e^{j2\pi\frac{k-k'}{M}\frac{l'}{N}}$

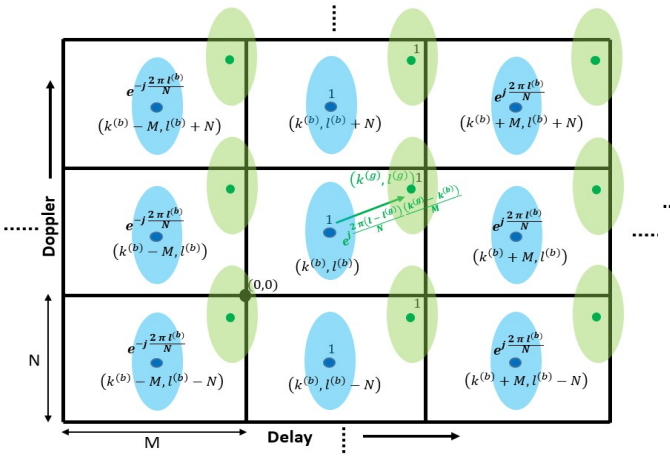

Fig. 2: Prediction in the crystalline regime. In the crystalline regime, the channel response to the green DD pilot signal can be predicted from the channel response to the blue DD pilot signal.

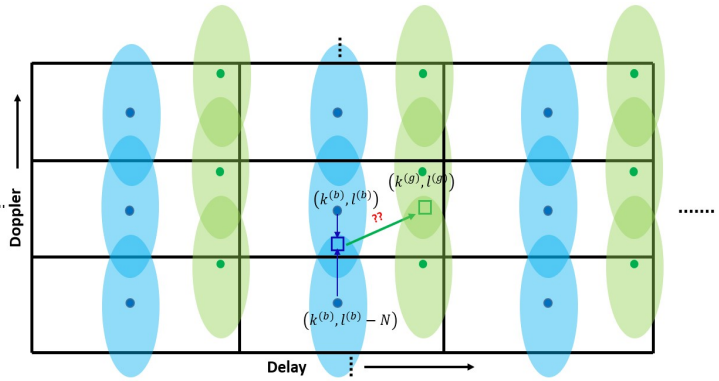

Fig. 3: Non-crystalline regime: Partially predictable Zak-OTFS I/O relation. The response to the green DD pilot signal can be predicted from the response to the blue DD pilot signal only for the part of the response which is free from any overlap/aliasing/self-interaction.

Similarly, the green pilot is a discrete DD domain impulse signal, given by

$$x^{(g)}[k,l] = \sum_{m,n\in\mathbb{Z}} e^{j2\pi\frac{nl^{(g)}}{N}} \delta[k-(k^{(g)}+nM)]\delta[l-(l^{(g)}+mN)] \quad (10)$$

The channel response to the blue pilot is given by

$$\begin{aligned} y^{(b)}[k,l] &= h_{\text{dd}}[k,l] \ *_\sigma \ x^{(b)}[k,l] \\ &= \sum_{n,m\in\mathbb{Z}} \Big[h_{\text{dd}}[k-(k^{(b)}+nM), l-(l^{(b)}+mN)] \\ &\qquad e^{j2\pi\frac{nl^{(b)}}{N}} \ e^{j2\pi\frac{(l-l^{(b)}-mN)}{N}\frac{(k^{(b)}+nM)}{M}} \Big]. \end{aligned} \quad (11)$$

We see from (11) that the total response is a sum of local responses to its constituent impulses where the response to the (n,m)-th impulse is given by

$$\begin{aligned} y^{(b)}_{\text{n,m}}[k,l] &= h_{\text{dd}}[k-(k^{(b)}+nM), l-(l^{(b)}+mN)] \\ &\qquad e^{j2\pi\frac{nl^{(b)}}{N}} \ e^{j2\pi\frac{(l-l^{(b)}-mN)}{N}\frac{(k^{(b)}+nM)}{M}}. \end{aligned} \quad (12)$$

Observe that the (n,m)-th response coincides up to multiplicative phases with the effective channel filter shifted by $(k^{(b)}+nM, l^{(b)}+mN)$.

Similarly, the response to the green pilot is a sum of local responses, where the (n,m)-th response is given by

$$\begin{aligned} y^{(g)}_{\text{n,m}}[k,l] &= \Big[h_{\text{dd}}[k-(k^{(g)}+nM), l-(l^{(g)}+mN)] \\ &\qquad e^{j2\pi\frac{nl^{(g)}}{N}} \ e^{j2\pi\frac{(l-l^{(g)}-mN)}{N}\frac{(k^{(g)}+nM)}{M}} \Big]. \end{aligned} \quad (13)$$

In Fig. 2, the response components to the blue (green) pilot are depicted as blue (green) ellipses. Observe that the (n,m)-th green response can always be calculated from the (n,m)-th blue response according to the following rule

$$\begin{aligned} y^{(g)}_{\text{n,m}}[k,l] &= y^{(b)}_{\text{n,m}}[k-(k^{(g)}-k^{(b)}), l-(l^{(g)}-l^{(b)})] \\ &\qquad e^{j2\pi n\frac{(l^{(g)}-l^{(b)})}{N}} \ e^{j2\pi\frac{(l-l^{(g)}-mN)}{N}\frac{(k^{(g)}-k^{(b)})}{M}}. \end{aligned} \quad (14)$$

where we use (12) and (13).[2] The problem is that the individual local responses are super-imposed at the receiver and, in general, cannot be separated from one another. However, when the crystallization conditions hold, the blue (green) ellipses do not overlap with each other which means that the local responses do not interact. As a consequence, in this situation. the total green response can be predicted from the total blue response using the local formula (14).

An important consequence of predictability is that the energy profile of the local responses is flat, that is, independent of the position of the pilot, due to the fact that

$$\left| y^{(g)}_{\text{n,m}}[k,l] \right| = \left| y^{(b)}_{\text{n,m}}[k-(k^{(g)}-k^{(b)}), l-(l^{(g)}-l^{(b)})] \right|, \quad (15)$$

In other words, in the crystalline regime, the Zak-OTFS I/O relation is *non-fading*.

C. Non-predictability of the Zak-OTFS I/O relation in the Non-crystalline Regime

When the crystallization conditions fail to hold the local channel responses interact (the ellipses overlap with one another), a phenomenon we refer to as DD domain aliasing. We now explain how non-predictability arises from DD domain aliasing.

We illustrate the aliasing phenomenon through an example. We again look to predict the response to a green pilot from the response to a blue pilot. However, we now assume that the Doppler spread of the effective channel filter is greater than the Doppler period. Under this assumption, local responses interact. Specifically, in our example the (n,m)-th response interacts with the $(n,m-1)$-th response. The interaction is depicted in Fig. 3 as an overlap between ellipses.

We consider a generic point (k,l) residing in the overlap of the $(0,0)$ and $(0,-1)$ blue ellipses, depicted in Fig. 3 as a square with a blue border. The response received at this point is

[2]To obtain the (n,m)-th response component to the green pilot from that of the blue pilot, we first align supports by translating the blue ellipses by $(k^{(g)}-k^{(b)})$ in delay, and by $(l^{(g)}-l^{(b)})$ in Doppler. We then multiply by a deterministic factor that is independent of the channel.

the super-position of two local responses $y^{(b)}_{0,0}[k,l]+y^{(b)}_{0,\text{-}1}[k,l]$, which is equal to

$$h_{\text{dd}}[\Delta k,\Delta l]e^{j2\pi\frac{\Delta l}{N}\frac{k^{(b)}}{M}}+h_{\text{dd}}[\Delta k,\Delta l+N]e^{j2\pi\frac{\Delta l+N}{N}\frac{k^{(b)}}{M}}, \quad (16)$$

where $\Delta k = k-k^{(b)}$ and $\Delta l = l-l^{(b)}$. We consider the parallel point (k',l') residing in the overlap of the $(0,0)$ and $(0,-1)$ green ellipses, depicted in Fig. 3 as a square with a green border. Here, $k' = k+(k^{(g)}-k^{(b)})$ and $l' = l+(l^{(g)}-l^{(b)})$. The response at this point is a super-position of two local responses $y^{(g)}_{0,0}[k',l']+y^{(g)}_{0,\text{-}1}[k',l']$, which is explicitly equal to

$$h_{\text{dd}}[\Delta k,\Delta l]e^{j2\pi\frac{\Delta l}{N}\frac{k^{(g)}}{M}}+h_{\text{dd}}[\Delta k,\Delta l+N]e^{j2\pi\frac{\Delta l+N}{N}\frac{k^{(g)}}{M}}. \quad (17)$$

We see that unless both the individual local terms in (16) are separately known, the green response in (17) cannot be predicted from the total blue response.

A consequence of non-predictability is that the energy profile of the received response depends on the DD location of the pilot signal. In our example, the energy profile of the received response to the blue pilot will, in general, be different than that for the green pilot. This is because the linear combination (16) can be different from (17), depending on the values of the phase coefficients in each case.

D. The Crystalline Decomposition

The previous discussion reveals that the channel response is unpredictable at points inside the overlap region between interacting ellipses. Outside this region, the response is in fact predictable. This observation suggests a decomposition of the effective channel filter into a predictable and a non-predictable component.

To see this, we consider a pilot at the origin. The channel response to this pilot is given by

$$h_{\text{dd}}[k,l]+\sum_{n,m\in\mathbb{Z},(n,m)\neq(0,0)} h_{\text{dd}}[k-nM,l-mN]\,e^{j2\pi\frac{nl}{N}} \quad (18)$$

where the first term is the local response to the impulse at the origin and the remaining terms are local responses to other impulses. The main observation is that if (k,l) satisfies $h_{\text{dd}}[k,l]\neq 0$ and $h_{\text{dd}}[k-nM,l-mN]\neq 0$ for some $(n,m)\neq(0,0)$, then the response at (k,l) is a superposition of at least two local responses - the local response to the $(0,0)$-th impulse and the local response to the (n,m)-th impulse, and, consequently, the effective channel filter tap at (k,l) cannot be estimated from the received response. We refer to such taps as *non-predictable* taps. Alternatively, it is easy to see that all other taps in the support can be estimated from the channel response. We refer to such taps as *predictable* taps. We denote the support set of all predictable taps by $\mathcal{P}$ and the complement support set of all non-predictable taps by $\mathcal{P}^c$.

Recall that given a set $\mathcal{S}$, the indicator function $\chi_{\mathcal{S}}[k,l]$ takes the value 1 for $(k,l)\in\mathcal{S}$, and the value 0 otherwise. We define the *crystalline decomposition* of the effective channel filter to be

$$h_{\text{dd}}[k,l] \;=\; \chi_{\mathcal{P}}[k,l]\,h_{\text{dd}}[k,l] \;+\; \chi_{\mathcal{P}^c}[k,l]\,h_{\text{dd}}[k,l], \quad (19)$$

where the first term is the predictable component and the second term is the non-predictable component. The crystallization condition corresponds to the special case where the non-predictable component vanishes, that is, $\mathcal{P}^c=\emptyset$.

E. Error in prediction of the Zak-OTFS I/O relation

In this section we analyze the prediction accuracy for different choices of the delay-Doppler period and shaping filter. The analysis is carried through simulations of a simple two-path channel, where, the first path has a delay of $\tau_1 = 0\,\mu s$ and a Doppler shift of $\nu_1 = 815$ Hz and the second path has a delay of $\tau_2 = 5\mu s$ and a Doppler shift of $\nu_2 = -815$ Hz. The normalized channel gains for the two paths are $h_1 = h_2 = 1/\sqrt{2}$. The system bandwidth is $B = 0.96$ MHz and the frame duration is $T = 1.6$ ms.

We consider two choices for the transmit/receive shaping filters. The first choice is a sinc filter

$$w_{rx}(\tau,\nu) = w_{tx}(\tau,\nu) = \sqrt{BT}\,sinc(B\tau)\,sinc(\nu T). \quad (20)$$

Observe that the slow decay of the sinc function amplifies DD domain aliasing, thereby, reducing predictability.

The second choice is a root raised cosine (RRC) filter

$$w_{rx}(\tau,\nu) = w_{tx}(\tau,\nu) = \sqrt{BT}\,rrc_{\beta_\tau}(B\tau)\,rrc_{\beta_\nu}(T\nu), \quad (21)$$

where the function $rrc_\beta(\cdot)$ for $0\leq\beta\leq 1$ is given by [11]

$$rrc_\beta(x) = \frac{\sin(\pi x(1-\beta))+4\beta x\cos(\pi x(1+\beta))}{\pi x\left(1-4\beta x\right)^2}. \quad (22)$$

Observe that the RRC function has faster decay than the sinc function, causing less DD domain aliasing, thereby increasing predictability. However, reduced aliasing comes at the cost of an increase in frame duration and bandwidth. In the simulation we use $\beta_\tau = 0.1$ and $\beta_\nu = 0.2$, corresponding to a $10\,\%$ increase in bandwidth, and a $20\,\%$ increase in duration.

We normalize the energy of the transmit/receive filters by setting

$$\iint |w_{tx}(\tau,\nu)|^2\,d\tau d\nu = \iint |w_{rx}(\tau,\nu)|^2\,d\tau d\nu = 1. \quad (23)$$

We now use a pilot signal at $(k^{(b)},l^{(b)}) = \left(\frac{M}{2},\frac{N}{2}\right)$ to estimate the effective channel filter $h_{\text{dd}}[k,l]$. Recall, that in the crystalline regime, local responses do not interact, therefore, the total response in the fundamental period coincides with the $(0,0)$-th local response $y^{(b)}_{0,0}[k,l]$ which is given by

$$y^{(b)}[k,l] = h_{\text{dd}}[k-M/2,l-N/2]\,e^{j\pi\frac{(l-N/2)}{N}}, \quad (24)$$

for $0\leq k<M$ and $0\leq l<N$. From this we conclude

$$h_{\text{dd}}[k,l] = y^{(b)}[k+M/2,l+N/2]\,e^{-j\pi\frac{l}{N}}, \quad (25)$$

for $-\frac{M}{2}\leq k<\frac{M}{2}$ and $-\frac{N}{2}\leq l<\frac{N}{2}$. We set the channel filter estimate to be

$$\widehat{h}_{\text{dd}}[k,l] \;=\; \begin{cases} y^{(b)}\left[k+\frac{M}{2},l+\frac{N}{2}\right]\,e^{-j\pi\frac{l}{N}} & ,-\frac{M}{2}\leq k<\frac{M}{2} \\ & -\frac{N}{2}\leq l<\frac{N}{2} \\ 0 & ,\text{otherwise.} \end{cases} \quad (26)$$

Using (26), we can predict the effective channel response to any pilot. The predicted channel response to a pilot located at $(k^{(g)},l^{(g)})$ is given by

$$\widehat{y}^{(g)}[k,l] = \widehat{h}_{\text{dd}}[k,l] \; *_\sigma \; x^{(g)}[k,l]. \tag{27}$$

The predicted response should be compared to the true channel response

$$y^{(g)}[k,l] = h_{\text{dd}}[k,l] \; *_\sigma \; x^{(g)}[k,l]. \tag{28}$$

The relative prediction error is given by

$$E(k^{(g)},l^{(g)}) \quad \triangleq \quad \frac{\sum\limits_{k=0}^{M-1}\sum\limits_{l=0}^{N-1}\left|\widehat{y}^{(g)}[k,l]-y^{(g)}[k,l]\right|^2}{\sum\limits_{k=0}^{M-1}\sum\limits_{l=0}^{N-1}\left|y^{(g)}[k,l]\right|^2}. \tag{29}$$

Figs. 4 and 5 depict the relative prediction error as a two-dimensional heat-map for three different points on the period curve: the points $\nu_p = 1.25$ KHz, representing a TDM approximation, $\nu_p = 30$ KHz representing a point in the crystalline regime and $\nu_p = 240$ KHz, representing an FDM approximation. Fig. 4 assumes sinc transmit/receive shaping filters, whereas Fig. 5 assumes root raised cosine transmit/receive filters.

When $\nu_p = 1.25$ KHz ($\tau_p = 800\,\mu s$), the Doppler period is smaller than the channel Doppler spread (1.63 KHz), causing aliasing along the Doppler dimension which, in turns, creates non-predictability and fading along delay. In this situation the relative prediction error is minimized at the pilot location and it increases rapidly towards 0 dB as we move away along the delay axis.

When $\nu_p = 240$ KHz, ($\tau_p = 4.16\,\mu s$), the delay period is smaller than the channel delay spread ($5\,\mu s$) causing aliasing along along the delay dimension, which, in turns, creates non-predictability and fading along Doppler. In this situation the relative prediction error is minimized at the pilot location, and it now increases rapidly towards 0 dB as we move away from the pilot location along the Doppler axis.

Finally, when $\nu_p = 30$ KHz, ($\tau_p = 33.3\,\mu s$), the crystallization condition holds and predictability is maintained. In this situation, the relative prediction error is small.

When we compare Fig. 4 to Fig. 5, we see that the choice of transmit/receive shaping filter makes a significant difference. Using root raised cosine filters lead to a relative prediction error of roughly -50 dB, whereas using sinc filters lead to a relative prediction error of roughly -20 dB.

III. Matrix-vector description of the I/O relation

Recall that the channel equation relating the received and the transmit TD signals is given by [10]

$$r_{\text{td}}(t) = \iint h(\tau,\nu) s_{\text{td}}(t-\tau)\, e^{j2\pi\nu(t-\tau)}\, d\tau\, d\nu \; + \; n_{\text{td}}(t), \tag{30}$$

where $h(\tau,\nu)$ is the delay-Doppler channel representation/spreading function and $n_{\text{td}}(t)$ is the AWGN noise term. In this section we study the I/O relation induced by (30) for Zak-OTFS, TDM and FDM and establish a matrix formulation in each case.

A. Zak-OTFS

Direct calculation reveals that the DD domain I/O relation induced by (30) is given by

$$y_{\text{dd}}[k,l] = \sum_{k',l'\in\mathbb{Z}} h_{\text{dd}}[k-k',l-l']\, x_{\text{dd}}[k',l']\, e^{j2\pi\frac{(l-l')}{N}\frac{k'}{M}} + n_{\text{dd}}[k,l], \tag{31}$$

where $h_{\text{dd}}[k,l]$ is the effective DD domain channel filter and $n_{\text{dd}}[k,l]$ is the discrete noise term obtained by sampling

$$n_{\text{dd}}(\tau,\nu) = w_{rx}(\tau,\nu) \; *_\sigma \; \mathcal{Z}_t(n_{\text{td}}(t)).$$

Here $\mathcal{Z}_t(\cdot)$ denotes the time-Zak transform (see Part I). Both sequences $x_{\text{dd}}[k,l]$ and $y_{\text{dd}}[k,l]$ are quasi-periodic. Hence, can be be reconstructed from samples within the fundamental period ($0 \le k < M, 0 \le l < N$). Fig. 6 illustrates how to reconstruct the local response to a DD domain pilot supported over the union of regions I, II, III and IV, from the received samples supported over the union of regions I', II', III and IV' within the fundamental period.

The MN received samples in the fundamental period constitute a sufficient statistic for the transmitted information symbols $x[k,l]$. Hence, the I/O relation (31) can be reduced to a finite matrix form relating the $MN \times 1$ vector of received samples $y_{\text{dd}}[k,l]$ to the $MN\times 1$ vector of transmitted symbols $x[k,l]$, where on both sides $k = 0,1,\cdots,M-1$ and $l = 0,1,\cdots,N-1$. In more detail, define the $MN\times 1$ vectors

$$\begin{aligned}\left(\mathbf{y}_{\text{dd}}\right)_{kM+l+1} &= y_{\text{dd}}[k,l]\\ \left(\mathbf{x}_{\text{dd}}\right)_{kM+l+1} &= x_{\text{dd}}[k,l]\\ \left(\mathbf{n}_{\text{dd}}\right)_{kM+l+1} &= n_{\text{dd}}[k,l],\end{aligned} \tag{32}$$

In addition, define the $MN \times MN$ matrix

$$\left(\mathbf{H}_{\text{dd}}\right)_{k'N+l'+1,kN+l+1} = H_{\text{dd}}[k'N+l', kN+l] \tag{33}$$

where the right hand side is given by (34) (see top of page 9). The matrix formulation of the DD domain I/O relation (31) is given by

$$\mathbf{y}_{\text{dd}} \quad = \quad \mathbf{H}_{\text{dd}}\mathbf{x}_{\text{dd}} \; + \; \mathbf{n}_{\text{dd}}. \tag{35}$$

B. TDM

Direct calculation reveals that the time domain I/O relation induced from (30), is given by (see Equation (24) in Part I)

$$y_{\text{td}}[k] \quad = \quad \sum_{k'\in\mathbb{Z}} x_{\text{td}}[k']\, h_{\text{td}}[k-k'\,;\,k'] \; + \; n_{\text{td}}[k] \tag{36}$$

where $h_{\text{td}}[n\,;\,k]$ is the effective TD channel filter and the noise term $n_{\text{td}}[k]$ is given by sampling at $t = k/B$ the continuous function

$$w_{rx}(t) \; \star \; n_{\text{td}}(t),$$

where $w_{rx}(t)$ is the matched filter at the receiver and $\star$ is linear convolution. Equation (36) can be expressed in a matrix form as follows. Arrange the transmitted information symbols as a BT column vector

$$(\mathbf{x}_{\text{td}})_{k+1} = x_{\text{td}}[k] = x[k], \tag{37}$$

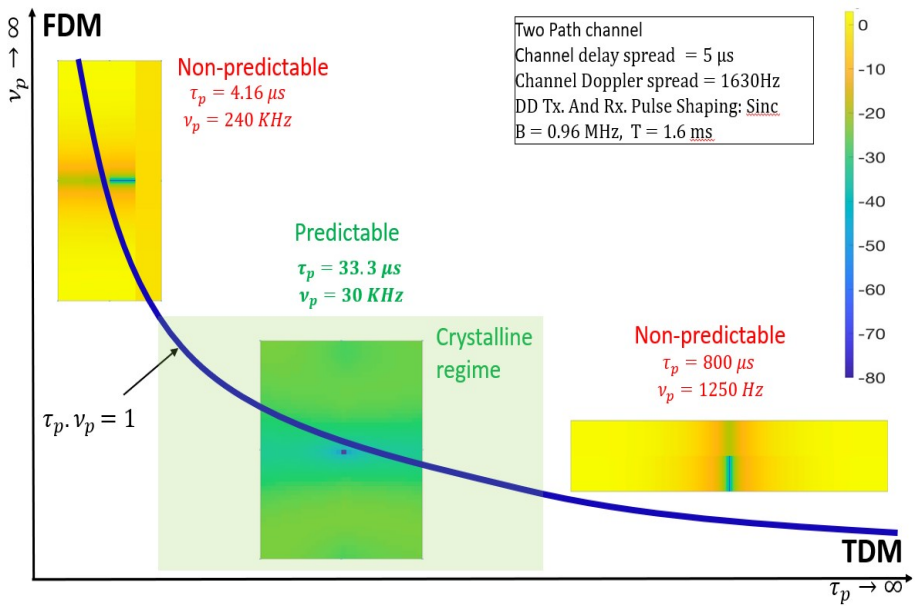

Fig. 4: Heatmap showing Relative Prediction Error (RPE), in dB, as a function of delay (horizontal axis), and Doppler (vertical axis) with sinc pulse shaping filters. RPE is significantly smaller in the crystalline regime when compared to that in the non-crystalline regime.

for $k = 0, 1, \cdots, BT-1$. Arrange the received TD samples as a $BT + K_1 + K_2$ column vector

$$(\mathbf{y}_{\text{td}})_{k+1+K_1} = y_{\text{td}}[k], \tag{38}$$

for $k = -K_1, \cdots, BT-1+K_2$, where the constants $K_1, K_2 \in \mathbb{Z}$ are determined by the channel delay spread. Similarly, arrange the received sampled noise as a $BT+K_1+K_2$ vector

$$(\mathbf{n}_{\text{td}})_{k+1+K_1} = n_{\text{td}}[k], \tag{39}$$

for $k = -K_1, \cdots, BT-1+K_2$. Finally, arrange the effective channel filter as a $(BT + K_1 + K_2) \times BT$ matrix

$$(\mathbf{H}_{\text{td}})_{k'+1,k+1} \;=\; h_{\text{td}}[k' - k - K_1\,;\,k]. \tag{40}$$

Putting everything together, we obtain the matrix relation

$$\mathbf{y}_{\text{td}} \;=\; \mathbf{H}_{\text{td}}\,\mathbf{x}_{\text{td}} \,+\, \mathbf{n}_{\text{td}}. \tag{41}$$

C. FDM

Direct calculation reveals that the frequency domain I/O relation induced from (30), is given by (see Equation (25) in Part I)

$$y_{\text{fd}}[k] \;=\; \sum_{k' \in \mathbb{Z}} x_{\text{fd}}[k']\, h_{\text{fd}}[k - k'\,;\,k'] \,+\, n_{\text{fd}}[k] \tag{42}$$

where $h_{\text{fd}}[n\,;\,k]$ is the effective FD channel filter and the noise term $n_{\text{fd}}[k]$ is obtained by sampling at $f = K/T$ the continuous function $w_{rx}(f) \star n_{\text{fd}}(f)$ where

$$n_{\text{fd}}(f) \;=\; \int n_{\text{td}}(t)\, e^{-j2\pi ft}\, dt.$$

Arrange the transmitted information symbols as a BT column vector

$$(\mathbf{x}_{\text{fd}})_{k+1} = x_{\text{fd}}[k] = x[k], \tag{43}$$

for $k = 0, 1, \cdots, BT-1$. Arrange the received samples as a $BT + L_1 + L_2$ column vector

$$(\mathbf{y}_{\text{fd}})_{k+1+L_1} = y_{\text{fd}}[k], \tag{44}$$

for $k = -L_1, \cdots, BT-1+L_2$, where the constants $L_1, L_2 \in \mathbb{Z}$ are determined by the channel Doppler spread. Similarly, arrange the received noise samples as a $BT+L_1+L_2$ column vector

$$(\mathbf{n}_{\text{fd}})_{k+1+L_1} = n_{\text{fd}}[k], \tag{45}$$

for $k = -L_1, \cdots, BT-1+L_2$. Finally, arrange the FD channel filter as a $(BT + L_1 + L_2) \times BT$ matrix

$$(\mathbf{H}_{\text{fd}})_{k'+1,k+1} = h_{\text{fd}}[k' - k - L_1\,;\,k], \tag{46}$$

for $k' = 0, 1, \cdots, BT+L_1+L_2-1$ and $k = 0, 1, \cdots, BT-1$. Putting everything together, we obtain the matrix relation

$$\mathbf{y}_{\text{fd}} \;=\; \mathbf{H}_{\text{fd}}\,\mathbf{x}_{\text{fd}} \,+\, \mathbf{n}_{\text{fd}}. \tag{47}$$

IV. Impact of Fading in the Crystalline Regime

In this Section, we compare uncoded BER performance of Zak-OTFS, TDM and FDM for the Veh-A channel model [**?**]. Since we are focusing on the impact of fading, we suppose that the input-output (I/O) relation is perfectly known. We study the performance of Zak-OTFS as we move along the hyperbola $\tau_p \cdot \nu_p = 1$ by choosing different Doppler periods ν_p. We shall demonstrate that the performance of Zak-OTFS is superior in the crystalline regime, that performance approaches TDM as the Doppler period shrinks, and that performance approaches FDM as the delay period shrinks.

Our Veh-A channel consists of six channel paths, and the delay-Doppler spreading function is given by

$$h(\tau, \nu) \;=\; \sum_{i=1}^{6} h_i\, \delta(\tau - \tau_i)\, \delta(\nu - \nu_i). \tag{48}$$

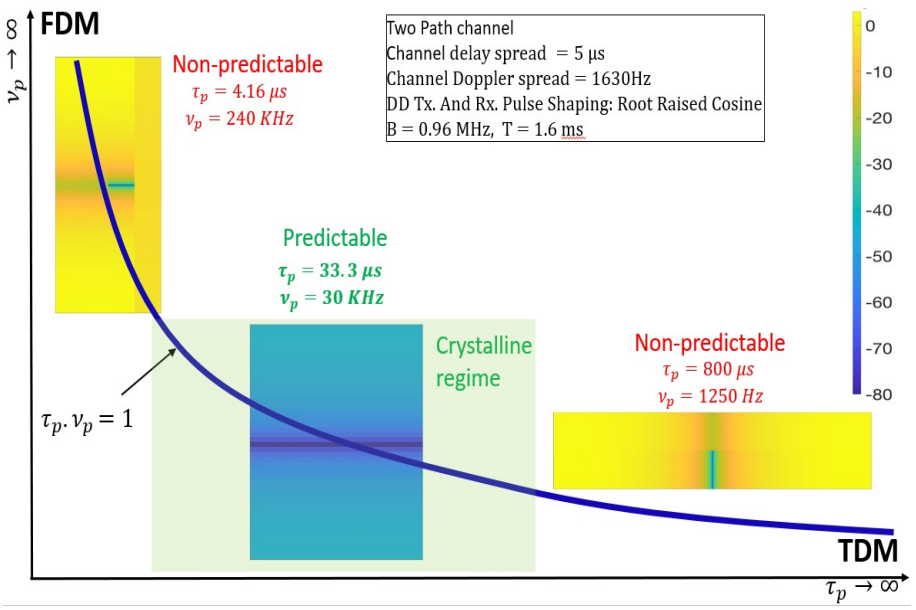

Fig. 5: Heatmap showing Relative Prediction Error (RPE), in dB, as a function of delay (horizontal axis), and Doppler (vertical axis) with root raised cosine (RRC) pulse shaping filters. In the crystalline regime, RRC pulse results in significantly smaller RPE when compared to sinc pulses, although at the cost of a higher OTFS frame duration and bandwidth.

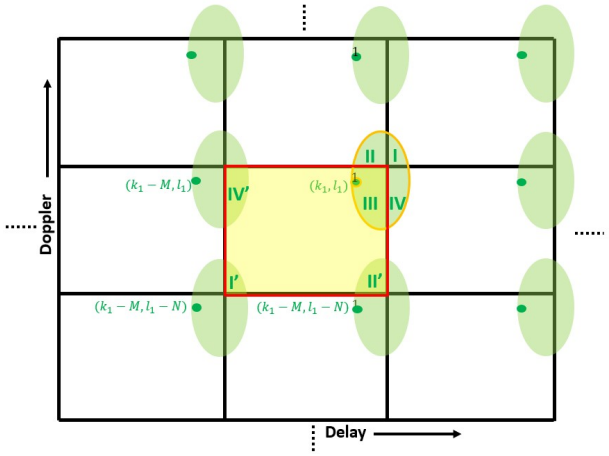

Fig. 6: Exploiting Quasi-periodicity: Generating complete response from response received in the fundamental DD period.

where h_i, τ_i, and ν_i respectively denote the gain, delay, and Doppler shift of the i-th channel path. Table-I lists the power-delay profile for the six channel paths. The maximum Doppler shift is $\nu_{max} = 815$ Hz, the Doppler spread is 1.63 KHz, and the delay spread is $\tau_{max} \triangleq \max_i \tau_i - \min_i \tau_i = 2.5\,\mu s$. The Doppler shift of the i-th path is modeled as $\nu_i = \nu_{max}\cos(\theta_i)$, where the variables $\theta_i, i = 1, 2, \cdots, 6$ are independent and distributed uniformly in the interval $[0\,,\, 2\pi)$. We fix the time duration T of a data frame to be $T = 1.6$ ms, and we fix the bandwidth B to be 0.96 MHz.

We now specify transmit and receive filters. For TDM, the pulse shaping filter at the transmitter, and the matched filter at the receiver, are both narrow TD sinc pulses of bandwidth B and TD pulse width proportional to $1/B$:

$$w_{tx}(t) = w_{rx}(t) = \sqrt{B}\, sinc(Bt). \tag{49}$$

For FDM, the transmit and receive filters are narrow FD sinc pulses with width proportional to $1/T$, so that the time-realization has duration T:

$$w_{tx}(f) = w_{rx}(f) = \sqrt{T}\, sinc(fT). \tag{50}$$

For Zak-OTFS, the DD domain transmit and receive filters are the product of a narrow pulse in the delay domain with width proportional to $1/B$ and a narrow pulse in the Doppler domain with width proportional to $1/T$:

$$w_{tx}(\tau, \nu) = w_{rx}(\tau, \nu) = \sqrt{BT}\, sinc(B\tau)\, sinc(T\nu). \tag{51}$$

We measure BER performance as a function of the received SNR, which is the ratio of the power of the information carrying signal to the power of the AWGN in the received TD signal. We normalize the complex channel gains by setting $\sum_{i=1}^{6} \mathbb{E}[|h_i|^2] = 1$. We define the transmitted signal power P_T to be the ratio of the energy of the transmitted signal to the frame duration T:

$$P_T \triangleq \frac{\int |s_{\text{td}}(t)|^2\, dt}{T}. \tag{52}$$

If N_0 is the noise power spectral density of the AWGN $n_{\text{td}}(t)$, then the noise power at the receiver is $N_0 B$, and the signal-to-noise ratio (SNR) is given by

$$\gamma \triangleq \frac{P_T}{N_0 B}. \tag{53}$$

We assume that the Zak-OTFS receiver knows the effective channel matrix $\mathbf{H}_{\text{dd}}$ perfectly, that the TDM receiver knows

$$H_{\rm dd}[k'N+l',kN+l] = \sum_{n=-\infty}^{\infty}\sum_{m=-\infty}^{\infty} h_{\rm dd}[k'-k-nM,l'-l-mN]\, e^{j2\pi nl/N} e^{j2\pi \frac{l'-l-mN}{N}\frac{k+nM}{M}},$$
$$k',k=0,1,\cdots,M-1,\ l',l=0,1,\cdots,N-1. \quad (34)$$

TABLE I: Power Delay Profile of Doubly-spread Veh-A Channel.

Path no. i	1	2	3	4	5	6
Rel. Delay τ_i (μs)	0	0.31	0.71	1.09	1.73	2.51
Rel. Power $\frac{\mathbb{E}[\lvert h_i\rvert^2]}{\mathbb{E}[\lvert h_1\rvert^2]}$ (dB)	0	-1	-9	-10	-15	-20

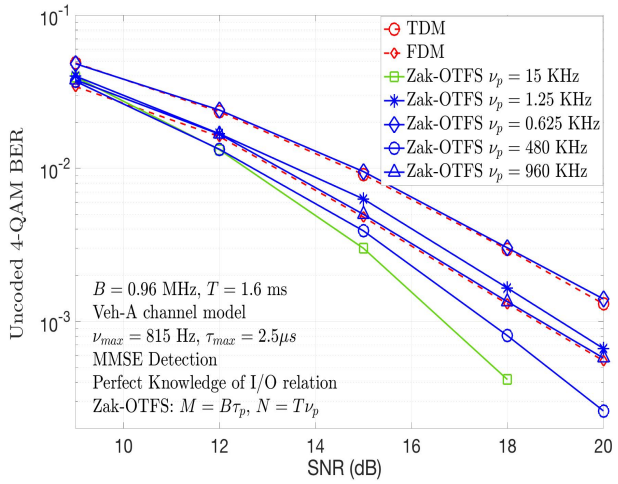

Fig. 7: BER performance of Zak-OTFS, TDM and FDM, on a doubly-spread Veh-A channel as we traverse the hyperbola $\tau_p \cdot \nu_p = 1$. Performance of Zak-OTFS is superior in the crystalline regime ($\nu_p = 15$ KHz), performance approaches TDM as the delay period grows ($\nu_p = 1.25$ and 0.625 KHz) and performance approaches FDM as the Doppler period grows ($\nu_p = 480$ and 960 KHz).

$\mathbf{H}_{\rm td}$ perfectly, and that the FDM receiver knows $\mathbf{H}_{\rm fd}$ perfectly. *Note that perfect knowledge of the I/O relation does not imply perfect knowledge of the channel delay-Doppler spreading function.* The matrix-vector I/O relations given by (35), (41) and (47) have the same structure as that of a MIMO system. Hence, we use a Linear Minimum Mean Squared Error (LMMSE) equalizer to detect the transmitted 4-QAM information symbols.

Fig. 7 illustrates how BER performance of Zak-OTFS changes as we move along the hyperbola $\tau_p \cdot \nu_p = 1$, choosing Doppler periods $\nu_p = 0.625, 1.25, 15, 480$, and 960 KHz. When $\nu_p = 15$ KHz, the crystallization conditions hold, the I/O relation is non-fading (see (15)), and BER performance is superior to both TDM and FDM. This is because the I/O relations for TDM and FDM exhibit fading on doubly spread channels.

As the Doppler period increases, first to 480 KHz, then to 960 KHz, the delay spread is no longer less than the delay period, and Zak-OTFS operates outside the crystalline regime. Aliasing occurs along the delay axis, and the received symbol energy varies (fades) along the Doppler axis. When $\nu_p = 960$ KHz, the BER performance of Zak-OTFS coincides with FDM, which is expected, given that FDM coincides with Zak-OTFS in the limit of vanishing delay period.

As the Doppler period decreases, first to 1.25 KHz, then to 0.625 KHz, the Doppler spread is no longer less than the Doppler period, and Zak-OTFS operates outside the crystalline regime. Aliasing now occurs along the Doppler axis, and the received symbol energy varies (fades) along the delay axis. When $\nu_p = 0.625$ KHz, the BER performance of Zak-OTFS almost coincides with TDM, which is expected, given that TDM coincides with Zak-OTFS in the limit of vanishing Doppler period.

A *delay-only* Veh-A channel is matched to TDM and mismatched to FDM. The channel is frequency-selective, not time-selective, hence FDM exhibits fading, and TDM does not. Fig. 8 compares BER performance of TDM and FDM with that of Zak-OTFS in a delay-only Veh-A channel (i.e., $\nu_i = 0, i = 1, 2, \cdots, 6$) for Doppler periods $\nu_p = 0.625, 15$, and 960 KHz. When $\nu_p = 0.625$ KHz, the delay period $\tau_p = 1.6$ ms is much larger than the delay spread ($2.5\,\mu s$), and BER performance is essentially the same as that of TDM. As long as we operate in the crystalline regime (for example, $\nu_p = 15$ KHz and $\tau_p = 66.66\,\mu s$), there is no fading, and BER performance changes very little with change in ν_p. Outside the crystalline regime, when the Doppler period is large and the delay period is smaller than the delay spread (for example, $\nu_p = 960$ KHz and $\tau_p = 1.04\,\mu s$), the BER performance degrades considerably.

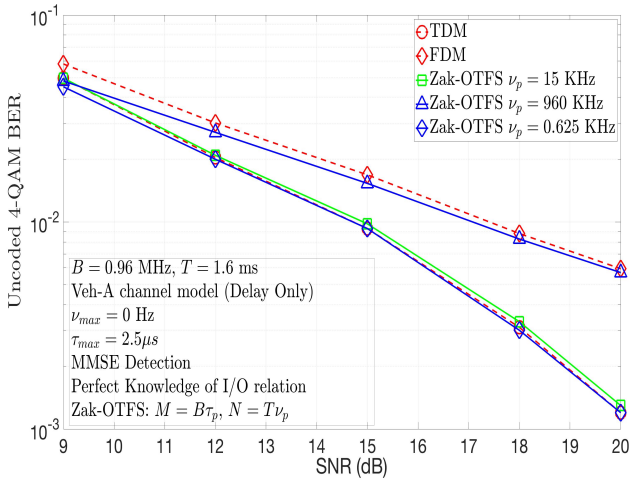

Fig. 8: BER performance of Zak-OTFS, TDM, and FDM on a delay-only Veh-A channel. The I/O relation for Zak-OTFS in the crystalline regime ($\nu_p = 15$ KHz) is non-fading, hence BER performance matches that of TDM. Frequency selective fading degrades the BER performance of FDM.

In contrast, a *Doppler-only* Veh-A channel is matched to FDM and mismatched to TDM. The channel is time-selective,

not frequency-selective, hence TDM exhibits fading and FDM does not. Fig. 9 compares BER performance of TDM and FDM with that of Zak-OTFS in a Doppler-only Veh-A channel (i.e., $\tau_{max} = 0$), for Doppler periods $\nu_p = 0.625, 15$ and 960 KHz. When $\nu_p = 960$ KHz, the Doppler period is much larger than the Doppler spread (1.63 KHz), and BER performance is essentially the same as that of FDM. Again, when we operate in the crystalline regime (for example $\nu_p = 15$ KHz and $\tau_p = 66.66\,\mu s$) there is no fading, and the BER performance changes very little with change in ν_p. Outside the crystalline regime (for example $\nu_p = 0.625$ KHz), the BER performance degrades considerably.

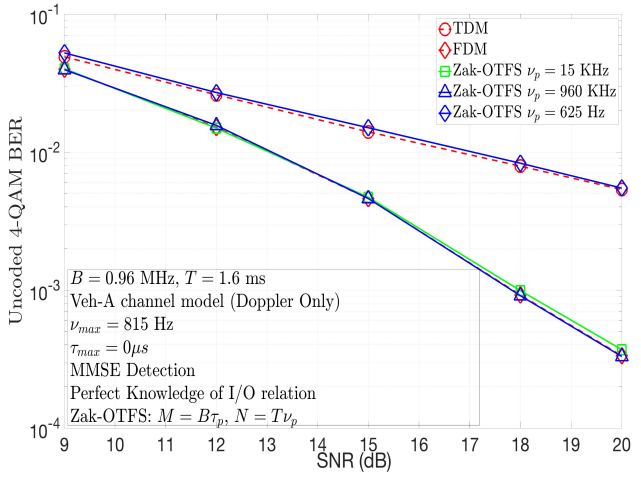

Fig. 9: BER performance of Zak-OTFS, TDM, and FDM, on a Doppler-only Veh-A channel. The I/O relation for Zak-OTFS in the crystalline regime ($\nu_p = 15$ KHz) is non-fading, hence BER performance matches that of FDM. Time selective fading degrades BER performance of TDM.

V. Model-free operation in the Crystalline regime

In this section, we compare uncoded BER performance of Zak-OTFS, TDM and FDM, in the typical scenario where the effective channel matrix is not known perfectly at the receiver. The effective channel matrix can be acquired from the I/O relation (see Section III), which in turn depends on the delay-Doppler spreading function $h(\tau, \nu)$. The *model-dependent* and *model-free* modes of operation correspond to different approaches to estimating the effective channel matrix.

Model-dependent operation: Here we impose a model on the delay-Doppler spreading function $h(\tau, \nu)$, typically by prescribing a finite number of paths, and constraining their delay and Doppler shifts. Given this model, the receiver estimates $h(\tau, \nu)$, then forms an estimate of the effective channel matrix. The accuracy of this estimate is limited by time and bandwidth constraints on the pilot signal, and by any mismatch between the channel model and the physical channel.

Model-free operation: Here the receiver estimates the effective channel filter taps directly, without reference to any model for the delay-Doppler spreading function $h(\tau, \nu)$. Model-free operation is possible in the crystalline regime with Zak-OTFS because the effective channel filter taps $h_{\text{dd}}[k, l]$ can be estimated from the response to a single pilot symbol (see Sections II-B and II-D).

The model-dependent mode of operation is a challenge for the Veh-A channel introduced in Section IV (with sinc pulse shaping filters). Since channel bandwidth $B = 0.96$ MHz, the delay domain resolution is $1/B \approx 1.04\,\mu s$, and the first three paths introduce delay shifts in the interval $[0\,,\, 0.71]\,\mu s$. These paths are not separable, and so cannot be estimated accurately.

However, model-free operation is still possible for TDM, FDM, and for Zak-OTFS in the crystalline regime ($\nu_p = 15$ KHz). We estimate the effective channel filter taps from the response to a pilot frame with a single high-energy pilot and no information symbols. We want to start from ground truth, so we suppose that the received pilot is not subject to AWGN. For Zak-OTFS we locate the single pilot at the center $(M/2\,,\, N/2)$ of the fundamental period, and we use (26) to estimate the effective filter taps. For TDM and FDM we locate the single pilot at $k = BT/2$, estimate the effective channel filter taps $h_{\text{td/fd}}[n\,;\, k = BT/2], n \in \mathbb{Z}$, and simply reuse this estimate for all other taps $h_{\text{td/fd}}[n\,;\, k]\ n \in \mathbb{Z}, k = 0, 1, \cdots, BT-1, k \neq BT/2$.

We attempt model-dependent operation, using the pilots described above to estimate the complex channel gain, delay and Doppler shift of each channel path. We interpret DD points with significant energy as the locations of channel path delays and Doppler shifts. Given the locations of delay and Doppler shifts, the received pilot signal depends linearly on the channel path gains, so we use least squares to estimate the vector of complex channel gains, then reconstruct $\mathbf{H}_{\text{dd}}$ using (34).

Fig. 10 compares BER performance of the model-dependent and model-free modes of Zak-OTFS in the crystalline regime. It also includes BER performance of TDM and FDM, which is poor, and which does not improve with increasing SNR. Since the channel is doubly spread, the I/O relation for TDM/FDM is non-stationary in the TD/FD, and it cannot be accurately predicted from the response to a single pilot at $k = BT/2$. Hence, the estimate of the effective channel matrix is inaccurate, and the probability of mis-detection is high. BER performance of the model-dependent mode of Zak-OTFS is slightly better, but it also exhibits a high error floor because insufficient frame bandwidth and duration precludes accurate estimation of channel path gains, delays and Doppler shifts.[3] BER performance of the model-free mode of Zak-OTFS is considerably better, only slightly worse than performance with perfect knowledge of the I/O relation.

Next, through Fig. 11 we explore what we lose when it is possible to learn the channel, but we choose to operate model-free. Again, we fix the time duration T of a data frame to 1.6 ms, and we fix the bandwidth B to 0.96 MHz. There

[3]For example, for the Veh-A channel, the first three paths introduce delay shifts which lie in the interval $[0\,,\, 0.71]\,\mu s$. With a channel bandwidth of $B = 0.96$ MHz, the delay domain resolution is $1/B \approx 1.04\,\mu s$ which is more than $0.71\,\mu s$ and therefore the first three dominant paths are not separable/resolvable. These non-separable paths cannot be estimated accurately.

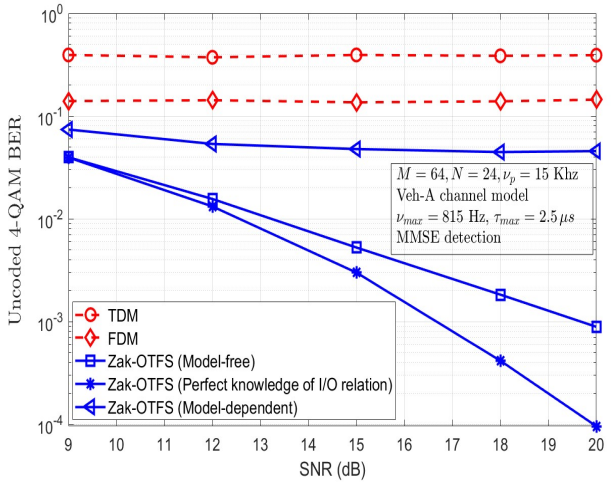

Fig. 10: BER performance of Zak-OTFS, TDM, and FDM on the Veh-A channel (Section IV), where signal bandwidth and duration is insufficient to estimate channel path delays and Doppler shifts. In the crystalline regime ($\nu_p = 15$ KHz), BER performance of the model-free mode of Zak-OTFS is only slightly worse than performance with perfect knowledge of the I/O relation.

are $M = 64$ delay bins and $N = 24$ Doppler bins. Again, the delay spread is roughly $2.5\,\mu s$, and the Doppler spread is 1.63 KHz. We consider a 5-path resolvable channel, where the delay shifts are integer multiples of the delay resolution $1/B$, and the Doppler shifts are integer multiples of the Doppler resolution $1/T$. To be precise, the path delays are $[0, 1, 2, 4, 7] \times 1/B$, the Doppler shifts are $[1, -2, -3, 3, 4] \times 1/T$, and the relative average power of the paths is $[0, -1, -9, -10, -13]$ dB. Again, we consider Zak-OTFS operating in the crystalline regime ($\nu_p = 15$ KHz).

Fig. 11 illustrates that for this resolvable channel, model-dependent performance coincides with that of performance with perfect knowledge of the I/O relation. This is expected since estimation of the channel spreading function is accurate when paths are resolvable. Further, model-free performance is only slightly inferior, despite the high Doppler spread of 4.375 KHz. Why the small degradation? Pulse shaping filters cause self-interaction/aliasing between the received pilot in the fundamental period and its quasi-periodic replicas. Though this aliasing is small in the crystalline regime, it is non-zero.[4]

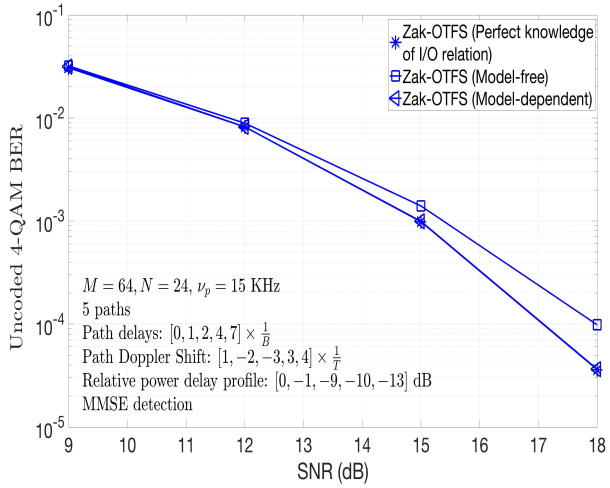

Fig. 11: BER performance of model-dependent and model-free Zak-OTFS on a 5-path resolvable channel, where the delay shifts are integer multiples of the delay resolution $1/B$, and the Doppler shifts are integer multiples of the Doppler resolution $1/T$. Model-dependent performance coincides with that of performance with perfect knowledge of the I/O relation. Model-free performance is only slightly inferior.

Next, through Fig. 12 we explore model-free performance in the crystalline regime as we traverse the hyperbola $\tau_p \cdot \nu_p = 1$, moving towards FDM. We consider the Veh-A channel introduced in Section IV, fixing the SNR at 16 dB. We measure performance for Doppler periods $\nu_p = 3.75, 7.5, 15$, and 30 KHz, as ν_{max} varies between 500 Hz and 4 KHz. For $\nu_p = 15, 30$ KHz, we are deep in the crystalline regime, since the Doppler period is significantly larger than the Doppler spread (which is at most 8 KHz), and the delay period $(66.66, 33.33\,\mu s)$ is significantly larger than the delay spread $(2.5\,\mu s)$. The spacing between quasi-periodic replicas limits self-interaction, even with sinc pulse shaping filters. The Zak-OTFS I/O relation is predictable, and we are able to accurately estimate the taps of the effective channel filter. Fig. 12 illustrates that increasing the spacing from 15 KHz to 30 KHz improves performance slightly. For both Doppler periods, the BER performance is excellent, and almost invariant to increasing Doppler spread.

BER performance changes as we reduce the Doppler period to 7.5 KHz. When the Doppler spread $2\nu_{max}$ is less

[4]In the Zak-OTFS I/O relation, the output DD signal is a twisted convolution of the effective discrete DD channel filter $h_{\text{dd}}[k, l]$ with the input information signal $x_{\text{dd}}[k, l]$, and therefore each information symbol is spread in the DD domain by an amount equal to the DD spread/width of $h_{\text{dd}}[k, l]$. Since $h_{\text{dd}}[k, l]$ is sampled from the effective continuous DD channel filter $h_{\text{dd}}(\tau, \nu)$ at integer multiples of $1/B$ and $1/T$ along the delay and Doppler domains respectively (see equation (36) in Part I), the DD spread/width of $h_{\text{dd}}[k, l]$ is directly related to the DD spread/width of $h_{\text{dd}}(\tau, \nu)$. Further, $h_{\text{dd}}(\tau, \nu) = w_{rx}(\tau, \nu) *_\sigma h(\tau, \nu) *_\sigma w_{tx}(\tau, \nu)$ is a twisted convolution of the channel DD spreading function and the transmit and receive pulse shaping filters (see equation (36) in Part I). The DD domain spread of $h_{\text{dd}}(\tau, \nu)$ is therefore the sum of the DD spread/width of $h(\tau, \nu)$ and the DD spread/width of the transmit and receive DD pulse shaping filters $w_{tx}(\tau, \nu)$ and $w_{rx}(\tau, \nu)$. Hence the maximum effective DD spread of $h_{\text{dd}}(\tau, \nu)$ is more than τ_{max} along the delay domain and more than $2\nu_{max}$ along the Doppler domain.

than 4 KHz, the interaction between the received DD pilot in the fundamental period and its quasi-periodic replicas is not significant. As the Doppler spread increases beyond 4 KHz, BER increases steadily. When the Doppler spread is the same as the Doppler period ($\nu_{max} = 3.75$ KHz), BER performance degrades completely. This is because the sinc pulse shaping filters leak energy outside their ideal delay width ($1/B = 1.04\,\mu s$) and Doppler width ($1/T = 625$ Hz). When we reduce the Doppler period to 3.75 KHz, the BER performance starts to degrade earlier at $\nu_{max} = 500$ Hz.

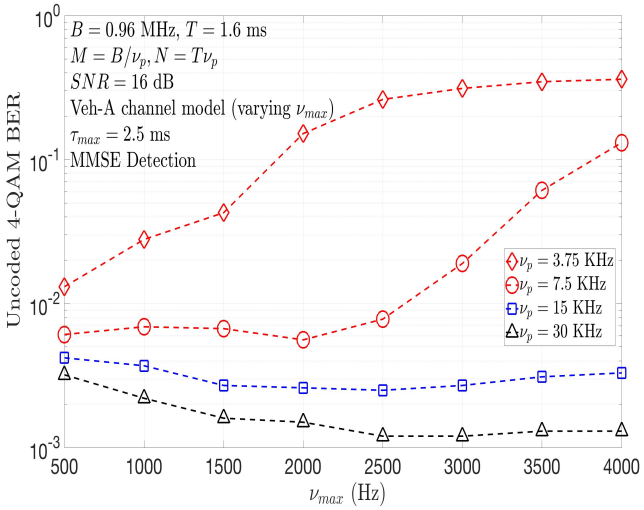

Fig. 12: . Model-free performance as a function of ν_{max} for the Veh-A channel introduced in Section IV. When the Doppler spread $2\nu_{max}$ is significantly less than ν_p, performance does not degrade as ν_{max} increases. When the Doppler spread $2\nu_{max}$ is close to ν_p, performance degrades because of Doppler domain aliasing. When operating deep in the crystalline regime, BER performance is consistently excellent over a wide range of Dopplers.

Next, through Fig. 13 we show that better pulse shaping filters extend the region of reliable model-free operation by reducing aliasing. We consider DD domain Root Raised Cosine (RRC) pulses (specified in (21)), parameterized by a roll-off parameter β_τ that controls localization in delay, and a parameter β_ν that controls localization in Doppler. The sinc pulse corresponds to $\beta_\tau = \beta_\nu = 0$, and localization improves as the parameters β_τ, β_ν, increase from 0 to 1. Better localization implies lesser DD domain aliasing and therefore accurate prediction/estimation of the I/O relation for a higher Doppler spread). We fix $\beta_\tau = 0.1$, and for $\beta_\nu = 0.3, 0.6$, and 0.9, we measure BER performance over the Veh-A channel as ν_{max} varies between 500 Hz and 4 KHz (fixed SNR of 16 dB, $M = 128$, $N = 12$, $\nu_p = 7.5$ KHz). Fig. 13 illustrates that increasing β_ν from 0 (a sinc pulse) to 0.3 extends the range of Doppler spreads for which BER performance is flat (from 4 KHz to 5 KHz). Increasing β_ν further, to $0.6, 0.9$, extends the range still further. The cost of introducing better filters is a reduction in spectral efficiency. When we replace a sinc filter by a RRC filter, we increase the frame duration by a factor $(1 + \beta_\nu)$, and we increase the frame bandwidth by $(1 + \beta_\tau)$.

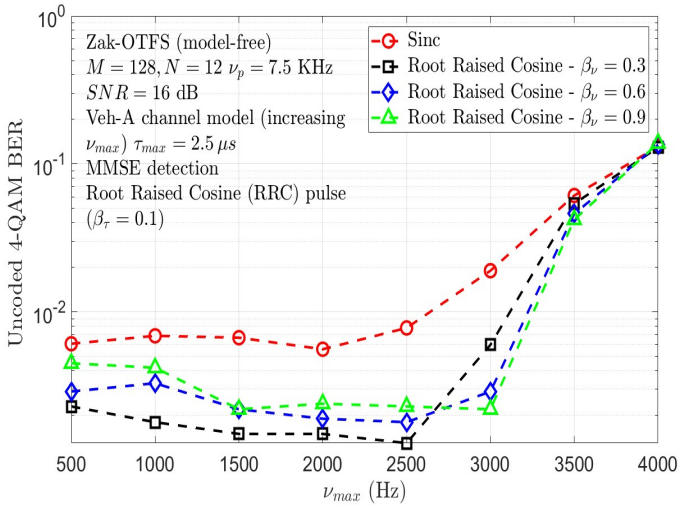

Fig. 13: Better pulse shaping filters extend the region of reliable model-free operation. BER performance, for the sinc pulse, and for RRC pulses with $\beta_\tau = 0.1$, and $\beta_\nu = 0.3, 0.6$, and 0.9, as ν_{max} varies between 500 Hz and 4 KHz.

VI. Multicarrier Approximations to Zak-OTFS

So far, all existing work on OTFS presumes a two-step modulation where the DD domain information symbols are first transformed to the time-frequency (TF) domain (using the Inverse Symplectic Finite Fourier transform or Inverse SFFT). The resulting TF symbols are then converted to a TD transmit signal using the Heisenberg transform (which is essentially an OFDM modulator) [5] (see Fig. 14). The compatibility of this two-step modulation with existing 4G/5G modems is one of the primary reasons why it was proposed.

This two-step modulation is fundamentally different from Zak-OTFS. One can think of it as a multicarrier approximation to Zak-OTFS, which we refer to here as MC-OTFS. Interestingly, the two-step implementation of MC-OTFS can be casted in the framework of the Zak-transform as shown in Fig. 15.

The simplicity of Zak-OTFS signal processing (shown in Fig. 1) results from the fact that channels and filters both act by twisted convolution and all signals in the signal processing chain are quasi-periodic. Since twisted convolution is associative, the I/O relation admits a simple structure: the output signal is the twisted convolution of the input signal with an effective DD channel filter (see (4)). This particular structure becomes predictable in the crystalline regime, implying that, in the crystalline regime, the complete I/O relation can be accurately estimated from the channel response to a single DD pilot. In contrast, as we will see now, the I/O relation of MC OTFS cannot be expressed as a simple twisted convolution and, consequently, it does not appear to satisfy any reasonable form of predictability. That is, it is not clear how to estimate the I/O relation from the observed channel response to a single DD pilot.

A careful study of the Zak-transform implementation of MC-OTFS reveals that its underlying signal processing is fundamentally different and more complicated than that of

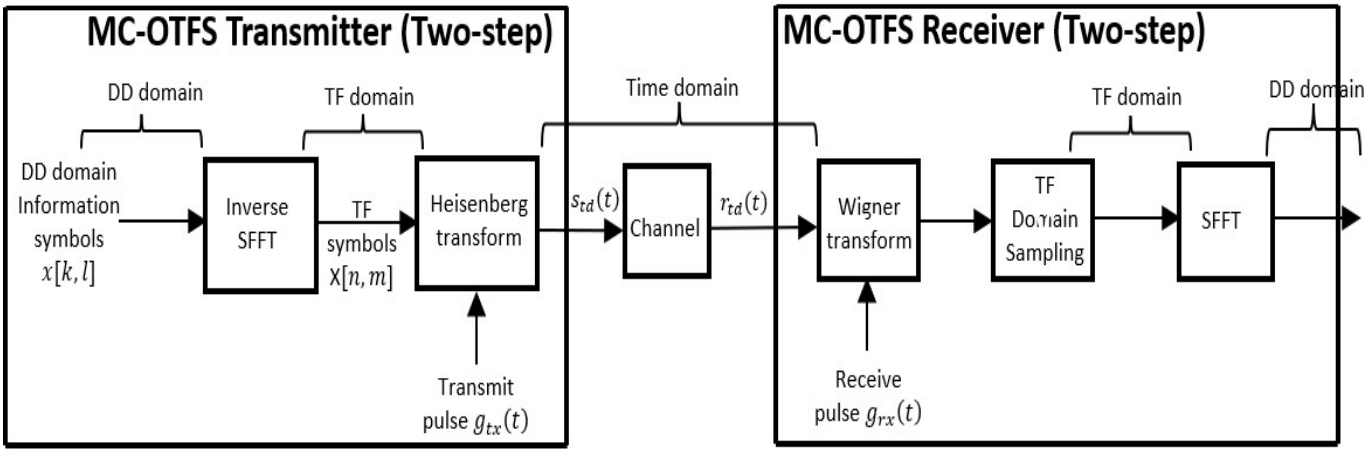

Fig. 14: Signal processing for MC-OTFS (two step).

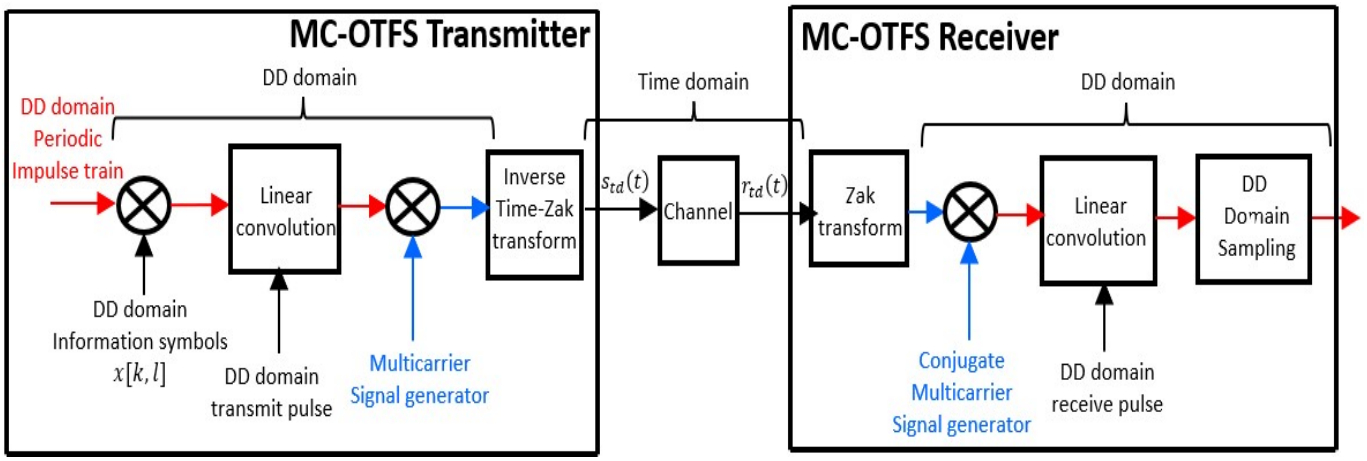

Fig. 15: Zak transform based equivalent signal processing for the two-step MC-OTFS.

Zak-OTFS. While the signal processing of Zak-OTFS comprises of pure cascade of twisted convolutions, the signal processing of MC-OTFS comprises of a mixture of linear convolution, multiplication by a quasi-periodic function and twisted convolution. As shown in Fig. 15, in MC-OTFS, the information is encoded as a discrete periodic function (instead of discrete quasi-periodic function as in Zak-OTFS) on the DD domain, given by

$$x(\tau,\nu) = \sum_{k,l\in\mathbb{Z}} x[k,l]\,\delta\left(\tau - \frac{k\tau_p}{M}\right)\,\delta\left(\nu - \frac{l\nu_p}{N}\right). \quad (54)$$

Next step is linear convolution (instead of twisted convolution as in Zak-OTFS) of the information with a transmit filter $w_{tx}(\tau,\nu)$ resulting in a periodic continuous DD domain function. The transmit pulse is usually taken to be the SFFT of a TF window (whose time and frequency support are the time duration and bandwidth of the MC-OTFS frame). Next step is to convert the periodic DD domain function to a quasi-periodic one. This is achieved by multiplication with a fixed quasi-periodic signal $G_{\text{dd}}(\tau,\nu)$ called the multi-carrier generator, which is nothing but the Zak transform of the TD transmit pulse $g_{tx}(t)$ in the two-step OTFS modulator. The transmitted signal is the inverse Zak transform of the resulting quasi-periodic signal. At the receiver, the Zak transform of the received signal which is quasi-periodic is converted into a periodic continous function by means of multiplication with the complex conjugate of the multi-carrier signal generator. This is then followed by linear convolution with a receive filter $w_{rx}(\tau,\nu)$ resulting in $y^{w_{rx}}(\tau,\nu)$ which is then sampled. Ignoring AWGN effect, the relation between the information function $x(\tau,\nu)$ and the received filtered function $y^{w_{rx}}(\tau,\nu)$ is given in Table-II (here $\star$ denotes linear convolution and $\cdot$ denotes multiplication). The counter part relation for Zak-OTFS is also given in this table. The main observation is that in Zak-OTFS, both the channel and the filters act through twisted convolution, hence due to associativity, the end-to-end signal processing is equivalent to a single twisted convolution with the effective channel filter $h_{\text{dd}}(\tau,\nu)$. In contrast, since the MC-OTFS I/O relation is a mix of linear convolution, multiplication and twisted convolution, it cannot be expressed as a simple action with some effective filter.

It is illuminating to observe, how MC-OTFS and Zak-OTFS evolve as we move towards TDM by shrinking the Doppler period ν_p. As we traverse the hyperbola $\tau_p \cdot \nu_p = 1$ along the limit $\nu_p \to 0$, the Zak transform converges to identity and Zak-OTFS converges to TDM whose I/O relation for delay-

TABLE II: I/O relation for Zak-OTFS and MC-OTFS

MC-OTFS I/O relation
$y^{w_{rx}}(\tau,\nu) \;=\; w_{rx}(\tau,\nu) \star \Big[G^{*}_{\text{dd}}(\tau,\nu)\cdot\Big(h(\tau,\nu) *_{\sigma} \Big\{G_{\text{dd}}(\tau,\nu)\cdot[w_{tx}(\tau,\nu)\star x(\tau,\nu)]\Big\}\Big)\Big]$
Zak-OTFS I/O relation
$y^{w_{rx}}_{\text{dd}}(\tau,\nu) \;=\; w_{rx}(\tau,\nu) *_{\sigma} h(\tau,\nu) *_{\sigma} w_{tx}(\tau,\nu) *_{\sigma} x_{\text{dd}}(\tau,\nu) \;=\; h_{\text{dd}}(\tau,\nu) *_{\sigma} x_{\text{dd}}(\tau,\nu)$

only channels is given by linear convolution with an effective channel filter, hence it is predictable. In contrast, as $\nu_p \to 0$, MC-OTFS does not converge to TDM. The limit in this case is a non-intuitive modulation whose I/O relation cannot be expressed as a simple linear convolution of an input signal with an effective TD channel filter.

Recall that Fig. 4 depicts the relative prediction error for Zak-OTFS as a two-dimensional heatmap, for three different points on the period hyperbola $\tau_p \cdot \nu_p = 1$. Fig. 16 depicts the relative prediction error for MC-OTFS, with respect to the same two-path channel, and the same transmit and receive filters (given by sinc pulses). The multi-carrier generator is taken to be the Zak transform of a rectangular TD pulse $g(t)$ of duration τ_p. The MC-OTFS I/O relation is derived from the continuous I/O relation given in Table-II. In this relation the input is a discrete periodic function obtained by extending $x[k,l]$ periodically with periods $M = B\tau_p$ along delay and $N = T\nu_p$ along Doppler respectively. The output is obtained by sampling $y^{w_{rx}}(\tau,\nu)$ at integer multiples of $1/B$ along the delay and $1/T$ along Doppler. As for Zak-OTFS, we estimate the MC-OTFS I/O relation from the received response to a pilot impulse located at $(M/2, N/2)$. The point $\nu_p = 1.25$ KHz, $\tau_p = 800\,\mu s$ represents the delay asymptote outside the crystalline regime, and the point $\nu_p = 240$ KHz, $\tau_p = 4.16\,\mu s$ represents the Doppler asymptote outside the crystalline regime. For these points, the relative prediction error is high for both Zak-OTFS and MC-OTFS. The point $\nu_p = 30$ KHz, $\tau_p = 33.3\,\mu s$ represents the crystalline regime. Relative prediction error for MC-OTFS is larger than for Zak-OTFS (the heatmap in Fig. 16 is a mixture of yellow and green, whereas the heatmap in Fig. 4 is mostly green).

Fig. 17 compares BER performance of Zak-OTFS and MC-OTFS on the Veh-A channel introduced in Section IV, and used in Fig. 10. We have changed the channel parameters slightly by increasing the Doppler shift ν_{max} from 815 Hz to 2 KHz. For MC-OTFS, the TF window and the multicarrier signal generator are the same as those used in Fig. 16 (see the above discussion). Given perfect knowledge of the I/O relation the performance of Zak-OTFS and MC-OTFS is quite similar. However, model-free performance of Zak-OTFS is greatly superior to that of MC-OTFS because of the Zak-OTFS I/O relation is more predictable than that of MC-OTFS.

Remark: In Fig. 17 reveals that with perfect knowledge of the I/O relation, both Zak-OTFS and MC-OTFS exhibit similar BER performance. Acquiring the I/O relation in the model-free approach amounts to estimating the channel response to a pilot at any arbitrary location. In Zak-OTFS, operating within the crystalline regime, allows to accurately predict a channel response to an arbitrary pilot from the the response to a single pilot, facilitating efficient acquisition of the complete I/O relation. In contrast, in MC-OTFS, due to the complex nature of its I/O relation, it seems that such prediction scheme is not possible, hence accurate acquisition of the complete I/O relation becomes less efficient when compared to that for Zak-OTFS.

Fig. 18 compares model-free performance of various modulation schemes as a function of increasing Doppler spread while operating at a fixed SNR of 16 dB. The simulations use the Veh-A channel introduced in Section IV. In particular, the figure illustrates that when ν_{max} is less than 1 KHz (Doppler spread less than 2 KHz), the BER performance of Zak-OTFS and MC-OTFS is essentially the same. However, as the Doppler spread increases, BER performance of MC-OTFS degrades steadily, while that of Zak-OTFS remains almost constant. The reason for that gap is because when the Doppler spread is high, the Zak-OTFS I/O relation is more predictable than that of MC-OTFS.

VII. The Radar Ambiguity Function

When we use a waveform to illuminate a radar scene, and we correlate the return with the transmitted waveform, the radar ambiguity function expresses the blurriness of the scene [12]. We begin by considering a single target and no reflectors, so that the received signal is given by

$$r_{\text{td}}(t) \;=\; h\, s_{\text{td}}(t-\tau)\, e^{j2\pi\nu(t-\tau)} \;+\; n_{\text{td}}(t), \qquad (55)$$

where $s_{\text{td}}(t)$ is the transmitted radar waveform and $n_{\text{td}}(t)$ is the AWGN at the radar receiver. We obtain the maximum likelihood (ML) estimate for the delay τ and Doppler ν using matched filtering at the radar receiver. Thus $(\widehat{\tau},\widehat{\nu}) = \arg\max_{\tau,\nu} |A_{r,s}(\tau,\nu)|$, where

$$A_{r,s}(\tau,\nu) \;\triangleq\; \int r_{\text{td}}(t)\, s^{*}_{\text{td}}(t-\tau)\, e^{-j2\pi\nu(t-\tau)} dt \quad (56)$$

is the cross-ambiguity function between the received and the transmitted waveform. The cross-ambiguity function is also relevant for the general radar scene where there are multiple targets and reflectors. Using (30) in the expression for $A_{r,s}(\tau,\nu)$ we get

$$\begin{aligned} A_{r,s}(\tau,\nu) \;&=\; h(\tau,\nu) *_{\sigma} A_{s,s}(\tau,\nu) \\ &\quad + \int n_{\text{td}}(t)\, s^{*}_{\text{td}}(t-\tau)\, e^{-j2\pi\nu(t-\tau)} dt, \end{aligned} \qquad (57)$$

where $*_{\sigma}$ denotes twisted convolution and $h(\tau,\nu)$ is the DD spreading function of the channel between the radar transmitter and receiver. In (57), $A_{s,s}(\tau,\nu)$ is the (auto-) ambiguity

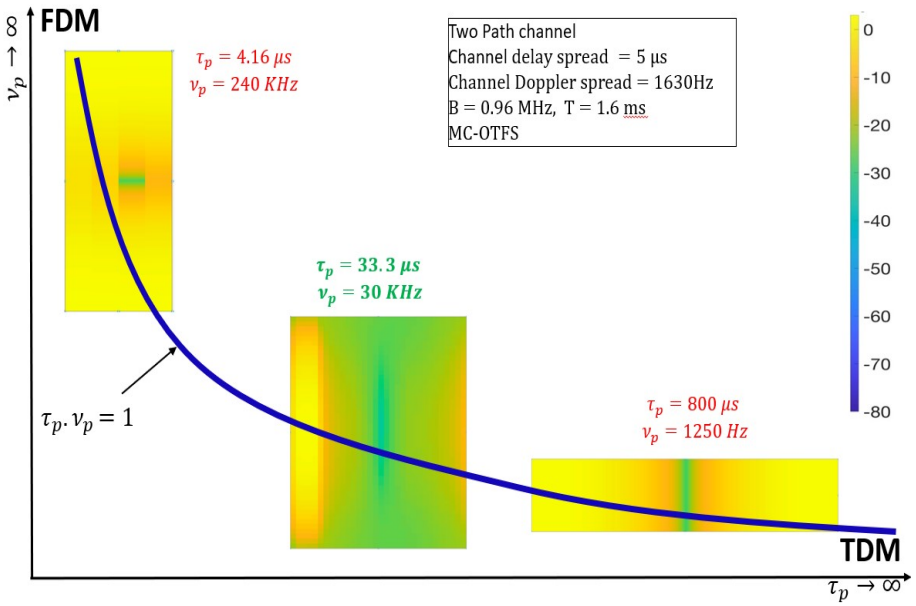

Fig. 16: Heatmap showing Relative Prediction Error (RPE), in dB, for MC-OTFS (with sinc pulses), as a function of delay (horizontal axis), and Doppler (vertical axis). RPE for MC-OTFS is larger than that for Zak-OTFS.

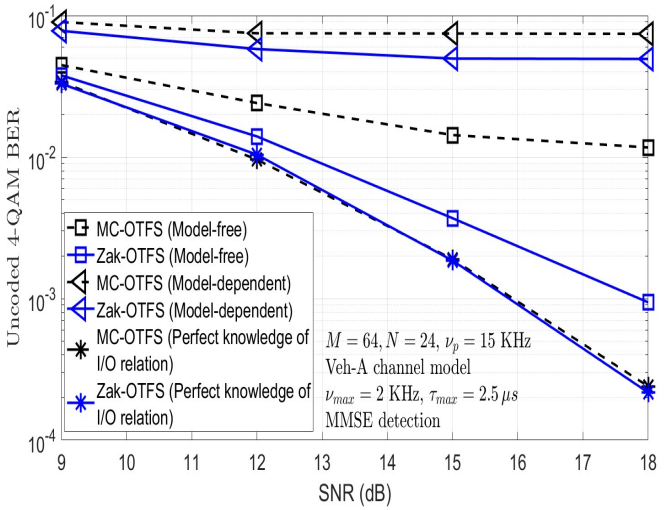

Fig. 17: Uncoded BER performance for Zak-OTFS and MC-OTFS on a Veh-A channel as a function of increasing SNR. The performance of model-dependent Zak-OTFS and MC-OTFS is poor since channel bandwidth and duration is not sufficient to estimate channel path gains, delays and Doppler shifts. Performance of model-free Zak-OTFS is superior to that of model-free MC-OTFS, because the Zak-OTFS I/O relation is more predictable.

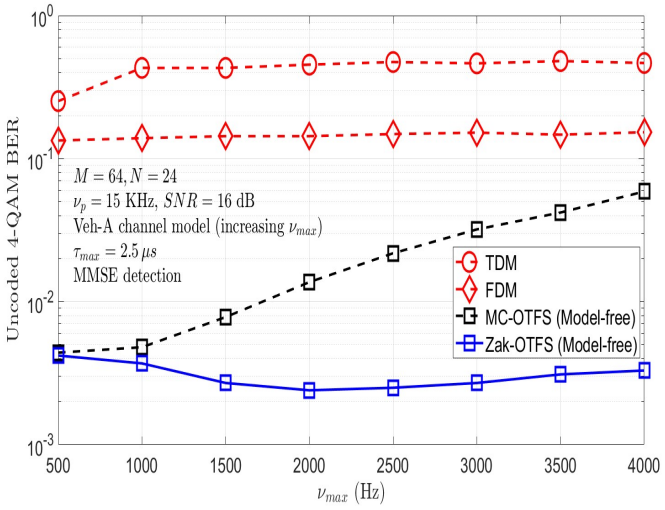

Fig. 18: Uncoded BER performance of model-free Zak-OTFS and model-free MC-OTFS as a function of ν_{max} for the Veh-A channel introduced in Section IV ($\nu_p = 15$ KHz, $M = 64, N = 24$). Zak-OTFS is more robust to high Doppler spreads.

function of the transmitted radar waveform $s_{\text{td}}(t)$ and is given by

$$A_{s,s}(\tau,\nu) \triangleq \int s_{\text{td}}(t)\, s^*_{\text{td}}(t-\tau)\, e^{-j2\pi\nu(t-\tau)} dt. \quad (58)$$

The ambiguity function $A_{s,s}(\tau,\nu)$ places fundamental limits on the *blur* which constrains our ability to estimate target distance (delay) and velocity (Doppler). Moyal's identity [13] captures the fundamental limits on blur by using the energy in the signal to provide a lower bound on the volume under the squared ambiguity surface.

$$\iint |A_{s,s}(\tau,\nu)|^2\, d\tau d\nu \;=\; \left(\int |s_{\text{td}}(t)|^2\, dt\right)^2. \quad (59)$$

Intuitively, the radar engineer aims to manipulate the ambiguity surface so that blur is concentrated in those regions that matter least for the operational task of the radar.

Next, we illustrate how the spread of the ambiguity function limits the resolution of the radar by considering a radar scene

with two targets. Thus

$$h(\tau,\nu) \quad = \quad \sum_{i=1}^{2} h_i \delta(\tau-\tau_i)\delta(\nu-\nu_i), \tag{60}$$

where (τ_1,ν_1), (τ_2,ν_2) are the delay-Doppler domain locations of the two targets. The noise-free cross-ambiguity is given by

$$h(\tau,\nu) *_\sigma A_{s,s}(\tau,\nu) = h_1 A_{s,s}(\tau-\tau_1,\nu-\nu_1)e^{j2\pi\nu_1(\tau-\tau_1)} + h_2 A_{s,s}(\tau-\tau_2,\nu-\nu_2)e^{j2\pi\nu_2(\tau-\tau_2)}$$

and we resolve the two targets by limiting the overlap between the two terms appearing on the R.H.S. In other words, we require the delay domain spread of $A_{s,s}(\tau,\nu)$ to be less than $|\tau_1-\tau_2|$ and the Doppler domain spread of $A_{s,s}(\tau,\nu)$ to be less than $|\nu_1-\nu_2|$.

A. Ambiguity functions for TDM and FDM waveforms

We consider a TDM pulse $s(t) = s_{\text{td}}(t) = \sqrt{B}\,sinc(Bt)$ with bandwidth B, that is localized around $t=0$. It follows from (58) that the ambiguity function in this case is given by

$$A^{\text{tdm}}_{s,s}(\tau,\nu) = \begin{cases} \left(1-\frac{|\nu|}{B}\right) e^{j\pi\nu\tau}\, sinc((B-|\nu|)\tau) & ,|\nu|<B \\ 0 & ,|\nu|\geq B \end{cases}. \tag{61}$$

For a fixed ν, we consider the term $sinc((B-|\nu|)\tau)$ as a function of τ and conclude that the delay spread of the TDM ambiguity function is about $1/B$. Hence, it is possible to separate two targets if their delays differ by more than $1/B$. On the other hand, the Doppler spread of the TDM ambiguity function is the bandwidth B. Hence, it is not possible to separate two targets with delays differing by less than $1/B$ unless the two Doppler shifts differ by more than B, and this is unlikely in most scenarios. The TDM waveform is localized in the TD, but not in the FD, and this is the reason it is unable to separate targets in the Doppler domain.

Next, we consider an FDM pulse $s(f) = s_{\text{fd}}(f) = \sqrt{T}\,sinc(fT)$ with duration T, that is localized around $f=0$. It follows from (58) that the ambiguity function in this case is given by

$$A^{\text{fdm}}_{s,s}(\tau,\nu) = \begin{cases} \left(1-\frac{|\tau|}{T}\right) e^{j\pi\nu\tau}\, sinc((T-|\tau|)\nu) & ,|\tau|<T \\ 0 & ,|\tau|\geq T \end{cases}. \tag{62}$$

Now, the Doppler spread is small (about $1/T$) and the delay spread is large (the duration T), hence it is not possible to separate two targets with Dopplers differing by less than $1/T$ unless the two delay shifts differ by more than T. This is unlikely in most scenarios. The FDM waveform is localized in the FD, but not in the TD, and this is the reason it is unable to separate targets in the delay domain.

The volume under the (squared) magnitude ambiguity surface is fixed by Moyal's identity (59), but it can be redistributed to enable resolution of radar targets. This possibility was known to P. M. Woodward more than 70 years ago. In his 1953 book [7], he described how the ambiguity function of a narrow Gaussian pulse (the red shaded ellipse in Fig. 20) can be redistributed into several DD domain functions/pulses

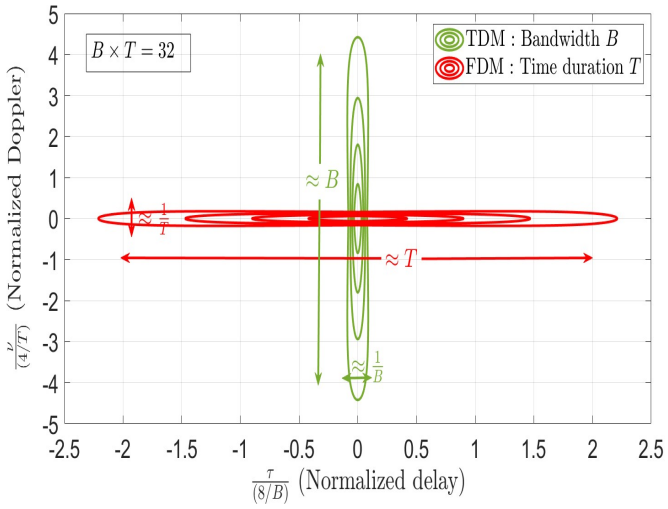

Fig. 19: Squared magnitude $|A_{s,s}(\tau,\nu)|^2$ of the ambiguity functions for TDM and FDM. The TDM carrier waveform is not able to separate targets in Doppler, and the FDM carrier waveform is not able to separate targets in delay.

(shown as black ellipses in Fig. 20). The trick is to modulate a train of narrow TD Gaussian pulses with a broad Gaussian envelope. There is a striking resemblance between Woodward's waveform and the Zak-OTFS carrier waveform (pulsone), which is a train of narrow pulses modulated by a sinusoid.

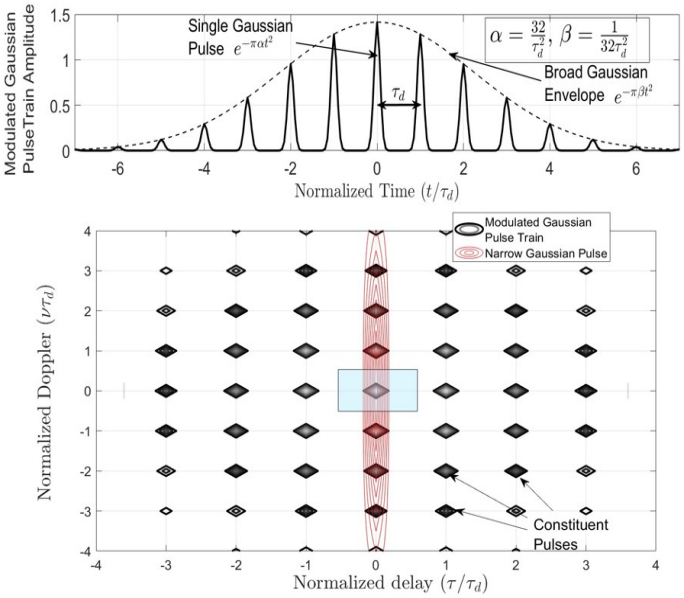

Fig. 20: Redistributing the squared magnitude $|A_{s,s}(\tau,\nu)|^2$ of the ambiguity function of a narrow Gaussian pulse. Modulating a train of narrow TD Gaussian pulses with a broad Gaussian envelope, produces an ambiguity function that is better able to separate targets in delay and Doppler.

B. Ambiguity function of Zak-OTFS pulsone

The Zak-OTFS carrier waveform is a pulsone in time which realizes a quasi-periodic pulse in the DD domain at some location (τ_0, ν_0) in the fundamental period. Recall from Part I, Table III that after filtering at the transmitter, the DD domain signal is given by

$$x_{\text{dd}}^{w_{tx}}(\tau,\nu) = w_{tx}(\tau,\nu) \; *_\sigma \sum_{n,m\in\mathbb{Z}} e^{j2\pi n\nu\tau_p}\delta(\tau - n\tau_p - \tau_0)\delta(\nu - m\nu_p - \nu_0). \quad (63)$$

Converting from the DD domain to the time domain by applying the inverse Zak transform, yields

$$s_{\text{td}}(t) \;=\; \sqrt{\tau_p}\int_0^{\nu_p} x_{\text{dd}}^{w_{tx}}(t,\nu)\, d\nu. \quad (64)$$

We assume the number of Doppler bins (N) is even, and we apply a sinc shaping filter $w_{tx}(\tau,\nu) = \sqrt{BT}sinc(B\tau)sinc(T\nu)$ to an impulse located at the origin i.e., $(\tau_0,\nu_0) = (0,0)$. The corresponding TD pulsone is given by

$$s_{\text{td}}(t) = \frac{\sqrt{B}}{\sqrt{N}} \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} sinc\left(B(t - n\tau_p - \tau_0)\right)\, e^{j2\pi n\nu_0\tau_p}. \quad (65)$$

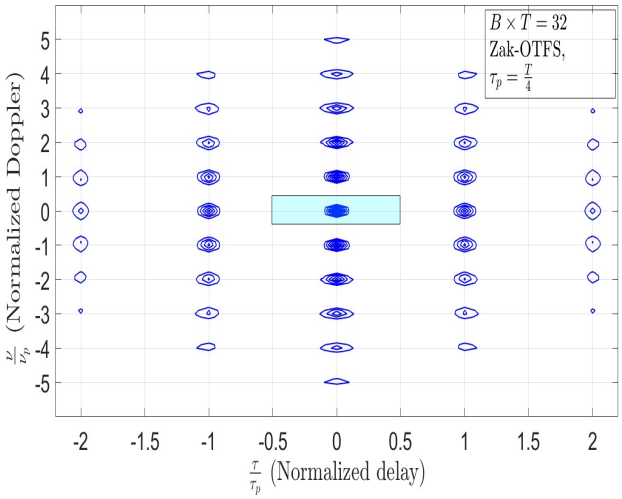

Fig. 21: Plot of the ambiguity function (squared magnitude) for the Zak-OTFS carrier waveform. Simultaneous delay and Doppler resolvability can be achieved. Unambiguous target estimation is achievable in the crystalline regime.

The corresponding ambiguity function is given in (66) (see top of next page), and is illustrated in Fig. 21 for the same time and bandwidth constraints as the TDM and FDM waveforms illustrated in Fig. 19.

The ambiguity function of the TD pulsone consists of narrow DD domain impulses separated by τ_p along the delay axis and ν_p along the Doppler axis. Each impulse has a spread of $1/B$ along the delay axis, and a spread of $1/T$ along the Doppler axis. Recall from (57), that in the absence of noise, the cross-ambiguity function is the twisted convolution of the delay-Doppler spreading function $h(\tau,\nu)$ and the ambiguity function $A_{s,s}(\tau,\nu)$. We can guarantee unambiguous delay–Doppler estimation by choosing τ_p to be greater than the delay spread, and ν_p to be greater than the Doppler spread. These are exactly the crystallization conditions guaranteeing predictability of the I/O relation in Zak-OTFS communication. Note that we can improve resolvability by increasing bandwidth B and duration T.

Given an arbitrary pulse shaping filter, it follows from (58) that the ambiguity function of the corresponding TD pulsone consists of pulses, where the pulse within the fundamental period (depicted in Fig. 21 as a light blue rectangle) is simply the twisted convolution of $w_{tx}(\tau,\nu)$ with its complex conjugate. In the above example, the pulse shaping filter is a sinc pulse, the Zak-OTFS pulsone is a train of sinc pulses, and the ambiguity function in Fig. 21 also consists of sinc pulses. In general, this structure enables precise design of ambiguity functions.

VIII. Conclusions

In this paper introduced a parametric family of waveforms, called Zak-pulsones, that can be matched to the delay and Doppler spreads of different propagation environments. We explained that a pulsone is a signal on the time domain which realizes a quasi-periodic localized function on the DD domain. The prototypical structure of a pulsone is a train of pulses in regular intervals modulated by a tone. We have put special emphasis on the case when the pulsone parameters matches the channel parameters, that is, the periods of the pulsone are greater than the corresponding spreads of the channel, a condition we refer to as the crystallization condition. Pulsones are universal and include as special cases conventional TDM and FDM waveform. A TDM waveform is a pulsone with infinite delay period and zero Doppler period, and, reciprocally, an FDM waveform is a pulsone with infinite Doppler period and zero delay period.

We have explained, why the I/O relation of sampled system based on pulsones becomes non-predictable and fading when the parameters of the pulsone do not satisfy the crystallization condition. Specifically, We have shown that the phenomena of non-predictability results from aliasing in the DD domain, which, in turns, occurs when one of the channel spreads is greater than the corresponding pulsone period - violation of the crystallization condition. In the context of radar sensing using pulsones, we have shown that inability to separate resolvable channel reflections (referred to as sensing ambiguity), occurs when one of the propagation spreads is greater than the corresponding pulsone period - again, violation of the crystallization condition.

Alternatively, we have explained why a sampled communication system based on pulsones yields superior performance when operated deep within the crystalline regime, i.e, when the spreads of the channel are considerably smaller than the pulsone periods. Specifically, we revealed that operation in the crystalline regime provides two mechanisms for improved performance: one mechanism minimizes DD domain aliasing through proper choice of the pulsone periods which results

$$A_{s,s}^{\text{otfs}}(\tau,\nu)=\begin{cases}\frac{\left(1-\frac{|\nu|}{B}\right)}{N}\sum\limits_{n_1=-\frac{N}{2}}^{\frac{N}{2}-1}\sum\limits_{n_2=-\frac{N}{2}}^{\frac{N}{2}-1}\left[e^{j\pi\nu(\tau-(n_1+n_2)\tau_p)}\,sinc\left((B-|\nu|)(\tau+(n_1-n_2)\tau_p)\right)\right] & ,\ |\nu|<B\\ 0 & ,\ |\nu|\geq B\end{cases}. \tag{66}$$

with increased predictability of the I/O relation, and a second mechanism maximizes reflection resolvability through proper choice of shaping filters, which results with increased diversity exploitation. In the context of radar sensing, minimizing DD domain aliasing translates to reducing ambiguity among resolvable reflections, and maximizing resolvability translates to increased resolution.

Another important implication of operating a sampled communication system based on pulsones within the crystalline regime is that the I/O relation, due to its increased predictability, can be learnt directly without the need to know the parameters of the underlying channel. This opens up the possibility of a model-free mode of operation, which is especially useful when channel estimation is out of reach.

The paper also includes a detailed comparison between Za-OTFS and its multi-carrier approximation, which we refer to as MC-OTFS, that has been the focus of almost all research attention so far. The two modulation schemes are compared both on theoretical grounds and on performance grounds. On the theoretical side, We have shown that the I/O relation of MC-OTFS is less predictable than that of Zak-OTFS ,which imply that as the Doppler spread increases, the BER performance of MC-OTFS is inferior to that of Zak OTFS. This suggests that MC-OTFS is less adapted to model free mode of operation than Zak-OTFS.

In conclusion, convergence of communications and sensing in 6G and beyond has focused research attention on the design of carrier waveforms that support both applications. Noticing that an environment, be it a radar scene or a communication medium is characterized by its delay and Doppler spreads. The underlying message of this paper is that in both contexts - radar and communication, it is beneficial to choose the periods of the pulsone to be greater than the spreads of the environment, i.e., one should operate in the crystalline regime.

REFERENCES

[1] H. Tataria, M. Shafi, A. F. Molisch, M. Dohler, H. Sjöland, and F. Tufvesson, "6G wireless systems: vision, requirements, challenges, insights, and opportunities," *Proceedings of the IEEE*, vol. 109, no. 7, pp. 1166-1199, July 2021.

[2] 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*, 2022. Available IEEE Early Access: doi: 10.1109/MBITS.2022.3216536.

[3] J. Zak, "Finite translations in solid state physics," *Phy. Rev. Lett.*, 19, pp. 1385-1387, 1967.

[4] A. J. E. M. Janssen, "The Zak transform: a signal transform for sampled time-continuous signals," *Philips J. Res.*, 43, pp. 23-69, 1988.

[5] R. Hadani, S. Rakib, M. Tsatsanis, A. Monk, A. J. Goldsmith, A. F. Molisch, and R. Calderbank, "Orthogonal time frequency space modulation," *Proc. IEEE WCNC'2017*, pp. 1-6, Mar. 2017.

[6] "Best Readings in Orthogonal Time Frequency Space (OTFS) and Delay Doppler Signal Processing," June 2022. https://www.comsoc.org/publications/best-readings/orthogonal-time-frequency-space-otfs-and-delay-doppler-signal-processing

[7] P. M. Woodward, *Probability and Information Theory with Applications to Radar*, Pergamon Press, 1953.

[8] N. Levanon and E. Mozeson, *Radar Signals*, Wiley, Hoboken, 2004.

[9] M.J. Skolnik, *An Introduction and Overview of Radar*, 3rd edition, McGraw-Hill, New York, 2008.

[10] P. A. Bello, "Characterization of randomly time-variant linear channels," *IEEE Trans. Comm. Syst.*, vol. 11, pp. 360-393, 1963.

[11] S. Haykin, *Digital Communication Systems*, John Wiley and Sons, 2014.

[12] L. Auslander and R. Tolmieri, "Radar ambiguity functions and group theory," *SIAM J. Mathematical Analysis*, vol. 16, no. 3, pp. 577-601, 1985.

[13] W. Moran, The mathematics of radar, in Twentieth Century Harmonic Analysis – A Celebration (NATO Science Series, NATO), J.S. Byrnes Ed., Norwell, MA: Kluwer, pp. 295-328, 2001.

[14] T. Thaj, E. Viterbo and Y. Hong,"General I/O relations and low-complexity universal MRC detection for all OTFS variants," *IEEE Access*, vol. 10, pp. 96026-96037, 2022.

[15] S. K. Mohammed,"Derivation of OTFS modulation from first principles," *IEEE Trans. Veh. Tech.*, vol. 70, no. 8, pp. 7619-7636, Aug. 2021.

[16] S. K. Mohammed,"Time-domain to delay-Doppler domain Conversion of OTFS signals in very high mobility scenarios," *IEEE Trans. Veh. Tech.*, vol. 70, no. 6, pp. 6178-6183, Jun. 2021.

[17] F. Lampel, A. Avarado and F. M. J. Willems,"On OTFS using the discrete Zak transform," *Proc. IEEE ICC'2022 Workshops*, May 2022.

[18] V. S. Bhat, G. Harshavardhan, and A. Chockalingam, "Input-output relation and performance of RIS-aided OTFS With fractional delay-Doppler," *IEEE Commun. Lett.*, vol. 27, no. 1, pp. 337-341, Jan. 2023.