Mirri *et al.*
Journal on Advances in Signal Processing (2026) 2026:24
https://doi.org/10.1186/s13634-025-01286-9

Journal on Advances in Signal Processing

RESEARCH Open Access

Zak-OTFS-based coded random access for uplink mMTC

Alessandro Mirri[1†], Venkatesh Khammammetti[2†], Beyza Dabak[2], Enrico Paolini[1], Krishna Narayanan[3] and Robert Calderbank[2*]

†Alessandro Mirri and Venkatesh Khammammetti have contributed equally to this work.

*Correspondence: robert.calderbank@duke.edu

[1] Department of Electrical, Electronic, and Information Engineering, University of Bologna, Bologna, Italy
[2] Department of Electrical and Computer Engineering, Duke University, Durham, USA
[3] Department of Electrical and Computer Engineering, Texas A&M University, College Station, USA

Abstract

This paper proposes a grant-free coded random access (CRA) scheme for uplink massive machine-type communications (mMTC), based on Zak-orthogonal time–frequency space (Zak-OTFS) modulation in the delay–Doppler domain. The scheme is tailored for doubly selective wireless channels, where conventional orthogonal frequency-division multiplexing (OFDM)-based CRA suffers from unreliable inter-slot channel prediction due to time–frequency variability. By exploiting the predictable nature of Zak-OTFS, the proposed approach enables accurate channel estimation across slots, facilitating reliable successive interference cancellation across user packet replicas. A fair comparison with an OFDM-based CRA baseline shows that the proposed scheme achieves significantly lower packet loss rates under high mobility and user density. Extensive simulations over the standardized Veh-A channel confirm the robustness and scalability of Zak-OTFS-based CRA, supporting its applicability to future mMTC deployments.

Keywords: Coded random access, Delay–Doppler communication, Internet of Things, Successive interference cancellation, Zak-OTFS

1 Introduction

Massive machine-type communications (mMTC) is a key use case in 5 G and beyond. It is characterized by a massive number of devices that sporadically and unpredictably transmit short data packets to a central base station (BS), typically under stringent constraints on energy consumption and signaling overhead, and often mild latency and reliability constraints [1–3]. To address this challenge of massive multiple access (MMA) [4, 5], grant-free random access schemes [6–8] have gained increasing attention, as they allow devices to transmit autonomously, whenever new data are available, without requiring explicit scheduling by the BS. While this uncoordinated approach reduces latency and protocol complexity, it also increases the likelihood of packet collisions among simultaneously transmitting users, potentially leading to information loss. As a result, achieving efficient and reliable communication under these constraints necessitates a carefully integrated design of the medium access control (MAC) and physical (PHY) layers. Successive interference cancellation (SIC)-based coded random access

Springer

(CRA) [9] has emerged as an effective and pragmatic integrated MAC/PHY approach for the design of grant-free random access.

Seminal works on integrated MAC/PHY designs were tailored to the collision channel [6, 7], which corresponds to the packet erasure channel at the PHY layer; hence, they did not account for a power constraint. In [10], Polyanskiy proposed a theoretical framework to model grant-free random access for the power-constrained additive white Gaussian noise (AWGN) channel, which is now popularly referred to as unsourced random access. He also derived an achievable probability of error for random coding as a function of signal-to-noise ratio (SNR), which has served as a benchmark for grant-free random access coding schemes. Polyanksiy's results demonstrate that the performance of naive random access schemes such as slotted ALOHA (SA) [11] is significantly far away from the random coding bound. Later works have shown that CRA-based schemes can be optimized for the AWGN channel and that they offer significant performance improvement over naive SA schemes [12, 13], particularly for a large number of users. Subsequent works have extended CRA-based schemes to flat Rayleigh fading channels [14, 15].

Despite these advances, existing CRA and SIC-based schemes are typically designed for simplified channel models (AWGN or flat Rayleigh fading), where SIC is generally performed at the MAC layer through packet-level interference cancellation. Such models fail to capture the time and frequency selectivity of realistic wireless channels. In particular, when the channel exhibits both significant delay and Doppler spreads, the effective channel response varies rapidly within a time–frequency resource block (slot). This variability poses a major challenge for SIC-based architectures operating at the PHY layer, which rely on accurate estimation of the input–output (I/O) relationship over the entire transmission slot. Modulation schemes such as orthogonal frequency-division multiplexing (OFDM) and single-carrier (SC) modulation are ill-suited for this setting: OFDM offers simplified equalization; however, it requires prior knowledge of the channel model for effective I/O acquisition, and SC suffers from requiring a significant number of pilots for estimating doubly spread channels.

In this paper, we engineer a CRA scheme based on Zak-orthogonal time–frequency space (Zak-OTFS) modulation [16] and show that our proposed scheme is significantly more effective than OFDM-based schemes for enabling PHY layer SIC. In Zak-OTFS modulation, the carrier waveform is a pulse in the delay–Doppler (DD) domain, formally a quasi-periodic localized function with specific periods along delay and Doppler. When the channel delay spread is less than the delay period, and the channel Doppler spread is less than the Doppler period, the response to a single Zak-OTFS carrier provides an image of the scattering environment, which is used to predict the effective channel. Since the scattering environment changes slowly and predictably, our approach enables accurate prediction of the channel, and hence, it makes SIC viable. While Zak-OTFS introduces complexity in equalization due to inter-symbol interference (ISI), it offers an important advantage over OFDM: The acquisition of the I/O relationship can be achieved in a model-agnostic manner. We demonstrate that CRA and SIC, when integrated with Zak-OTFS, significantly outperforms OFDM-based schemes in doubly selective channels, thereby enabling efficient and reliable grant-free random access for mMTC.

The main contributions of this paper are summarized as follows:

- We propose a novel grant-free CRA scheme designed directly in the DD domain, leveraging Zak-OTFS modulation to address the challenges of random access over doubly selective wireless channels. Unlike conventional approaches that operate in the time–frequency (TF) domain, the proposed scheme exploits the predictable structure of the Zak-OTFS modulation in the DD domain to enable replica decoding and effective PHY layer SIC across slots.
- We analyze the impact of pulse shaping filter design on Zak-OTFS performance, comparing Gaussian and sinc filters. This design choice enables balanced trade-offs in slot-level processing across different system performance metrics, depending on the specific system settings.
- We verify that under high-mobility scenarios, the performance of the OFDM-based scheme degrades significantly due to its reliance on inter-slot channel prediction, which compromises the effectiveness of SIC. In contrast, the proposed Zak-OTFS-based approach maintains robust and reliable SIC processing, enabled by its near-invariant I/O relationship across the entire frame.
- We perform extensive simulations under realistic Veh-A channel conditions, confirming that the proposed Zak-OTFS-based CRA scheme consistently achieves lower packet loss rate (PLR) than the OFDM baseline, particularly in high-mobility and high-user-density scenarios.

These results highlight the practical advantages of designing CRA schemes in the DD domain and establish Zak-OTFS as a viable modulation framework for scalable and reliable uplink mMTC.

2 Preliminaries and background

The MAC layer in a wireless network manages contention resolution, resource selection, and, in some cases, interference cancellation, while the PHY layer governs how information is encoded, modulated, and transmitted over the wireless channel. Meeting the stringent performance requirements of mMTC, particularly in terms of scalability under latency and reliability constraints, requires a joint and well-coordinated MAC/PHY design. This section is therefore divided into two parts: The first subsection reviews key principles and recent developments in random access protocols at the MAC layer, with particular emphasis on contention resolution and the role of SIC; and the second provides the necessary PHY layer background on Zak-OTFS modulation, which serves as the foundation of the PHY layer model adopted in this work.

2.1 MAC layer: random access and SIC

Random access protocols have historically played a central role in enabling uncoordinated communication in shared wireless networks. A foundational example is SA, where time is divided into slots and active users independently attempt transmission of a single information packet without any coordination. The analysis typically relies on the *collision channel model*, which assumes that a packet is successfully received only if it is the sole transmission in a given slot; otherwise, all colliding packets are lost. It is further

assumed that collisions are always detectable at the BS and that, at the end of each slot, a feedback message is broadcast to all users. This feedback enables each user to determine whether its packet was successfully received or if a retransmission is required in subsequent time slots. While simple and fully decentralized, this one-shot transmission strategy suffers from a high probability of collisions, especially as the number of active users grows, leading to poor throughput and reduced reliability.

To overcome these limitations and improve the system performance, more advanced random access schemes introduced temporal redundancy, where each user transmits multiple replicas of the same packet in randomly selected time slots. This strategy exploits slot diversity, increasing the probability that at least one replica avoids a collision and is successfully received. An early example of this idea is diversity slotted ALOHA (DSA) [17], which demonstrated improved reliability over traditional SA by reducing the probability of complete packet loss due to collisions.

A major breakthrough came with the introduction of SIC, which enables the receiver to iteratively decode packets and subtract their reconstructed contributions from the received signal. This process allows for the resolution of packet collisions that would otherwise lead to data loss. This principle forms the foundation of modern CRA protocols such as contention resolution diversity slotted ALOHA (CRDSA) [6] and irregular repetition slotted ALOHA (IRSA) [7], which model the access process as a sparse bipartite graph and adopt belief-propagation-inspired decoding techniques originally developed for erasure codes. In these schemes, each user transmits a predefined [6] or randomly selected [7] number of replicas in independently chosen time slots within a contention window, commonly referred to as a frame. This structure is illustrated in Fig. 1a, where the frame is divided into N_s discrete time slots, each representing a potential transmission opportunity for a packet replica. The BS then leverages the graph structure of the collisions to iteratively apply SIC, progressively removing interference and recovering additional packets.

For CRA schemes to operate effectively and support reliable SIC, the conventional assumptions of the collision channel model are typically extended to include ideal interference subtraction across slots. Under this model, perfect SIC is assumed, meaning that once a packet is successfully decoded, its contribution can be perfectly removed from the received signal in all time slots where its replicas were transmitted, regardless of the actual channel conditions. Because interference cancellation is assumed to be perfect,

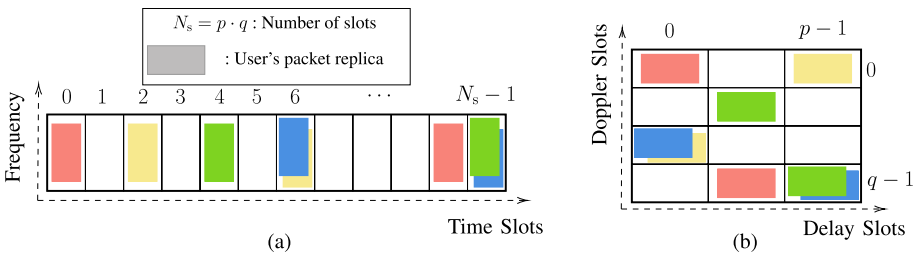

Fig. 1 Comparison between (**a**) conventional coded random access in the time–frequency domain and (**b**) proposed Zak-OTFS-based coded random access in the delay–Doppler domain. In both cases, colored rectangles represent user-specific packet replicas transmitted in randomly selected slots within the corresponding resource grid

SIC is typically modeled at the MAC layer, where packet-level removal is performed and directly coupled to medium access considerations.

However, while this assumption enables tractable analysis and protocol design, it neglects critical PHY layer impairments encountered in practical wireless deployments. In reality, the wireless channel coefficients for the same user may vary across the frame, depending on the relationship between the channel coherence time and the frame duration. This effect is particularly pronounced in high-mobility environments, such as industrial or vehicular networks, where Doppler shifts and delay spreads significantly reduce the channel coherence time (potentially down to the slot duration, which is much shorter than the frame duration), violating the assumption of invariant channels across replicas. As a result, replicas of the same packet are likely to experience different channel conditions, leading to residual interference errors and degrading the effectiveness of SIC. Under these practical conditions, interference cancellation must necessarily be performed at the PHY layer, where its performance can be severely affected by the short channel coherence time [18].

To overcome these challenges, the channel coefficients must be accurately estimated at the BS not only in the slots where a packet is successfully decoded, but also in all slots where interference needs to be removed through PHY layer SIC. Accordingly, recent research has explored more advanced BS architectures and sophisticated decoding algorithms that go beyond the idealized SIC assumption, enabling more effective interference cancellation at the PHY layer under realistic channel conditions. In particular, equipping the BS with a massive multiple-input multiple-output (MIMO) architecture introduces spatial degrees of freedom that can be effectively exploited to enhance collision resolution and boost the performance of SIC-based decoding. Recent studies, such as [19–22], have demonstrated that spatial processing in massive MIMO systems significantly improves the effectiveness of SIC, even under rapidly varying channel conditions that change on a slot-by-slot basis [23, 24]. However, as noted, these approaches concentrate the majority of the system burden at the BS, in terms of both architecture and receiver-side signal processing, without directly addressing potential improvements on the transmission side, such as the selection of a more suitable PHY layer waveform.

In this work, we adopt a complementary perspective by shifting the focus from conventional TF domain PHY processing to a DD representation, leveraging the Zak-OTFS waveform as the modulation framework. Rather than relying on traditional CRA schemes in the TF domain, as depicted in Fig. 1a, we propose a novel CRA strategy designed in the DD domain, as illustrated in Fig. 1b and further detailed in Sect. 3. This shift is motivated by the observation that wireless channels, often highly dynamic and unpredictable in the TF domain, tend to exhibit a more structured and nearly static behavior in the DD domain. In particular, the Zak-OTFS representation makes the channel approximately invariant over the entire duration of DD frame [16, 25], even in high-mobility scenarios. As a result, assumptions such as ideal MAC layer SIC, typically violated in TF settings, become significantly more realistic in the DD domain, enabling more robust and efficient grant-free random access in challenging mMTC environments.

2.2 PHY layer: Zak-OTFS

Zak-OTFS modulation is a variant of DD modulation [26–28] that multiplexes information symbols in the DD domain. In DD modulation, the delay period τ_p and the Doppler period ν_p are divided into M and N uniform intervals, respectively, such that the delay resolution is $\frac{\tau_p}{M}$ and the Doppler resolution is $\frac{\nu_p}{N}$. Also, we assume the wireless channel is doubly spread so that $\tau_p \nu_p = 1$.

In Zak-OTFS modulation, the information symbols $x[k, l]$, with indices $k = 0, 1, \cdots, M-1$ and $l = 0, 1, \cdots, N-1$, are arranged on an $M \times N$ grid and mapped to the DD lattice $\Lambda_{\text{dd}} \triangleq \{(k\tau_p/M, l\nu_p/N) \,|\, k, l \in \mathbb{Z}\}$. The quasi-periodic DD domain signal[1] $x_{\text{dd}}[k + nM, l + mN]$ is generated such that

$$x_{\text{dd}}[k + nM, l + mN] = \begin{cases} x[k, l], & \text{if } n = m = 0, \\ x[k, l] e^{j2\pi nl/N}, & \text{otherwise}, \end{cases} \tag{1}$$

where $n, m \in \mathbb{Z}$. This discrete DD domain signal is then mapped to the continuous domain as

$$x_{\text{dd}}(\tau, \nu) = \sum_{k,l \in \mathbb{Z}} x_{\text{dd}}[k, l] \delta\left(\tau - \frac{k\tau_p}{M}\right) \delta\left(\nu - \frac{l\nu_p}{N}\right), \tag{2}$$

where $\tau, \nu \in \mathbb{R}$ denote the continuous delay and Doppler variables, respectively. The signal in (2) is then passed through a transmit shaping filter $w_{\text{tx}}(\tau, \nu)$ resulting in

$$x_{\text{dd}}^{\text{tx}}(\tau, \nu) = w_{\text{tx}}(\tau, \nu) *_\sigma x_{\text{dd}}(\tau, \nu), \tag{3}$$

where $*_\sigma$ denotes the twisted convolution[2] operation. Note that the choice of transmit filter $w_{\text{tx}}(\tau, \nu)$ is a design degree of freedom that can be used to balance different measures of system performance [29, 30].

The doubly spread wireless channel is modeled through its DD impulse response $h_{\text{phy}}(\tau, \nu) = \sum_{\ell=1}^{L} h_\ell \delta(\tau - \tau_\ell) \delta(\nu - \nu_\ell)$, where h_ℓ, τ_ℓ, and ν_ℓ are the channel gain, delay spread, and Doppler spread of the ℓth path, respectively, and L is the total number of paths. The transmit signal $x_{\text{dd}}^{\text{tx}}(\tau, \nu)$ propagates through this DD domain channel, resulting in the received signal

$$y_{\text{dd}}(\tau, \nu) = h_{\text{phy}}(\tau, \nu) *_\sigma x_{\text{dd}}^{\text{tx}}(\tau, \nu) + n(\tau, \nu), \tag{4}$$

where $n(\tau, \nu)$ represents additive noise in the DD domain, corresponding to a time domain noise with power spectral density N_0.

In Zak-OTFS demodulation, the received DD domain signal $y_{\text{dd}}(\tau, \nu)$ is first filtered using the receive pulse $w_{\text{rx}}(\tau, \nu)$, yielding

$$y_{\text{dd}}^{\text{rx}}(\tau, \nu) = w_{\text{rx}}(\tau, \nu) *_\sigma y_{\text{dd}}(\tau, \nu). \tag{5}$$

[1] Note that the time Zak transform exists only for quasi-periodic signals in the DD domain. Consequently, only quasi-periodic DD domain signals admit a valid representation in the time domain.

[2] $a(\tau, \nu) *_\sigma b(\tau, \nu) = \iint a(\tau', \nu') b(\tau - \tau', \nu - \nu') e^{j2\pi\nu'(\tau - \tau')} \, d\nu' d\tau'$.

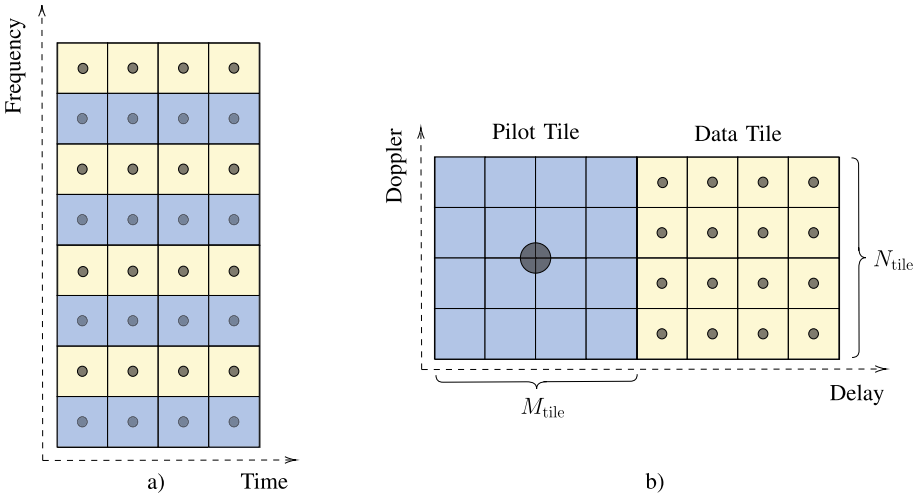

Fig. 2 Structure of a single slot in the coded random access scheme, representing one slot within the overall frame illustrated in Fig. 1. **(a)** Time–frequency domain representation of the slot, where pilot and data symbols are spread across time and frequency resources, following a structure typical of OFDM-based transmissions. **(b)** Delay–Doppler domain representation of the slot, comprising two tiles: a pilot tile (blue) and a data tile (yellow), each of size $M_{\text{tile}} \times N_{\text{tile}}$. The pilot tile includes a single point pilot symbol centered in the grid (large circle), while the data tile carries modulated data symbols (small circles)

The output $y_{\text{dd}}^{\text{rx}}(\tau, \nu)$ is then sampled at $\tau = k\tau_p/M$ and $\nu = l\nu_p/N$, resulting in the discrete received symbols

$$\begin{aligned} y[k,l] &= y_{\text{dd}}^{\text{rx}}\left(\tau = \frac{k\tau_p}{M}, \nu = \frac{l\nu_p}{N}\right) \\ &= \sum_{k',l' \in \mathbb{Z}} h_{\text{eff}}[k-k', l-l'] x_{\text{dd}}[k',l'] e^{j2\pi\left(\frac{(l-l')k'}{M}\right)}, \end{aligned} \tag{6}$$

where $k = 0, 1, \cdots, M-1$ and $l = 0, 1, \cdots, N-1$. The discrete effective channel response $h_{\text{eff}}[k,l]$ is obtained by sampling the continuous effective channel

$$h_{\text{eff}}(\tau, \nu) = w_{\text{rx}}(\tau, \nu) *_\sigma h_{\text{phy}}(\tau, \nu) *_\sigma w_{\text{tx}}(\tau, \nu). \tag{7}$$

The effective channel $h_{\text{eff}}[k,l]$ can be accurately predicted in Zak-OTFS, provided that the channel delay spread is less than the delay period and the channel Doppler spread is less than the Doppler period. (This is the *crystalline condition* introduced in [16, 25].)

In the next section, we present the proposed uplink MMA scheme for time-varying channels, integrating both the MAC and PHY layer aspects discussed above.

3 Proposed uplink MMA scheme

3.1 OTFS-based framing and random access protocol

We consider a CRA scheme operating in the DD domain based on Zak-OTFS modulation. In the proposed setup, each user transmits a complete OTFS frame in the uplink. The frame is divided into multiple slots, with each slot occupying M_{slot} and N_{slot} resources along the delay and Doppler axes, respectively. The full frame spans $M \times N = p \cdot M_{\text{slot}} \times q \cdot N_{\text{slot}}$ resources, where p and q denote the number of slots arranged along the delay and Doppler dimensions, as shown in Fig. 1b. The fundamental

unit of an OTFS slot is the *tile*, defined as a region of size $M_{\text{tile}} \times N_{\text{tile}}$ in the DD grid (see Fig. 2b). Each tile is allocated either to pilot or data (payload) transmission. A single OTFS slot is formed by concatenating two adjacent tiles, i.e., one for pilot and one for data. This concatenation can occur along either the delay or Doppler axis; for simplicity, we adopt concatenation along the delay axis, which yields slot dimensions of $M_{\text{slot}} \times N_{\text{slot}} = 2M_{\text{tile}} \times N_{\text{tile}}$. While the choice of axis affects indexing, it does not impact system operation as long as the crystalline condition is satisfied.

In the proposed protocol, each active user independently selects r non-overlapping OTFS slots with uniform probability within the frame for transmission. Selecting a slot[3] entails choosing both a pilot tile and its associated data tile. The user then transmits an identical replica of its packet in each selected slot, i.e., the same pilot symbol is placed in each pilot tile, and the same modulated data symbols are used in each corresponding data tile. This replica consistency is crucial for enabling coherent SIC at the receiver, as it ensures that all copies carry identical information. Overall, each active user occupies $r \cdot M_{\text{slot}} \cdot N_{\text{slot}}$ DD resources, distributed across r slots over the entire Zak-OTFS frame, which spans a bandwidth $B = \nu_p \cdot M$ and duration $T = \tau_p \cdot N$.

3.2 Signal processing at the transmitter

Consider an uplink MMA system where K_{a} users aim to communicate simultaneously with the BS. Let $x_{k_{\text{a}}}[k,l]$, where $k = 0, 1, \cdots, M-1$ and $l = 0, 1, \cdots, N-1$, denote the information transmitted by the k_{a}th user over the entire DD frame. Let $s_{k_{\text{a}}}[\tilde{k}, \tilde{l}]$, with $\tilde{k} = 0, 1, \cdots, 2M_{\text{tile}} - 1$ and $\tilde{l} = 0, 1, \cdots, N_{\text{tile}} - 1$, represent the symbols transmitted by the k_{a}th user within a single slot[4] (see Sect. 3.1 for details). For $\tilde{k} = 0, 1, \cdots, M_{\text{tile}} - 1$ and $\tilde{l} = 0, 1, \cdots, N_{\text{tile}} - 1$, $s_{k_{\text{a}}}[\tilde{k}, \tilde{l}]$ corresponds the pilot tile. In this region, the user transmits a single pilot symbol at the center position $(k_p, l_p) = (M_{\text{tile}}/2, N_{\text{tile}}/2)$ with all other symbols set to zero.[5] For $\tilde{k} = M_{\text{tile}}, M_{\text{tile}} + 1, \cdots, 2M_{\text{tile}} - 1$ and $\tilde{l} = 0, 1, \cdots, N_{\text{tile}} - 1$, $s_{k_{\text{a}}}[\tilde{k}, \tilde{l}]$ defines the data tile, where the user transmits modulated information symbols. As illustrated in Fig. 1b, the DD frame is partitioned into $N_{\text{s}} = p \cdot q$ slots. These slots are indexed sequentially from 0 to $N_{\text{s}} - 1$ and arranged into a $p \times q$ slot index matrix following a row-wise order. This matrix representation facilitates efficient indexing and systematic access to slot locations during transmission and reception. Let $a \in \{0, 1, \cdots, N_{\text{s}} - 1\}$ denote a linear slot index. Its corresponding matrix indices (i, j) are given by

$$\begin{aligned} i &= \left\lfloor \frac{a}{q} \right\rfloor, \ \ i \in \{0, 1, \cdots, p-1\}, \\ j &= (a)_q, \ \ j \in \{0, 1, \cdots, q-1\}, \end{aligned} \tag{8}$$

where $(\cdot)_q$ denotes the modulus-q operation. Let $\mathcal{A} = \{a_0, a_1, \ldots, a_{r-1}\} \subset \{0, 1, \ldots, N_s - 1\}$ be the set of r selected slot indices. Each slot index $a_z \in \mathcal{A}$ maps to coordinates (i_z, j_z) computed as

[3] In this paper, we use the term *slot* to refer specifically to an OTFS slot.

[4] We refer to this sequence of symbols as the user packet.

[5] This configuration resembles a point pilot surrounded by a guard band, which serves to prevent interference (between the pilot and data regions) caused by channel spreading.

Table 1 DD Filters [31]

DD Filters	$w_{tx}(\tau,\nu) = w_1(\tau)w_2(\nu)$
Sinc Filter	$w_1(\tau) \triangleq \sqrt{B}\,\mathrm{sinc}(B\tau)$
	$w_2(\nu) \triangleq \sqrt{T}\,\mathrm{sinc}(T\nu)$
Gaussian Filter	$w_1(\tau) \triangleq \left(\frac{2\alpha_\tau B^2}{\pi}\right)^{\frac{1}{4}} e^{-\alpha_\tau B^2\tau^2}$
	$w_2(\nu) \triangleq \left(\frac{2\alpha_\nu T^2}{\pi}\right)^{\frac{1}{4}} e^{-\alpha_\nu T^2\nu^2}$

$$i_z = \left\lfloor \frac{a_z}{q} \right\rfloor, \quad j_z = (a_z)_q,$$

where $z = 0, 1, \cdots, r-1$. The k_{a}th user packet $s_{k_{\mathrm{a}}}[\tilde{k}, \tilde{l}]$ is replicated across the r selected slots to effectively distribute the transmission across multiple DD resources. Therefore, the corresponding transmit signal over the DD frame, $x_{k_{\mathrm{a}}}[k,l]$, is expressed as

$$\begin{aligned} x_{k_a}[k,l] = &\sum_{z=0}^{r-1} s_{k_{\mathrm{a}}}[k - 2M_{\mathrm{tile}} i_z,\ l - N_{\mathrm{tile}} j_z] \\ &\cdot \mathbb{1}_{[2M_{\mathrm{tile}} i_z,\, 2M_{\mathrm{tile}}(i_z+1))}(k) \cdot \mathbb{1}_{[N_{\mathrm{tile}} j_z,\, N_{\mathrm{tile}}(j_z+1))}(l), \end{aligned} \tag{9}$$

where $\mathbb{1}_{[c,d)}(x)$ is the indicator function:

$$\mathbb{1}_{[c,d)}(x) = \begin{cases} 1 & \text{if } x \in [c,d) \\ 0 & \text{otherwise.} \end{cases}$$

The function $s_{k_{\mathrm{a}}}[m', n']$ is defined only for $0 \le m' < 2M_{\mathrm{tile}}$ and $0 \le n' < N_{\mathrm{tile}}$, and the values (i_z, j_z) ensure that $s_{k_{\mathrm{a}}}[m', n']$ is mapped to the desired slot. The replication strategy in Eq. (9) ensures that each user's transmission spans the entire OTFS frame, enhancing diversity and robustness against channel impairments such as delay and Doppler spreads. Following the Zak-OTFS signal processing described in Sect. 2.2, the user's DD domain signal $x_{k_{\mathrm{a}}}[k,l]$ is processed via Eqs. (1, 2, and 3) to produce the continuous transmit signal $x^{\mathrm{tx}}_{\mathrm{dd},k_{\mathrm{a}}}(\tau,\nu)$. In this paper, we employ two DD domain transmit filters $w_{\mathrm{tx}}(\tau,\nu)$, as detailed in Table 1[6].

At the receiver, we use a matched filter in the DD domain, which is defined as [31]

$$w_{rx}(\tau,\nu) = e^{j2\pi\nu\tau} w^*_{tx}(-\tau,-\nu). \tag{10}$$

Closed-form expressions for the effective channel $h_{\mathrm{eff}}[k,l]$ under each filter choice can be obtained using Table 1, Eqs. (7) and (10). We refer the reader to [29] for details.

[6] We do not consider the Gaussian-sinc filter [30] due to its mathematical complexity. Numerical results show that, for the proposed MMA scheme, its performance is close to the Sinc filter and not better than the Gaussian filter. Similarly, the RRC filter is omitted due to the lack of a tractable closed-form expression for $h_{\mathrm{eff}}[k,l]$ [29]

3.3 Slot-level signal processing at the receiver

Assuming all active users are frame-synchronized, they transmit within a shared uplink OTFS frame. Consequently, the received DD domain symbols at the BS, denoted by $y[k', l']$, for $k' = 0, 1, \cdots, M-1$ and $l' = 0, 1, \cdots, N-1$, consist of the superposition of the contributions from all K_{a} active users, i.e.,

$$\begin{aligned} y[k', l'] &= \sum_{k_{\mathrm{a}}=1}^{K_{\mathrm{a}}} \tilde{y}_{k_{\mathrm{a}}}[k', l'] + n[k', l']; \\ \tilde{y}_{k_{\mathrm{a}}}[k', l'] &\triangleq \sum_{k'', l'' \in \mathbb{Z}} h_{\mathrm{eff},k_{\mathrm{a}}}[k'-k'', l'-l''] x_{\mathrm{dd},k_{\mathrm{a}}}[k'', l''] e^{j2\pi\left(\frac{(l'-l'')k''}{M}\right)} + n[k', l']. \end{aligned} \tag{11}$$

Here, $\tilde{y}_{k_{\mathrm{a}}}[k', l']$ denotes the received contribution from the k_{a}th user, where $h_{\mathrm{eff},k_{\mathrm{a}}}[k', l']$ is the corresponding user-specific effective channel response, and $x_{\mathrm{dd},k_{\mathrm{a}}}[k', l']$ is the quasi-periodic DD domain signal transmitted by the k_{a}th user, for all $k', l' \in \mathbb{Z}$. At the receiver, signal processing is performed on a per-slot basis[7], where the channel response is estimated from the pilot tile and then used to equalize the symbols in the associated data tile. Assuming that the repetition pattern (or preambles) of all active users are known, the BS begins by identifying and processing the *singleton slots*, i.e., slots in which no collision occurs and only a single user transmits its packet[8].

Using the knowledge of the preambles, the BS identifies that the $\widetilde{k}_{\mathrm{a}}$th user's signal contribution, denoted as $\tilde{y}_{\widetilde{k}_{\mathrm{a}}}[k^{\dagger}, l^{\dagger}]$ for $k^{\dagger} = 0, 1, \cdots, 2M_{\mathrm{tile}} - 1$ and $l^{\dagger} = 0, 1, \cdots, N_{\mathrm{tile}} - 1$, is received within a singleton slot indexed by $z' \in \{0, 1, \ldots, N_s - 1\}$. The BS extracts the user's received symbols from the global received frame as

$$\begin{aligned} \tilde{y}_{\widetilde{k}_{\mathrm{a}}}[k^{\dagger}, l^{\dagger}] &= y[k^{\dagger} + r',\ l^{\dagger} + c'], \\ &\qquad \text{for } 0 \le k^{\dagger} < 2M_{\mathrm{tile}},\ 0 \le l^{\dagger} < N_{\mathrm{tile}}, \end{aligned} \tag{12}$$

where the slot coordinates are given by $i' = \left\lfloor \frac{z'}{q} \right\rfloor$, $j' = \left(z'\right)_q$, $r' = 2M_{\mathrm{tile}} \cdot i'$ and $c' = N_{\mathrm{tile}} \cdot j'$. Next, the BS isolates the pilot tile by selecting the subset of $\tilde{y}_{\widetilde{k}_{\mathrm{a}}}[k^{\dagger}, l^{\dagger}]$ with indices $0 \le k^{\dagger} < M_{\mathrm{tile}}$ and $0 \le l^{\dagger} < N_{\mathrm{tile}}$. Using the cross-ambiguity operation [35], the estimated I/O relationship for the $\widetilde{k}_{\mathrm{a}}$th user in the singleton slot is given by

$$\begin{aligned} \widehat{h}_{\mathrm{eff},\widetilde{k}_{\mathrm{a}}}[k^{\dagger}, l^{\dagger}] &= A_{\tilde{y}_{\widetilde{k}_{\mathrm{a}}}, s_{\widetilde{k}_{\mathrm{a}}}}[k^{\dagger}, l^{\dagger}] \\ &= \sum_{k_1=0}^{M_{\mathrm{tile}}-1} \sum_{l_1=0}^{N_{\mathrm{tile}}-1} \tilde{y}_{\widetilde{k}_{\mathrm{a}}}[k_1, l_1] s^{*}_{\mathrm{dd},\widetilde{k}_{\mathrm{a}}}\left[k_1 - k^{\dagger}, l_1 - l^{\dagger}\right] e^{-j2\pi \frac{l^{\dagger}(k_1 - k^{\dagger})}{MN}}, \end{aligned} \tag{13}$$

for $0 \le k^{\dagger} < M_{\mathrm{tile}},\ 0 \le l^{\dagger} < N_{\mathrm{tile}}$, where $s_{\mathrm{dd},\widetilde{k}_{\mathrm{a}}}[\tilde{k}, \tilde{l}]$ denotes the transmitted pilot signal of the $\widetilde{k}_{\mathrm{a}}$th user in the tile. In DD modulation systems such as OTFS 1.0 or Zak-OTFS (OTFS 2.0), all symbols within the frame experience nearly the same effective channel [25, 36]. Therefore, the channel estimated from the pilot tile can be reliably used to

[7] Per-slot decoding and DD domain analysis have been conducted in [32–34], primarily within the context of MC-OTFS/OTFS 1.0 under an orthogonal multiple access (OMA) framework. In contrast, this work addresses randomized transmissions in a grant-free setting.

[8] Alternatively, the BS may operate in a fully blind manner, without prior knowledge of users' repetition patterns. However, this approach typically incurs higher decoding complexity.

detect the data symbols in the associated data tile. The resulting I/O relationship in the data tile can be represented in matrix form as

$$\tilde{\mathbf{y}} = \mathbf{H}\mathbf{s} + \tilde{\mathbf{n}}, \tag{14}$$

where $\tilde{\mathbf{y}}, \mathbf{s}, \tilde{\mathbf{n}} \in \mathbb{C}^{M_{\text{tile}}N_{\text{tile}}\times 1}$ represent the received, transmitted, and noise vectors, respectively. $\tilde{y}_{\tilde{k}_a}[k^\dagger, l^\dagger]$ and $n[k^\dagger, l^\dagger]$ correspond to the $((k^\dagger - M_{\text{tile}})N + l^\dagger + 1)$th entries of $\tilde{\mathbf{y}}$ and $\tilde{\mathbf{n}}$, respectively, for $M_{\text{tile}} \le k^\dagger < 2M_{\text{tile}}$ and $0 \le l^\dagger < N_{\text{tile}}$. Similarly, $s_{\tilde{k}_a}[\tilde{k}, \tilde{l}]$ corresponds to the $((\tilde{k} - M_{\text{tile}})N + \tilde{l} + 1)$th entry of $\mathbf{s}$ for $M_{\text{tile}} \le \tilde{k} < 2M_{\text{tile}}$ and $0 \le \tilde{l} < N_{\text{tile}}$. Also, $\mathbf{H} \in \mathbb{C}^{M_{\text{tile}}N_{\text{tile}}\times M_{\text{tile}}N_{\text{tile}}}$ is the effective channel matrix, whose entries $H[(k^\dagger - M_{\text{tile}})N + l^\dagger + 1, (\tilde{k} - M_{\text{tile}})N + \tilde{l} + 1]$ are derived as in Mohammed et al. [25]. The minimum mean squared error (MMSE) equalization of the received vector $\tilde{\mathbf{y}}$ is performed as

$$\begin{aligned} \widehat{\mathbf{s}} &= \mathbf{W}_{\text{MMSE}}\ \tilde{\mathbf{y}}; \\ \mathbf{W}_{\text{MMSE}} &\triangleq \left(\widehat{\mathbf{H}}\widehat{\mathbf{H}}^H + \frac{1}{E_{\text{d}}}\mathbf{R}_n\right)^{-1}\widehat{\mathbf{H}}^H, \end{aligned} \tag{15}$$

where $E_{\text{d}} = \mathbb{E}[|s_{\text{dd},\tilde{k}_a}[\tilde{k}, \tilde{l}]|^2]$ is the average energy per information symbol, $\widehat{\mathbf{H}}$ is the estimated channel matrix obtained from the effective channel response in eq. (13), and $\mathbf{R}_n$ is the noise covariance matrix, whose entries depend on the receiver filter (see [29] for derivation). The estimated symbol vector $\widehat{\mathbf{s}}$ is then decoded using minimum Euclidean distance detection [37], yielding the estimated transmitted symbols $\check{s}_{\tilde{k}_a}[\widetilde{k}, \widetilde{l}]$, for $M_{\text{tile}} \le \widetilde{k} < 2M_{\text{tile}}$ and $0 \le \widetilde{l} < N_{\text{tile}}$. Using the decoded symbols $\check{s}_{\tilde{k}_a}[\tilde{k}, \tilde{l}]$, the channel estimate $\widehat{h}_{\text{eff},\tilde{k}_a}[k, l]$ from (13), and the Zak-OTFS I/O relationship in (6), the contribution of the $\widetilde{k}_a$th user to the received signal is estimated as

$$\widehat{y}_{\tilde{k}_a}[k^\dagger, l^\dagger] \approx \sum_{k,l\in\mathbb{Z}} \widehat{h}_{\text{eff},\tilde{k}_a}[k^\dagger - k, l^\dagger - l]\check{s}_{\text{dd},\tilde{k}_a}[k, l]e^{j2\pi\left(\frac{(l^\dagger - l)k}{M}\right)}. \tag{16}$$

From the received signal model in eq. (11) and the reconstructed signal in eq. (16), the BS performs SIC by removing the contribution of the $\widetilde{k}_a$th user from all its occupied slots

$$\begin{aligned} y^{(1)}[k', l'] = y[k', l'] - \sum_{z'=0}^{r-1} \widehat{y}_{\tilde{k}_a}[k' - 2M_{\text{tile}}i_{z'},\ l' - N_{\text{tile}}j_{z'}] \\ \cdot \mathbb{1}_{[2M_{\text{tile}}i_{z'},\, 2M_{\text{tile}}(i_{z'}+1))}(k') \cdot \mathbb{1}_{[N_{\text{tile}}j_{z'},\, N_{\text{tile}}(j_{z'}+1))}(l'). \end{aligned} \tag{17}$$

We remark that the SIC processing is performed at the PHY layer, as it directly uses the decoded payload and the estimated channel to cancel the user's contribution from the overall DD domain frame. After applying (17), the BS re-evaluates the frame to identify newly formed singleton slots. The decoding and SIC procedures are repeated iteratively, each time canceling the contributions of successfully decoded users, until either all user packets have been decoded or no further singleton slots remain. The complete transceiver processing of the proposed Zak-OTFS-based MMA scheme is summarized

Table 2 Transceiver processing steps of the proposed Zak-OTFS-based MMA scheme

Transceiver operation	Description
User Packet Construction	Form user packet $s_{k_a}[\widetilde{k},\widetilde{l}]$ with a central pilot symbol in the pilot tile and modulated data symbols in the data tile
Slot Repetition Mapping	Replicate $s_{k_a}[\widetilde{k},\widetilde{l}]$ across selected slot indices to form the full-frame DD domain signal $x_{k_a}[k,l]$ (see 9)
DD Domain Filtering (Tx)	Apply the transmit filter $w_{\text{tx}}(\tau,\nu)$ in the DD domain to generate the continuous-time signal $x^{\text{tx}}_{\text{dd},k_a}(\tau,\nu)$ (see 3)
Channel Propagation	The signal propagates through a doubly spread channel with an effective response $h_{\text{eff},k_a}[\widetilde{k},\widetilde{l}]$ (see 4)
DD Domain Filtering (Rx)	Apply the matched filter $w_{\text{rx}}(\tau,\nu)$ at the receiver (see 5, 10)
DD Domain Sampling	Sample the received signal in the DD domain to obtain $y[k',l']$, representing the superimposed contributions from all active users (see 11)
Slot Extraction	Identify the singleton slot z', and extract the received signal $\widetilde{y}_{\widetilde{k}_a}[k^\dagger,l^\dagger]$ based on the known repetition patterns (see 12)
Channel Estimation	Obtain an estimate of the effective channel $\widehat{h}_{\text{eff},\widetilde{k}_a}[k^\dagger,l^\dagger]$ through cross-ambiguity processing (see 13)
MMSE Equalization	Apply MMSE data equalization as $\widehat{\mathbf{s}} = \mathbf{W}_{\text{MMSE}}\,\widetilde{\mathbf{y}}$ (see 15)
Symbol Decoding	Perform minimum distance decoding on $\widehat{\mathbf{s}}$ to obtain the estimated symbols $\widetilde{s}_{\widetilde{k}_a}[\widetilde{k},\widetilde{l}]$
Signal Reconstruction	Reconstruct the user contribution $\widehat{y}_{\widetilde{k}_a}[k',l']$ by combining the decoded symbols with the channel estimate (see 16)
Successive Interference Cancellation	Subtract the reconstructed signals from $y[k',l']$ and iteratively decode singleton slots until none remains (see 17)

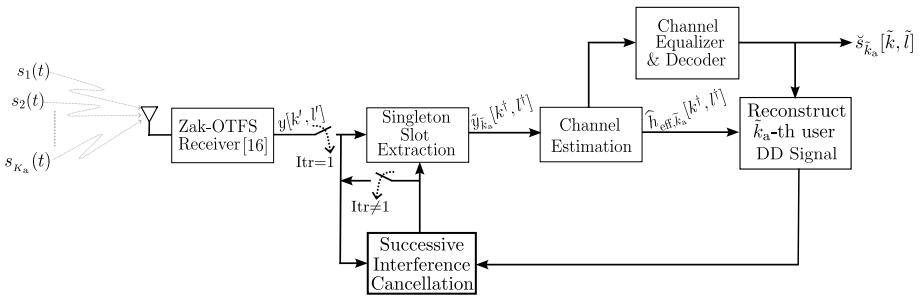

Fig. 3 Receiver block diagram of the proposed Zak-OTFS coded random access scheme, illustrating the processing from signal reception to user decoding with SIC

in Table 2. Moreover, Fig. 3 provides a pictorial representation of the receiver block diagram, illustrating all slot-level signal processing performed at the BS, from time domain signal reception to the user decoding and SIC operations.

3.4 OFDM processing

In this work, we select cyclic prefix (CP)-OFDM with pilot insertion as a benchmark scheme to compare against the proposed Zak-OTFS approach. To ensure a fair comparison, the CP-OFDM system is configured with a subcarrier spacing of $\Delta f = \nu_p$, and a CP length matched to the maximum delay spread of the channel. To enhance robustness against Doppler effects, pilot and data symbols are interleaved across alternating subcarriers. Under this configuration, a *slot*, as defined consistently with Zak-OTFS, spans

$M_{\text{sub}} = 2M_{\text{tile}}$ subcarriers and $N_{\text{OFDM}} = N_{\text{tile}}$ time indices. The resulting CP-OFDM slot structure is illustrated in Fig. 2a, in direct parallel to its Zak-OTFS counterpart.

To ensure parity with Zak-OTFS in terms of overall resource usage, the CP-OFDM system is configured with $N = p \cdot q \cdot N_{\text{OFDM}}$ time indices per frame, matching Zak-OTFS in terms of time–bandwidth product. Note that this configuration follows the structure of conventional CRA schemes based on time slot division. Decoding at the BS begins after the reception of all N time indices, aligning with the frame-level processing strategy of Zak-OTFS. Leveraging the known preamble, the BS first identifies singleton slots, i.e., those associated with the $\widetilde{k}_{\text{a}}$th active user free from collision. Channel estimation is first performed using the pilot symbols, followed by MMSE-based interpolation. Subsequently, data detection is carried out via MMSE equalization.

Similar to the iterative, PHY layer SIC approach used in Zak-OTFS, the BS exploits the estimated channels and the correlation structure [38] to predict the contributions of each $\widetilde{k}_{\text{a}}$th user's signal across other slots containing replicas. These contributions are iteratively canceled from the received signal, progressively revealing additional singleton slots. This process continues until all user data have been decoded or no further singleton slots can be identified. In this way, CP-OFDM follows a decoding paradigm closely aligned with that of Zak-OTFS, enabling a fair performance comparison under common assumptions. Accordingly, our focus is on evaluating each modulation scheme's ability to support channel prediction across data carriers. We remark that the comparison between Zak-OTFS and OFDM is absolutely fair, since both systems are configured with the same time–bandwidth product, transmit power, and number of information symbols per frame, ensuring identical spectral efficiency and resource usage. Therefore, any observed performance differences stem solely from the intrinsic properties of the modulation schemes.

4 Performance evaluation

4.1 Transmission setup

Each user encodes a k-bit payload using a (n, k, t) binary Bose–Chaudhuri–Hocquenghem (BCH) code capable of correcting up to t errors. A portion of the k bits is reserved for a cyclic redundancy check (CRC), which is used to validate the decoded packet. The resulting code word is padded with one zero bit and mapped to a quadrature phase-shift keying (QPSK) constellation with Gray mapping. Two BCH code configurations are used depending on the slot size: *i*) A $(n = 31, k = 16, t = 3)$ code for small slots ($M_{\text{tile}} = 4$, $N_{\text{tile}} = 4$); *ii*) A $(n = 511, k = 255, t = 31)$ code for larger slots ($M_{\text{tile}} = 16$, $N_{\text{tile}} = 16$). In both cases, the code rate is approximately $R_{\text{c}} \approx 0.5$. Furthermore, to ensure fairness in the overall use of resources between Zak-OTFS and OFDM, the OFDM subcarrier spacing is set to $\Delta f = \nu_p$ and the OFDM symbol duration to $T_{\text{sym}} = \tau_p$, matching the DD domain slot's time and frequency span. Two different frame lengths are considered: $N_{\text{s}} = 128$ and $N_{\text{s}} = 200$ slots. For the Zak-OTFS-based scheme, the $N_{\text{s}} = 128$ configuration is arranged as a $p = 8$ by $q = 16$ grid in the DD domain, while the $N_{\text{s}} = 200$ configuration uses a $p = 10$ by $q = 20$ grid, in both cases preserving a square-shaped frame structure.

The pilot SNR per slot is given as

$$\mathrm{SNR_P} = \frac{E_\mathrm{P}}{N_0 BT}, \tag{18}$$

where $E_\mathrm{P} = E_\mathrm{p} \cdot N_\mathrm{P}$ is the pilot energy per user per slot, E_p is the energy per pilot symbol, N_P is the number of pilot symbols per slot, N_0 is the noise power spectral density, B is the frame bandwidth, and T is the frame duration. In Zak-OTFS, only one pilot symbol is transmitted in the pilot tile ($N_\mathrm{P} = 1$), while in OFDM, all symbols in the pilot tile are used as pilots ($N_\mathrm{P} = M_\mathrm{tile} \cdot N_\mathrm{tile}$). Similarly, the data SNR per slot is given as

$$\mathrm{SNR_D} = \frac{E_\mathrm{D}}{N_0 BT}, \tag{19}$$

where $E_\mathrm{D} = E_\mathrm{d} \cdot N_\mathrm{D}$ is the data energy per user per slot, E_d is the energy per data symbol, and N_D is the total number of data symbols per slot. For both Zak-OTFS and OFDM configurations, the entire data slot region is dedicated to payload transmission, thus $N_\mathrm{D} = M_\mathrm{tile} \cdot N_\mathrm{tile}$. The total data and pilot energy budgets are kept equal ($E_\mathrm{P} = E_\mathrm{D}$), ensuring that $\mathrm{SNR_P} = \mathrm{SNR_D} = \mathrm{SNR}$ for both CRA schemes (Zak-OTFS and OFDM).

4.2 Channel model

We adopt the six-path Veh-A channel model [39] with the power delay profile shown in Table 3. We normalized the path energies such that $\sum_{i=1}^{6} \mathbb{E}[|h_i|^2] = 1$. The ith path Doppler shift is modeled as $\nu_i = \nu_\mathrm{max} \cos(\theta_i)$, where θ_i is an independent and identically distributed (i.i.d.) random variable uniformly distributed over $[0, 2\pi)$, and $\nu_\mathrm{max} = 815$ Hz is the maximum Doppler shift assumed throughout the simulations.

4.3 Numerical results

We organize the numerical results into three parts to evaluate the proposed uplink MMA scheme from different perspectives: (1) We examine an ideal interference-free scenario, where a single user transmits in a single slot without collisions for the proposed Zak-OTFS-based receiver operating in the DD domain and the conventional OFDM-based receiver in the TF domain (see Fig. 2a and b). (2) We investigate the impact of channel variability on the effectiveness of SIC for both Zak-OTFS and OFDM-based CRA schemes (see Sect. 3). (3) We present a frame-level performance analysis based on Monte Carlo simulations of the full MMA uplink system, for both the Zak-OTFS and the conventional OFDM-based CRA schemes. For cases (1) and (2), we evaluate performance in terms of PLR versus SNR, while for case (3), we assess the PLR as a function of the number of active users per frame, denoted by K_a.

Table 3 Power delay profile of the Veh-A channel [39]

Path number *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

4.3.1 Single-user performance analysis

The goal is to analyze the influence of key system parameters, including the number of delay and Doppler bins per tile ($M_{\text{tile}}, N_{\text{tile}}$), the Doppler period ν_p, and the pulse shaping filter.

Figure 4 shows the results for a compact configuration with $M_{\text{tile}} = N_{\text{tile}} = 4$, leading to slot sizes $M_{\text{slot}} = 8$ and $N_{\text{slot}} = 4$ for Zak-OTFS, and $M_{\text{sub}} = 8$ subcarriers and $N_{\text{OFDM}} = 4$ OFDM symbols for the OFDM-based scheme. The Doppler period is fixed to $\nu_p = 30$ kHz, aligned with the OFDM subcarrier spacing Δf, and a (31, 16, 3) BCH code is adopted. In this configuration, for the Veh-A channel considered, we observe that Zak-OTFS suffers from limited resources in the DD domain, which degrades channel estimation accuracy and results in a pronounced PLR floor between 10^{-1} and 10^{-2} for both pulse shapes. In contrast, OFDM benefits from a coarse subcarrier spacing, which mitigates Doppler-induced inter-carrier interference (ICI) and yields superior reliability.

Figure 5 presents results for a larger configuration, with $M_{\text{tile}} = N_{\text{tile}} = 16$, resulting in slot sizes $M_{\text{slot}} = 32$ and $N_{\text{slot}} = 16$. A (511, 255, 31) BCH code is used, and two values of Doppler period are considered, namely $\nu_p = 5$ kHz and $\nu_p = 30$ kHz. For $\nu_p = 30$ kHz (solid curves), all schemes benefit from the improved resolution. OFDM achieves extremely low PLR for SNR ≥ -5 dB. Zak-OTFS also shows marked improvements compared to the small slot case. Sinc pulse shaping outperforms the Gaussian one at medium SNR, but still exhibits an error floor around 10^{-2} due to residual channel estimation errors. Conversely, the Gaussian pulse, although requiring a higher SNR to enter the waterfall region, provides superior performance at high SNR, thanks to its improved channel estimation capabilities. This, however, comes at the cost of increased equalization complexity and a more gentle PLR slope.

When the Doppler period is reduced to $\nu_p = 5$ kHz (dashed curves), OFDM performance degrades significantly. The reduced subcarrier spacing leads to more severe ICI, causing a PLR floor between 10^{-1} and 10^{-2}. Zak-OTFS, by contrast, maintains stable and reliable performance for both pulse shapes, as the selected Doppler and delay period

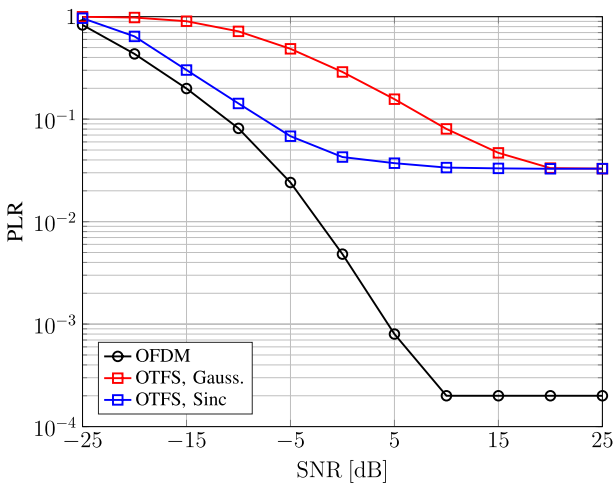

Fig. 4 PLR versus SNR for a single-user, interference-free scenario with slot configuration $M_{\text{slot}} = 8, N_{\text{slot}} = 4$ and Doppler period $\nu_p = 30$ kHz. The performance of Zak-OTFS with Gaussian and sinc pulse shaping is compared to conventional OFDM using a (31, 16, 3) BCH code

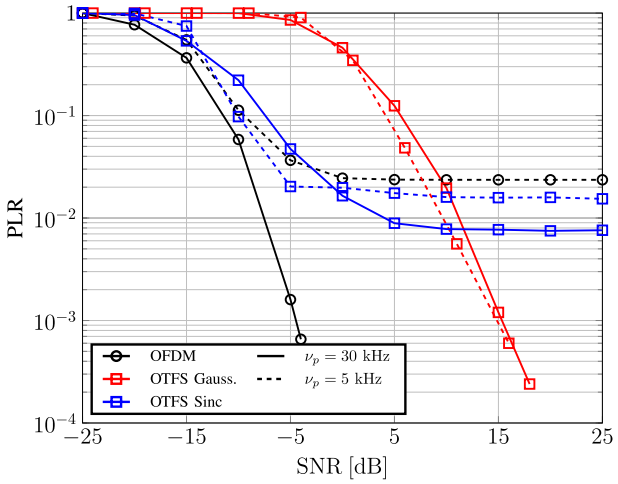

Fig. 5 PLR versus SNR in a single-user, interference-free scenario with slot configuration $M_{\text{slot}} = 32$, $N_{\text{slot}} = 16$. Two Doppler period values are considered: $\nu_p = 30$ kHz (solid lines) and $\nu_p = 5$ kHz (dashed lines). A (511, 255, 31) BCH code is employed. Zak-OTFS performance with Gaussian and sinc pulses is compared to conventional OFDM

parameters ensure operation within the crystalline regime. In particular, the conditions $\nu_p > 2\nu_{\max} = 1630$ Hz and $\tau_p = 1/\nu_p = 200\ \mu s > \tau_{\max} = 2.51\ \mu s$ guarantee that the channel interaction remains structured and predictable.[9] This states that OFDM suffers from ICI with a lower subcarrier spacing, but the Zak-OTFS performance remains almost constant, under crystalline condition, even at a lower resolutions.

Across both scenarios, a noticeable SNR gap emerges between Zak-OTFS and OFDM in the waterfall region, particularly when the Gaussian filter is employed. Nonetheless, these results highlight that Zak-OTFS, especially when operated with sufficiently high resolution, can achieve reliable performance even under severe mobility conditions. In contrast, OFDM remains competitive in static or low-mobility environments, thanks to its efficient equalization and simplified signal processing.

4.3.2 Impact of channel variability on SIC performance

We now present a simplified numerical experiment to illustrate how rapidly time-varying channels affect the performance of PHY layer SIC under conventional OFDM-based processing, and how DD domain processing using Zak-OTFS enables a more robust and effective application of SIC. The scenario features a transmission frame and two active users: an "uncollided" user and a "collided" user. The uncollided user is assumed to have at least one replica transmitted without interference and at least one overlapping with other transmissions, e.g., the yellow user in Fig. 1a, which has a non-interfered replica in slot 2 and an interfered one in slot 6. In contrast, the collided user has no replica transmitted alone, and its packet can only be recovered via SIC, after decoding and subtracting the contributions of interfering packets from the uncollided user, as shown for the blue user in Fig. 1a, which is interfered in both slots 6 and $N_s - 1$.

[9] $2\nu_{\max}$ and $\tau_{\max}$ denote the Doppler and delay spreads of the channel, respectively.

Focusing on the collided user, successful decoding requires effective PHY layer SIC in the slots affected by collisions. This, in turn, demands reliable estimation of the interfering uncollided user's channel not only in the slot where it was decoded, but also in all other slots where its replicas were transmitted. For example, if the yellow user is decoded in slot 2, its channel must also be estimated in slot 6 to enable interference cancellation. However, due to the presence of interference, it is generally not possible to directly estimate the channel via pilot symbols in the collided slot. As a result, the BS must resort to MMSE-based channel prediction, initialized from the reliable estimate obtained in the decoded slot, and attempt to extrapolate the channel evolution across time to the remaining replicas.

Under high-mobility conditions, however, conventional OFDM-based processing faces substantial challenges due to rapid channel variations across slots, which inherently reduces the channel coherence time. These variations degrade the accuracy of the predicted channel estimates, especially as the temporal distance between slots increases. Consequently, the farther a replica lies from the slot where the channel was estimated, the less reliable the prediction becomes. This degradation significantly impairs the performance of SIC, resulting in residual errors, and lowering the probability of successfully recovering collided user packets. In contrast, Zak-OTFS-based processing in the DD domain inherently mitigates rapid time variations. In this domain, the channel remains approximately static across the entire frame. Thus, once a user is successfully decoded, its channel estimate can be reliably reused in all other slots, allowing for accurate and consistent SIC application, even in high-mobility environments.

To illustrate this effect numerically, we consider a small slot configuration analogous to Fig. 4, with $M_{\text{slot}} = 8$, $N_{\text{slot}} = 4$, and Doppler period $\nu_p = 30$ kHz. Figure 6 reports the PLR versus SNR for both uncollided and collided users, comparing OFDM-based and Zak-OTFS-based CRA schemes, with Gaussian and sinc pulse shaping filters applied

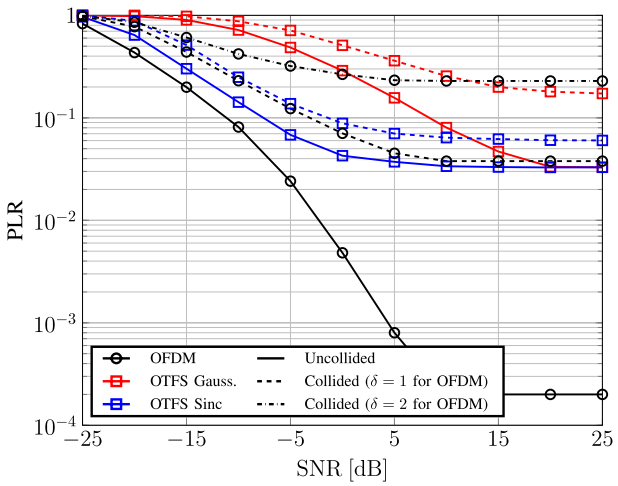

Fig. 6 PLR versus SNR for a simplified two-user scenario with $M_{\text{slot}} = 8$, $N_{\text{slot}} = 4$ and Doppler period $\nu_p = 30$ kHz. The performance of both uncollided and collided users is shown for Zak-OTFS (with Gaussian and sinc pulse shaping) and conventional OFDM. For OFDM, results are reported for different values of δ, the temporal distance between the reference and the collided slots, highlighting the degradation in SIC performance due to channel prediction errors under high-mobility conditions

to the latter. The uncollided user performance replicate those in Fig. 4. In the OFDM-based case, we introduce the variable δ, which denotes the temporal distance, measured in slot indices, between the slot where the uncollided user's packet is decoded and the slot where its replicas must be canceled to enable decoding of the collided user (e.g., in the example mentioned above, $\delta = 4$). Since channel prediction quality deteriorates with time, the PLR increases as δ grows. We report results for $\delta = 1$ (adjacent slots) and $\delta = 2$, assuming a maximum Doppler frequency of $\nu_{\max} = 815$ Hz. The results in Fig. 6 reveal a clear degradation in SIC performance with increasing δ, confirming that OFDM suffers from time selectivity. In contrast, Zak-OTFS maintains a consistently low PLR for the collided user, despite its slightly inferior performance for the uncollided case in this setting.

In Fig. 7, we also consider a larger slot configuration, as in Fig. 5, with $M_{\text{slot}} = 32$, $N_{\text{slot}} = 16$, and $\nu_p = 30$ kHz. The PLR for the collided user in this setup corresponds to the solid curves in Fig. 5. For the OFDM-based CRA scheme, SIC performance deteriorates dramatically even for $\delta = 1$, due to the longer slot duration. As a result, channel prediction becomes unreliable, effectively disabling SIC and preventing recovery of the collided packet. Also, for Zak-OTFS with Gaussian pulse shaping, the collided user experiences only minor performance degradation compared to the uncollided case, owing to accurate channel estimation and effective interference cancellation. On the other hand, with sinc pulses, even if the PLR remain within acceptable bounds with an error floor between 10^{-1} and 10^{-2} the performance degrades more significantly due to its poorer estimation capabilities, which compromise the effectiveness of SIC.

4.3.3 Frame-level performance of the MMA uplink system

In this final section, we evaluate the frame-level performance of the overall MMA uplink system under high-mobility conditions. The simulations are conducted using the larger slot configuration with $M_{\text{slot}} = 32$, $N_{\text{slot}} = 16$, and a Doppler period of $\nu_p = 30$ kHz at a

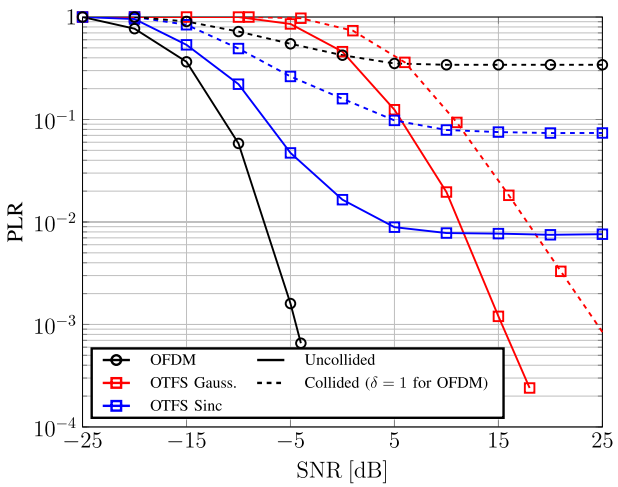

Fig. 7 PLR versus SNR for a two-user scenario with slot configuration $M_{\text{slot}} = 32$, $N_{\text{slot}} = 16$ and Doppler period $\nu_p = 30$ kHz. The PLR for collided users is shown for Zak-OTFS (using Gaussian and sinc pulse shaping) and conventional OFDM. Results highlight the severe degradation of SIC in the OFDM-based scheme, while Zak-OTFS maintains robust performance thanks to channel stability in the DD domain

SNR of 25 dB. Each user transmits $r = 3$ replicas, as often assumed in CRA literature [40, 41]. Figure 8 shows the PLR as a function of the number of simultaneously active users per frame, K_a, for a frame of $N_s = 128$ slots, for both the conventional OFDM-based and the proposed Zak-OTFS-based CRA schemes. For Zak-OTFS, results are reported for both Gaussian and sinc pulse shaping filters.

As for the OFDM-based scheme, the performance is reported under the assumption that no SIC is applied. As discussed above and validated via simulations, the application of SIC is ineffective in this high-mobility regime. The strong temporal channel variability renders inter-slot channel prediction highly inaccurate (i.e., the channel coherence time is approximately equal to the slot duration), preventing the recovery of collided packets. Therefore, the displayed results correspond to the DSA baseline, where performance benefits stem only from packet repetition. However, this approach allows for low PLR only when a small number of users is active, and rapidly saturates as K_a increases.

In contrast, the Zak-OTFS-based CRA scheme exhibits strong robustness to high mobility, enabling effective SIC and successful decoding of multiple user packets. As a result, it achieves much better scalability. In particular, when using the Gaussian pulse shaping filter, the Zak-OTFS system is able to support a significantly higher number of users, up to around $K_a = 60$, while maintaining a PLR below 10^{-3}. The observed error floor is mainly attributed to so-called *unsolvable collisions*, i.e., events where two or more users fully overlap in all their replica transmissions, leaving no opportunity for successful decoding. The Zak-OTFS scheme employing the sinc pulse filter shows both a higher error floor and a visible performance gap in the waterfall region compared to the Gaussian-filtered counterpart. This degradation is due to the reduced accuracy of channel estimation with the sinc filter, an issue clearly observable in Fig. 7. The poorer estimation quality directly impacts SIC effectiveness, leading to a degradation in overall performance.

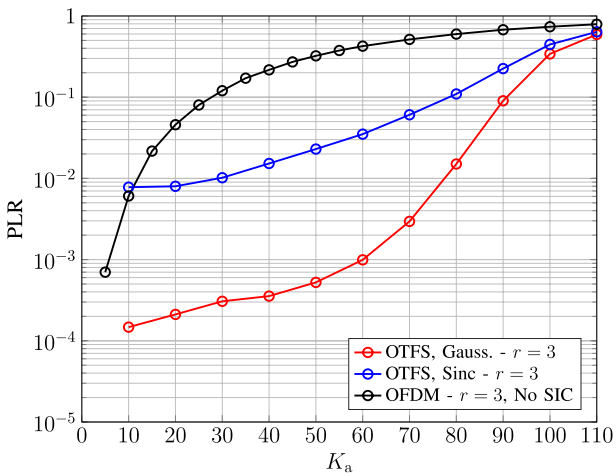

Fig. 8 PLR versus number of active users per frame, K_a, in a full-frame MMA uplink scenario with $N_s = 128$ slots, slot configuration $M_{slot} = 32$, $N_{slot} = 16$, and Doppler period $\nu_p = 30$ kHz. Each user transmits $r = 3$ replicas of the same packet. The SNR is fixed to 25 dB. Zak-OTFS-based CRA performance is shown for Gaussian and sinc pulse shaping, and compared against the OFDM-based scheme with no SIC

To better highlight the scalability of our scheme and its ability to support a large number of users (a fundamental aspect of mMTC services), Fig. 9 reports the frame-level uplink performance for both the OFDM-based and Zak-OTFS-based schemes (with Gaussian and sinc pulse shaping) under two configurations: $N_s = 128$ slots (as previously shown in Fig. 8) and $N_s = 200$ slots. For the OFDM-based scheme, performance remains limited due to the poor effectiveness of SIC, which becomes largely ineffective under conventional TF processing in the presence of high channel variability. Consequently, the best results are still obtained with DSA, where increasing the number of slots yields only a slight improvement.

In contrast, the Zak-OTFS-based scheme benefits significantly from a larger frame size, as more users can be successfully decoded at the slot level, enabling more reliable interference cancellation through SIC. For example, at a target PLR of 10^{-2}, typical for mMTC scenarios, the Zak-OTFS-based CRA scheme with Gaussian pulse shaping supports approximately $K_a \simeq 75$ users for $N_s = 128$ slots (red circle curve), and approximately $K_a \simeq 125$ users for $N_s = 200$ slots (red squared curve), corresponding to a scalability boost of roughly 55%. For sinc pulse shaping, performance remains limited in both configurations due to channel estimation constraints. Nevertheless, a noticeable scalability improvement is observed at a moderate PLR of 10^{-1}, yielding a boost of approximately 50%.

5 Conclusion

In this work, we introduced a novel CRA scheme operating in the DD domain and leveraging Zak-OTFS modulation for uplink MMA. The proposed framework takes advantage of the predictable structure of Zak-OTFS modulation to enable effective PHY layer SIC, maintaining robust performance even in the presence of high user mobility, conditions under which conventional OFDM-based CRA schemes suffer significant degradation and residual SIC errors. Through extensive numerical evaluations, we demonstrated

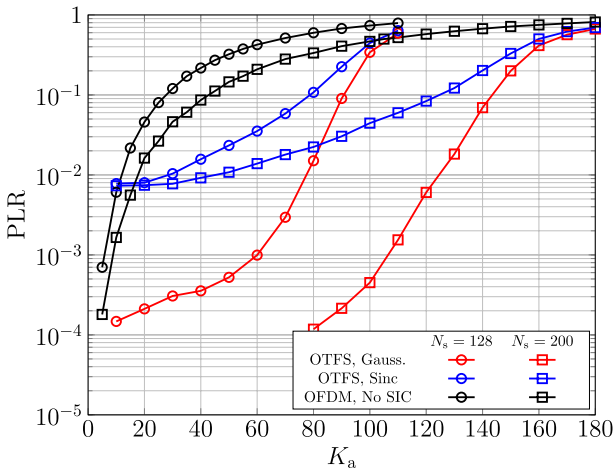

Fig. 9 PLR versus number of active users per frame, K_a, in a full-frame MMA uplink scenario with $N_s \in \{128, 200\}$ slots, slot configuration $M_{slot} = 32$, $N_{slot} = 16$, and Doppler period $\nu_p = 30$ kHz. Each user transmits $r = 3$ replicas of the same packet. The SNR is fixed to 25 dB. Zak-OTFS-based CRA performance is shown for Gaussian and sinc pulse shaping, and compared against the OFDM-based scheme with no SIC

that the proposed solution offers notable gains in scalability and reliability, consistently achieving lower packet loss rates across a broad spectrum of user densities. These improvements are further amplified by increasing the frame size, aligning with the requirements of typical mMTC scenarios.

As part of future work, we plan to extend and integrate the proposed framework with additional components aimed at further enhancing overall system performance. For instance, novel algorithms for multiple preamble detection should be developed to reconstruct users' repetition patterns more effectively. Moreover, advanced multi-packet reception (MPR) capabilities can be achieved through the deployment of a massive MIMO BS, which provides additional spatial degrees of freedom to enable user separation and channel estimation even in collided slots. We also aim to investigate alternative unsourced random access schemes, such as coded compressed sensing and random spreading techniques beyond CRA, and study their interplay with Zak-OTFS modulation.

This broader exploration will contribute to a more comprehensive understanding of how Zak-OTFS can influence the design and performance of future communication systems, ultimately supporting the development of more scalable, efficient, and robust next-generation wireless networks.

Author Contributions

R.C. and K.N. conceived this work. A.M. and V.K. made substantial contributions to the conception and design of the work. They also contributed to the analysis and interpretation of data, as well as to the creation of new software used in the work. In addition, A.M. and V.K. were primarily responsible for drafting and revising the manuscript. B.D. assisted A.M. and V.K. with data interpretation and software development. E.P., K.N., and R.C. reviewed and proofread the manuscript. All authors read and approved the final version of the manuscript.

Funding

The Duke team is supported in part by the National Science Foundation under grants 2342690 and 2148212, and is supported in part by funds from federal agency and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program. The Duke team is also supported in part by the Air Force Office of Scientific Research under grants FA 8750-20-2-0504 and FA 9550-23-1-0249. Texas A&M team is supported in part by the National Science Foundation under grant CNS-2148354. University of Bologna team is supported by the European Union - Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART").

Declarations

Competing interests

The authors declare no conflict of interest.

Received: 25 July 2025 Accepted: 17 December 2025
Published online: 07 January 2026

References

1. C. Bockelmann, N. Pratas, H. Nikopour, K. Au, T. Svensson, C. Stefanovic, P. Popovski, A. Dekorsy, Massive machine-type communications in 5G: physical and MAC-layer solutions. IEEE Commun. Mag. **54**(9), 59–65 (2016)
2. C. Bockelmann et al., Towards massive connectivity support for scalable mMTC communications in 5G networks. IEEE Access **6**, 28 969-28 992 (2018)
3. N. H. Mahmood et al., White paper on critical and massive machine type communication towards 6G, 6G Research Visions, (11), (2020)
4. Y. Wu, X. Gao, S. Zhou, W. Yang, Y. Polyanskiy, G. Caire, Massive access for future wireless communication systems. IEEE Wireless Commun. **27**(4), 148–156 (2020)
5. X. Chen, D.W.K. Ng, W. Yu, E.G. Larsson, N. Al-Dhahir, R. Schober, Massive access for 5G and beyond. IEEE J. Sel. Areas Commun. **39**(3), 615–637 (2021)
6. E. Casini, R. De Gaudenzi, O. del Rio Herrero, Contention resolution diversity slotted ALOHA (CRDSA): an enhanced random access scheme for satellite access packet networks. IEEE Trans. Wireless Commun. **6**(4), 1408–1419 (2007)

7. G. Liva, Graph-based analysis and optimization of contention resolution diversity slotted ALOHA. IEEE Trans. Commun. **59**(2), 477–487 (2011)
8. E. Paolini, G. Liva, M. Chiani, Coded slotted ALOHA: a graph-based method for uncoordinated multiple access. IEEE Trans. Inf. Theory **61**(12), 6815–6832 (2015)
9. M. Berioli, G. Cocco, G. Liva, A. Munari, Modern random access protocols. Found. Trends Netw. **10**(4), 317–446 (2016)
10. Y. Polyanskiy, A perspective on massive random-access, In: 2017 IEEE International Symposium on Information Theory , Aachen, Germany, pp. 2523–2527 (2017)
11. L.G. Roberts, ALOHA packet system with and without slots and capture. ACM SIGCOMM Comput. Commun. Rev. **5**(2), 28–42 (1975)
12. A. Vem, K. Narayanan, J.-F. Chamberland, J. Cheng, A user-independent successive interference cancellation based coding scheme for the unsourced random access Gaussian channel. IEEE Trans. Commun. **67**(12), 8258–8272 (2019)
13. V. Tralli, E. Paolini, IRSA-based random access over the Gaussian channel. IEEE Trans. Inf. Theory **70**(6), 4117–4139 (2024)
14. F. Clazzer, E. Paolini, I. Mambelli, Č. Stefanović, Irregular repetition slotted ALOHA over the Rayleigh block fading channel with capture, In: 2017 IEEE International Conference on Communications, Paris, France, (2017)
15. G. Liva, E. Paolini, Č. Stefanović, A. Graell i Amat, Coded slotted Aloha over the on-off fading channel: performance bounds, In: 2019 53rd Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, USA, (2019)
16. S.K. Mohammed, R. Hadani, A. Chockalingam, R. Calderbank, OTFS-A mathematical foundation for communication and radar sensing in the delay-doppler domain. IEEE BITS Inf. Theory Mag. **2**(2), 36–55 (2022)
17. G.L. Choudhury, T.S. Rappaport, Diversity ALOHA-A random access scheme for satellite communications. IEEE Trans. Commun. **31**(3), 450–457 (1983)
18. I. Hmedoush, J. Cao, C. Adjih, S. Sharma, K. Deka, *An Introduction to Modern Random Access Protocols for IoT Communications* (Springer, Singapore, 2025), pp.403–427. https://doi.org/10.1007/978-981-96-3185-8_13
19. L. Valentini, M. Chiani, E. Paolini, Massive grant-free access with massive MIMO and spatially coupled replicas. IEEE Trans. Commun. **70**(11), 7337–7350 (2022)
20. L. Valentini, M. Chiani, E. Paolini, Interference cancellation algorithms for grant-free multiple access with massive MIMO. IEEE Trans. Commun. **71**(8), 4665–4677 (2023)
21. L. Valentini, A. Mirri, E. Paolini, Feedback-aided coded random access with intentional power unbalance. IEEE Trans. Commun. **73**(1), 230–244 (2025)
22. A. Mirri, L. Valentini, I. Leyva-Mayorga, M. Chiani, E. Paolini, P. Popovski, Coded random access schemes for critical mMTC with multiple latency deadlines, IEEE Trans. Commun., (2025), early access
23. E. Björnson, J. Hoydis, L. Sanguinetti, Massive MIMO networks: spectral, energy, and hardware efficiency. Found. Trends Signal Process. **11**(3–4), 154–655 (2017)
24. J.H. Sørensen, E. De Carvalho, Č Stefanovic, P. Popovski, Coded pilot random access for massive MIMO systems. IEEE Trans. Wireless Commun. **17**(12), 8035–8046 (2018)
25. S.K. Mohammed, R. Hadani, A. Chockalingam, R. Calderbank, OTFS-Predictability in the delay-doppler domain and its value to communication and radar sensing. IEEE BITS Inf. Theory Mag. **3**(2), 7–31 (2023)
26. S. Mohammed, R. Hadani, A. Chockalingam, O.T.F.S. Modulation, *Theory and Applications* (Wiley, UK, 2024)
27. Q. Deng, Y. Ge, Z. Ding, A unifying view of OTFS and its many variants. IEEE Commun. Surv. Tutor. **27**(6), 3561–3586 (2025)
28. Y. Hong, T. Thaj, E. Viterbo, *Delay-Doppler communications: principles and applications* (Academic Press, 2022)
29. A. Das, F. Jesbin, A. Chockalingam, Closed-form expressions for I/O relation in Zak-OTFS with different delay-doppler filters. IEEE Trans. Veh. Technol. **74**(9), 1–16 (2025)
30. A. Das, F. Jesbin, A. Chockalingam, A Gaussian-sinc pulse shaping filter for Zak-OTFS, (2025), Preprint at https://arxiv.org/abs/2502.03904
31. J. Jayachandran, R. K. Jaiswal, S. K. Mohammed, R. Hadani, A. Chockalingam, R. Calderbank, Zak-OTFS: Pulse shaping and the tradeoff between time/bandwidth expansion and predictability, (2024), Preprint at arXiv:2405.02718
32. V. Khammammetti, S.K. Mohammed, OTFS-based multiple-access in high doppler and delay spread wireless channels. IEEE Wireless Commun. Lett. **8**(2), 528–531 (2019)
33. R. M. Augustine, A. Chockalingam, Interleaved time-frequency multiple access using OTFS modulation, In: IEEE 90th Vehicular Technology Conference (VTC2019-Fall), pp. 1–5 (2019)
34. V. Khammammetti, S.K. Mohammed, Spectral efficiency of OTFS based orthogonal multiple access with rectangular pulses. IEEE Trans. Veh. Technol. **71**(12), 12 989-13 006 (2022)
35. M. Ubadah, S. K. Mohammed, R. Hadani, S. Kons, A. Chockalingam, R. Calderbank, Zak-OTFS for integration of sensing and communication, (2024), Preprint at https://arxiv.org/abs/2404.04182
36. R. Hadani, S. Rakib, M. Tsatsanis, A. Monk, A. J. Goldsmith, A. F. Molisch, R. Calderbank, Orthogonal time frequency space modulation, In: 2017 IEEE wireless communications and networking conference (WCNC), pp. 1–6 (2017)
37. J.G. Proakis, M. Salehi, *Digital communications* (McGraw-hill, 2008)
38. Y.S. Cho, J. Kim, W.Y. Yang, C.G. Kang, *MIMO-OFDM wireless communications with MATLAB* (Wiley, 2010)
39. I.-R. Recommendation, Guidelines for evaluation of radio transmission technologies for IMT-2000, Rec. ITU-R M. 1225, (1997)
40. J. Haghighat, T.M. Duman, An energy-efficient feedback-aided irregular repetition slotted ALOHA scheme and its asymptotic performance analysis. IEEE Trans. Wireless Commun. **22**(12), 9808–9820 (2023)
41. A. Graell i Amat, G. Liva, Finite-length analysis of irregular repetition slotted ALOHA in the waterfall region. IEEE Commun. Lett. **22**(5), 886–889 (2018)

Publisher's Note