Density Evolution Analysis of Structured IRSA With Multi-Packet Reception and Orthogonal Pilots

Alessandro Mirri, *Graduate Student Member, IEEE*, and Enrico Paolini, *Senior Member, IEEE*

***Abstract*— We present an analytical framework based on density evolution (DE) for irregular repetition slotted ALOHA (IRSA) operating over orthogonal time-frequency resources. A structured IRSA scheme is introduced, in which each active user transmits multiple replicas across a selected subset of slot-subcarriers within a frame. The proposed DE formulation explicitly accounts for multi-packet reception (MPR) and orthogonal pilot selection, extending classical DE analysis to capture the combined effects of structured repetition and MPR-enabled decoding. Monte Carlo simulations, carried out under both idealized and realistic physical-layer conditions, validate the analysis and confirm its accuracy in characterizing and optimizing next-generation random access protocols.**

***Index Terms*— Coded random access, density evolution, massive machine-type communications, multi-packet reception, structured IRSA.**

I. Introduction

GRANT-FREE random access is a key enabler of massive machine-type communications (mMTC), where numerous devices sporadically transmit short packets to a central base station (BS) without prior scheduling [1], [2]. Traditional grant-free protocols like slotted ALOHA [3] suffer from packet collisions that severely limit performance. To address this issue, coded random access (CRA) schemes have been proposed [4], which exploit packet repetitions and successive interference cancellation (SIC) to recover collided packets. Notable examples include contention resolution diversity slotted ALOHA (CRDSA) [5], irregular repetition slotted ALOHA (IRSA) [6], and the coded slotted ALOHA (CSA) framework [7].

To achieve high scalability under reliability and latency constraints, CRA schemes must be carefully designed. In IRSA, this design centers on optimizing the probability distribution of users' packet repetition rates, a problem naturally linked to codes on sparse graphs and typically analyzed through density evolution (DE) [8]. Most existing designs rely on simplified channel models, such as the collision channel [6], [7] and its multi-packet reception (MPR) extension [9]. More recent studies have incorporated more realistic physical (PHY)-layer features, e.g., Rayleigh block-fading, massive multiple-input multiple-output (MIMO), and practical detection techniques such as maximal ratio combining [10]. However, these works typically fix the medium access control (MAC)-layer repetition strategy, optimizing only the degree distribution under an unstructured, single-level repetition model where each user's replicas are spread across time-domain slots within a frame.

In this work, we extend the DE analysis to structured IRSA (S-IRSA) over a time-frequency resource grid, enabling multi-level packet repetition. Here, "structured" is used to emphasize that each active user transmits replicas in subset of resources, selecting a set of slots and then multiple distinct subcarriers within each of the selected slots. The proposed DE formulation explicitly incorporates both MPR at the resource level and orthogonal pilot selection, providing a rigorous analytical framework for structured random access. We validate the analysis through simulations under both a simplified mathematical model and a realistic PHY-layer setup including Rayleigh block-fading, channel estimation, multi-user detection (MUD), and practical SIC. The main contributions are summarized as follows: (i) Development of a DE framework for S-IRSA that captures joint time-frequency diversity under MPR-enabled decoding with pilot orthogonality; (ii) Validation through a simplified model showing excellent agreement with theory; (iii) Symbol-level simulations confirming the DE framework's accuracy under realistic PHY-layer conditions.

Notation: Throughout the letter, bold uppercase and lowercase letters denote matrices and vectors, respectively; $\boldsymbol{v}[i]$ and $\boldsymbol{M}[i,j]$ indicate the i-th entry of $\boldsymbol{v}$ and the (i,j) element of $\boldsymbol{M}$. The conjugate transpose is indicated by $(\cdot)^H$, and $\|\cdot\|$ denotes the Euclidean norm. For a set $\mathcal{S}$, $|\mathcal{S}|$ is its cardinality, and for a positive integer I, we define $[I] = \{1,\ldots,I\}$. We denote the probability that a random variable A takes the value a, i.e., $\mathbb{P}\{\mathsf{A}=a\}$, as $P(a)$. Similarly, we use $P(a|b,c)$ to indicate the conditional probability $\mathbb{P}\{\mathsf{A}=a|\mathsf{B}=b,\mathsf{C}=c\}$, and $\mathbb{P}\{\mathcal{E}\}$ to denote the probability that an event $\mathcal{E}$ occurs.

II. Preliminaries and Background

A. Density Evolution Over the Collision Channel

We review the DE equations of IRSA over the collision channel [6]. The system can be modeled as a bipartite graph in which K_{a} user nodes, also called variable nodes (VNs), are connected to N_{s} slot nodes, also referred to as check nodes (CNs). A VN representing an active user that transmits r packet replicas is connected by r edges to the corresponding r CNs associated with the selected slots. The repetition degree r is a random variable characterized by the probability generating function (PGF) $\Lambda(x) = \sum_r \Lambda_r x^r$, and it is independently drawn by each active user. The probability that an edge is connected to a degree-r VN is then $\lambda_r = (\Lambda_r r)/\Lambda'(1)$, where $\Lambda'(1) = \sum_r r\Lambda_r$ is the average user repetition degree.

On the slot side, a CN has c edges corresponding to the number of users that chose the slot to transmit a replica. Let ρ_c denote the probability that an edge is connected to a degree-c CN, which can be expressed as $\rho_c = (\Psi_c c)/(\sum_h \Psi_h h)$,

Received 28 January 2026; accepted 5 March 2026. Date of publication 10 March 2026; date of current version 19 March 2026. This work was 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"). The associate editor coordinating the review of this letter and approving it for publication was P. Giard. *(Corresponding author: Enrico Paolini.)*

The authors are with CNIT/WiLab, DEI, University of Bologna, 47522 Cesena, Italy (e-mail: alessandro.mirri7@unibo.it; e.paolini@unibo.it).

Digital Object Identifier 10.1109/LCOMM.2026.3672440

where Ψ_c is the probability that c users transmit in the considered slot. The VN degree distribution $\Lambda(x)$ (i.e., the IRSA distribution) can be optimized by the system designer. In contrast, the CN degree distribution $\Psi(x) = \sum_c \Psi_c x^c$ is determined by the system load G and the average user repetition rate $\Lambda'(1)$. In particular, for a large user population K, it is common to assume that the number of transmissions per slot follows a Poisson distribution

$$\Psi_c = \frac{\big(G\Lambda'(1)\big)^c}{c!} \exp\big(-G\Lambda'(1)\big), \tag{1}$$

which in turn leads to

$$\rho_c = \frac{\big(G\Lambda'(1)\big)^{c-1}}{(c-1)!} \exp\big(-G\Lambda'(1)\big). \tag{2}$$

When adopting a collision channel model, its main assumptions can be summarized as follows: (i) If more than one user transmits in a slot, the receiver cannot successfully decode any of the packets. (ii) If exactly one user transmits in a slot, the packet is decoded successfully with zero error probability and can be perfectly removed from the slot. (iii) Whenever a packet is successfully decoded, the interference caused by its replicas in other slots can be perfectly canceled (perfect SIC). Under the collision channel assumptions, let $q_\ell^{(r)}$ denote the probability that an edge connected to a degree-r VN remains unknown at the end of decoding iteration ℓ. Similarly, let $p_\ell^{(c)}$ be the probability that an edge connected to a degree-c CN is still unknown at the end of iteration ℓ. Using the edge-oriented probabilities λ_r and ρ_c, the average probabilities q_ℓ and p_ℓ can be expressed as

$$q_\ell = \sum_r \lambda_r q_\ell^{(r)} \quad \text{and} \quad p_\ell = \sum_c \rho_c p_\ell^{(c)}. \tag{3}$$

For a degree-r VN, an edge is revealed as soon as at least one of the other edges connected to the same VN has been recovered (i.e., it is sufficient that a single replica of the packet within the frame is successfully decoded). Consequently, the probability that a packet remains unrecovered is

$$q_\ell^{(r)} = p_{\ell-1}^{r-1}. \tag{4}$$

Similarly, for a degree-c CN, an edge is revealed only when all the other edges connected to that CN have already been revealed, in line with the collision channel model (i.e., a packet in a slot can be decoded if and only if it is the sole remaining transmission in that slot). Accordingly, the probability that a packet in a slot has not yet been canceled by SIC is

$$p_\ell^{(c)} = 1 - (1 - q_\ell)^{c-1}. \tag{5}$$

Combining equations from (2) to (5) yields the DE recursion for the collision channel

$$p_\ell = 1 - \exp\left(-G \sum_r r\Lambda_r p_{\ell-1}^{r-1}\right), \tag{6}$$

with initial condition $p_0 = 1$ (all edges are initially unknown). The asymptotic packet loss rate at iteration ℓ is

$$Q_\ell = \sum_r \Lambda_r p_\ell^r, \tag{7}$$

and the asymptotic load threshold of an IRSA distribution $\Lambda(x)$ is defined as

$$G^\star = \sup\big\{G > 0 \,:\, Q_\ell \to 0 \text{ as } \ell \to \infty\big\}. \tag{8}$$

Note that $Q_\ell \to 0$ if and only if $p_\ell \to 0$. Here, "asymptotic" refers to the assumption of statistical independence of messages along the graph edges, which holds in the limit $K \to \infty$, $N_\mathrm{s} \to \infty$, with a constant ratio K/N_s.

B. S-IRSA: Mathematical System Model

We consider a grant-free, frame-synchronous S-IRSA protocol in which K_a active users (out the overall population K, typically $K_\mathrm{a} \ll K$) contend for access over a frame composed of N_s time slots and F frequency subcarriers. The set of all time-frequency resources is defined as $\mathcal{R} = \{(i,j) : i \in [N_\mathrm{s}],\ j \in [F]\}$, so that each resource $n = (i,j) \in \mathcal{R}$ corresponds to a slot-subcarrier pair. Each user $k \in [K_\mathrm{a}]$ transmits a single information packet using a structured repetition strategy.

First, user k draws a time-domain repetition degree r_k from a given distribution $\Lambda(x)$, as in conventional IRSA, and uniformly selects a set $\mathcal{S}_k \subseteq [N_\mathrm{s}]$ of r_k distinct slots out of the N_s available ($|\mathcal{S}_k| = r_k$). Then, for each selected slot $i \in \mathcal{S}_k$, the user independently chooses a set $\mathcal{F}_{k,i} \subseteq [F]$ of f subcarriers from the F available ones ($|\mathcal{F}_{k,i}| = f$), where the number of frequency-domain replicas f is fixed and identical for all users.[1] The overall set of resources employed by user k is therefore $\mathcal{V}_k = \big\{(i,j) \in \mathcal{R} : i \in \mathcal{S}_k,\ j \in \mathcal{F}_{k,i}\big\}$, so that each user transmits exactly $|\mathcal{V}_k| = r_k \cdot f$ replicas, one per selected resource. Each replica $v \in \mathcal{V}_k$ carries a pilot sequence chosen from a pool of N_P orthogonal ones, according to a mapping $p_k : \mathcal{V}_k \to [N_\mathrm{P}]$, where $p_k(n)$ denotes the index of the pilot used by user k in resource $n = (i,j)$.

For each resource $n \in \mathcal{R}$, let $\mathcal{A}_n = \{\, k : n \in \mathcal{V}_k \,\}$ denote the set of users transmitting on that resource, and, for each $q \in [N_\mathrm{P}]$, $\mathcal{P}_{n,q} = \{\, k \in \mathcal{A}_n : p_k(n) = q \,\}$ denote the subset of users employing pilot q. Each resource n supports a MPR capability $T \geq 1$, representing the maximum number of simultaneously decodable packets within that resource. A resource n enables successful decoding of all packet replicas transmitted within it if

$$|\mathcal{A}_n| \leq T \quad \text{and} \quad |\mathcal{P}_{n,q}| \leq 1 \ \forall q \in [N_\mathrm{P}], \tag{9}$$

i.e., the number of active users does not exceed the MPR capability, and pilot collisions are avoided. If either $|\mathcal{A}_n| > T$ or $|\mathcal{P}_{n,q}| \geq 2$ for some n, q, no packet can be recovered in that resource until SIC is performed. Decoding proceeds iteratively across the frame. At each iteration, the receiver scans all resources, decodes according to (9), and collects the set of newly decoded users $\mathcal{D}$. For each decoded user $\ell \in \mathcal{D}$, its contribution is perfectly removed via SIC from every resource $n \in \mathcal{V}_\ell$ where it was active

$$\mathcal{A}_n \leftarrow \mathcal{A}_n \setminus \{\ell\}, \quad \mathcal{P}_{n,q} \leftarrow \mathcal{P}_{n,q} \setminus \{\ell\} \ \forall q \in [N_\mathrm{P}]. \tag{10}$$

The SIC process continues iteratively until no further users can be decoded, i.e., $\mathcal{D} = \emptyset$.

C. S-IRSA: Realistic PHY-Layer Model

Each user transmits a message $\mathsf{W} \in [2^\kappa]$, encoded at rate R_c, modulated using quadrature phase-shift keying (QPSK), and mapped into N_D payload symbols. Each replica includes N_P pilot symbols randomly selected from a pool of N_P orthogonal pilots, followed by the N_D payload symbols. The access and repetition strategy follows the model described in Section II-B. The wireless channel is modeled as Rayleigh

[1]This choice allows mathematical tractability of the DE analysis.

block fading, constant within each resource (slot-subcarrier pair) and independent across both users and resources. At the receiver, the M-antenna MIMO BS observes each resource $n \in \mathcal{R}$ as $[\boldsymbol{P}_n, \boldsymbol{Y}_n] \in \mathbb{C}^{M \times (N_\mathrm{P}+N_\mathrm{D})}$, with

$$\boldsymbol{P}_n = \sum_{k \in \mathcal{A}_n} \boldsymbol{h}_{n,k} \boldsymbol{s}_{n,k} + \boldsymbol{Z}_n, \boldsymbol{Y}_n = \sum_{k \in \mathcal{A}_n} \boldsymbol{h}_{n,k} \boldsymbol{x}_k + \boldsymbol{W}_n, \quad (11)$$

where $\boldsymbol{h}_{n,k} \sim \mathcal{CN}(\boldsymbol{0}, \boldsymbol{I}_M)$ denotes the vector of channel coefficients of user k on resource n, $\boldsymbol{s}_{n,k}$ is the pilot sequence, $\boldsymbol{x}_k$ is the user's payload (identical across replicas), and $\boldsymbol{Z}_n$ and $\boldsymbol{W}_n$ are additive white Gaussian noise (AWGN) matrices. The BS scans the frame slot by slot, processing all subcarriers.

For each resource, channel estimation is performed for every orthogonal pilot $q \in [N_\mathrm{P}]$ according to

$$\widehat{\boldsymbol{h}}_{n,q} = \frac{\boldsymbol{P}_n \boldsymbol{s}_q^H}{\|\boldsymbol{s}_q\|^2} \,. \quad (12)$$

A pilot q is declared active if $|\widehat{\boldsymbol{h}}_{n,q}|^2 > \eta$.[2] We assume each detected pilot corresponds to a single active user; any pilot collision results in decoding failure [10]. When the number of detected users satisfies $|\widehat{\mathcal{A}}_n| \leq T$, the receiver performs an MUD algorithm which consists of estimating, for each payload symbol $t \in [N_\mathrm{D}]$, all users' contributions according to

$$\{\widehat{\boldsymbol{x}}_k[t]\}_{k \in \widehat{\mathcal{A}}_n} = \underset{\{x_k \in \mathcal{Q}\}}{\operatorname{argmin}} \sum_{m=1}^{M} \left| \boldsymbol{Y}_n[m,t] - \sum_{k \in \widehat{\mathcal{A}}_n} \widehat{\boldsymbol{h}}_{n,k}[m] \, x_k \right|^2 . \quad (13)$$

where $\mathcal{Q}$ represents the QPSK symbols constellation. The MPR capability T is as a deterministic, resource-independent parameter that limits the affordable MUD complexity and should be regarded as a receiver design choice rather than an intrinsic PHY-layer parameter. Notably, the realistic PHY-layer MUD processing reflects the decoding condition in (9).

When a user ℓ is successfully decoded, its contribution is removed from all associated resources $n \in \mathcal{V}_\ell$ via SIC

$$\boldsymbol{P}_n \leftarrow \boldsymbol{P}_n - \widehat{\boldsymbol{h}}_{n,\ell} \boldsymbol{s}_{n,\ell}, \boldsymbol{Y}_n \leftarrow \boldsymbol{Y}_n - \widehat{\boldsymbol{h}}_{n,\ell} \boldsymbol{x}_\ell. \quad (14)$$

SIC is first applied at the slot level ("intra-slot" SIC) exploiting frequency-domain replicas and then at the frame level ("inter-slot" SIC) once no further users can be decoded within the slot. For resources where user ℓ was not directly decoded, the channel is estimated using the known payload [11]

$$\widehat{\boldsymbol{h}}_{n,\ell} = \frac{\boldsymbol{Y}_n \boldsymbol{x}_\ell^H}{\|\boldsymbol{x}_\ell\|^2}. \quad (15)$$

Reliable "payload-aided" subtraction is enabled by the favorable propagation effect [12], which allows for spatial separation among users' payloads as the number of receive antennas at the BS increases. After each SIC iteration, pilot detection and MUD are re-run on affected resources until no additional users are recovered.

III. Density Evolution of S-IRSA

We extend the DE framework introduced in Section II-A for the collision channel to the S-IRSA scheme, incorporating structured replica transmission over time-frequency resources, an MPR-enabled channel, and orthogonal pilot sequences as described in Section II-B. As it will be shown in Section IV, the DE analysis also accurately predicts the performance of the realistic PHY-layer model of Section II-C, assuming N_P orthogonal pilots and an MPR capability T. The system is modeled as a bipartite graph, where the K_a active users correspond to VNs, the N_s slots correspond to CNs, and edges represent users' transmissions across slots.

Consider a slot of degree c, i.e., c users transmit replicas over its associated resources. For a given edge (user replica), let I denote the number of unresolved interferers in that resource at iteration ℓ, with $0 \leq \mathsf{I} \leq c-1$. Let ϵ_ℓ denote the event that a user replica transmitted in a resource of the slot is not correctly decoded at iteration ℓ. Then, the probability that an edge remains unresolved at the end of iteration ℓ, $p_\ell^{(c)}$, is

$$p_\ell^{(c)} = \left(P(\epsilon_\ell|c)\right)^f, \quad (16)$$

where f is the number of frequency-domain replicas per user per slot. The probability of failure within a single resource, conditioned on c arrivals in the slot, is

$$P(\epsilon_\ell|c) = \sum_{i=0}^{c-1} P(\epsilon_\ell|c,i)\, P(i|c), \quad (17)$$

where i denotes the number of unresolved interferers in the resource, $P(i|c)$ their occurrence probability, and $P(\epsilon_\ell|c,i)$ the corresponding conditional error probability.

Due to the MPR capability T, we have

$$P(\epsilon_\ell|c) = \begin{cases} \sum_{i=0}^{c-1} P(\epsilon_\ell|c,i) P(i|c), & c \leq T, \\ \sum_{i=0}^{T} P(\epsilon_\ell|c,i) P(i|c) + \sum_{i=T+1}^{c-1} P(i|c), & c > T. \end{cases} \quad (18)$$

Whenever the number of unresolved users exceeds T, decoding fails, i.e., $P(\epsilon_\ell|c,i) = 1$. Given $i \leq T$ unresolved interferers, a decoding failure occurs if at least one of them selects the same pilot as the reference user. Assuming N_P orthogonal pilot sequences, the corresponding probability is

$$P(\epsilon_\ell|c,i) = 1 - \left(1 - \frac{1}{N_\mathrm{P}}\right)^i = 1 - \left(\frac{N_\mathrm{P}-1}{N_\mathrm{P}}\right)^i. \quad (19)$$

Moreover, the probability of observing i unresolved interferers in a given resource can be derived by conditioning on the total number k of unresolved interferers within the corresponding slot, namely,

$$P(i|c) = \sum_{k=i}^{c-1} P(i|k,c)\, P(k|c). \quad (20)$$

Here, $P(k|c)$ denotes the probability that k interferers remain unresolved within the slot, while $P(i|k,c)$ represents the probability that exactly i out of these k interferers occupy the considered resource. For a slot of degree c, the number of unresolved interferers k follows a binomial distribution

$$P(k|c) = \binom{c-1}{k} (q_\ell)^k (1-q_\ell)^{c-1-k}, \quad (21)$$

where q_ℓ denotes the probability that a generic edge remains unknown at iteration ℓ. Given k interferers within the slot, the number occupying a particular resource follows a binomial

[2] The threshold η balances misdetection and false alarm events and can be set adopting the maximum-likelihood criterion.

TABLE I

ASYMPTOTIC LOAD THRESHOLD $G^\star$ [users/slot] OF S-IRSA AS A FUNCTION OF THE MPR CAPABILITY T AND THE NUMBER OF ORTHOGONAL PILOTS N_P. THE HIGHEST THRESHOLD FOR EACH T IS HIGHLIGHTED IN BOLD

$\Lambda(x)$	$T=1$	$T=2$	$T=5$	$T=10$	$T\to\infty$
$N_\mathrm{P}=4$					
x^2	2.298	4.087	**7.513**	9.643	9.689
x^3	2.272	3.545	6.123	8.626	9.310
$0.5x^2+0.5x^3$	2.424	3.929	6.926	9.603	10.007
$0.86x^2+0.14x^8$	**2.710**	**4.194**	7.099	**10.088**	**11.109**
$N_\mathrm{P}=64$					
x^2	2.298	**4.460**	**9.554**	**16.804**	157.017
x^3	2.272	3.771	7.311	12.315	149.669
$0.5x^2+0.5x^3$	2.424	4.210	8.410	14.354	161.204
$0.86x^2+0.14x^8$	**2.710**	4.438	8.320	13.744	**178.574**

distribution with success probability f/F, where f subcarriers are selected out of a total of F. Hence, we have

$$P(i|k)=\binom{k}{i}\left(\frac{f}{F}\right)^{i}\left(1-\frac{f}{F}\right)^{k-i}. \tag{22}$$

By successively substituting (22) and (21) into (20), then (20) and (19) into (18), and finally (18) into (16), we derive the recursion p_ℓ for the two MPR regimes. For example, for a configuration with $F=4$ subcarriers and $f=2$ frequency-domain replicas per user,[3] the recursion becomes

$$p_\ell^{(c)} = \begin{cases} 1+\left(1-\dfrac{q_\ell}{2N_\mathrm{P}}\right)^{2(c-1)} - 2\left(1-\dfrac{q_\ell}{2N_\mathrm{P}}\right)^{c-1} & c\le T \\ \left[1-\left(1-\dfrac{q_\ell}{2}\right)^{c-1}\displaystyle\sum_{k=0}^{T-1}\binom{c-1}{k}\left(\dfrac{q_\ell(N_\mathrm{P}-1)}{N_\mathrm{P}(2-q_\ell)}\right)^{k}\right]^{2} & c>T. \end{cases} \tag{23}$$

By combining (2), (3), and (4), previously derived for the collision channel, with (23), we obtain the complete density evolution recursion for the proposed scenario. Finally, applying (7) and (8) yields the asymptotic load threshold.

In the limiting cases $T=1$ and $T\to\infty$, the recursion admits simplified closed-form expressions. For single-packet reception ($T=1$), the BS can decode at most one user per resource due to the limited MUD capability. In this regime, corresponding to $c>T$ for all $c>1$, the recursion reduces to

$$p_\ell = 1+\exp\left(-\frac{G}{R}\left(q_\ell-\frac{q_\ell^2}{4}\right)\right)-2\exp\left(-\frac{G}{R}\frac{q_\ell}{2}\right), \tag{24}$$

where $R=1/\Lambda'(1)$. For unbounded MPR ($T\to\infty$), any number of users can be potentially decoded per resource, corresponding to an ideal scenario in which MUD is no longer a computational limitation. In this case, where the MPR constraint is effectively removed (i.e., $c\le T$ always), the failure probability depends solely on pilot collisions across frequency-domain replicas, and the recursion simplifies to

$$p_\ell = 1+\exp\left(-\frac{G}{R}\frac{q_\ell}{N_\mathrm{P}}\left(1-\frac{q_\ell}{4N_\mathrm{P}}\right)\right)-2\exp\left(-\frac{G}{R}\frac{q_\ell}{2N_\mathrm{P}}\right). \tag{25}$$

[3]The analysis can also be extended to higher values of frequency usage (i.e., $f\in\{3,4\}$); however, simulations show that no performance improvement with respect to $f=2$ is achieved in these cases.

IV. NUMERICAL RESULTS

A. Simulation Setup

For the numerical simulations of S-IRSA, we consider a time–frequency frame with $N_\mathrm{s}=50$ slots and $F=4$ orthogonal subcarriers. Each user transmits $f=2$ frequency-domain replicas, while the number of time-domain replicas r is drawn according to $\Lambda(x)$. In the realistic PHY-layer simulations, each user employs a $(n=255,\kappa=99,t=23)$ Bose-Chaudhuri-Hocquenghem (BCH) code. Part of the k information bits is reserved for cyclic redundancy check (CRC)-based validation of decoded packets, with one null-padding bit appended. After QPSK modulation, each user transmits $N_\mathrm{D}=128$ payload symbols per resource. A single pilot sequence, randomly selected from a pool of $N_\mathrm{P}=64$ orthogonal Hadamard sequences with unit-normalized energy, is appended to each payload. The BS is equipped with $M=8$ antennas. In the adopted simulation setup, the realistic PHY-layer-based S-IRSA scheme operates in a high-SNR, interference-limited regime, which is the relevant condition for validating the proposed DE. For the DE analysis, the degree distribution $\Lambda(x)$ maximizing the asymptotic load is optimized using the *differential evolution* algorithm [13] with population size $P=20$, maximum degree $D=8$, crossover probability $\mathrm{CR}=0.5$, differential weight $W=1.2$, and maximum generations $J_\mathrm{max}=50$. The optimized distributions are obtained for the limiting cases where closed-form expressions of the DE recursion are available.

B. Performance Evaluation

1) Asymptotic Load Threshold Analysis: Table I reports the asymptotic load threshold $G^\star$ [users/slot] of S-IRSA with structured time-frequency repetitions, as a function of the MPR capability T and the number of orthogonal pilots N_P. For a small pilot pool ($N_\mathrm{P}=4$), the optimized distribution $\Lambda(x)=0.86x^2+0.14x^8$ maximizes $G^\star$ for low MPR capability (i.e., $T\in\{1,2\}$). For moderate values of $T\simeq N_\mathrm{P}$ (e.g., $T=5$), a concentrated distribution $\Lambda(x)=x^2$ is preferable. For higher T, the distribution optimized via differential evolution again achieves the highest $G^\star$. For a larger pilot set size ($N_\mathrm{P}=64$), more users can be resolved per resource, leading to higher asymptotic loads for all $T>1$. In this setting, for moderate MPR capability (i.e., $1<T<N_\mathrm{P}$), performance is limited by MPR. As a result, the concentrated distribution with fewer replicas, $\Lambda(x)=x^2$, is the best. Conversely, the optimized distribution reaches the best performance for very large T, where pilot availability becomes the dominant limiting factor and a higher number of replicas allows to mitigate pilot collisions. In the MPR-unrestricted regime, system performance is solely limited by pilot collisions, and increasing N_P directly improves the asymptotic load threshold.

2) Performance Comparison Between IRSA and S-IRSA: Table II compares baseline IRSA and S-IRSA under fair conditions by enforcing the same total resources ($N_\mathrm{s}\cdot F$) and the same effective number of packet replicas per user. To this end, the asymptotic load is normalized by the number of subcarriers F, and identical power budgets are ensured by

TABLE II

PERFORMANCE COMPARISON OF IRSA AND S-IRSA UNDER FAIR RESOURCE ALLOCATION, FOR $T = 1$. THE ASYMPTOTIC LOAD THRESHOLD $G^\star/F$[users/resource] IS REPORTED

Scheme - $\Lambda(x)$	$\Lambda'(1) \cdot f$	$G^\star/F$
IRSA - x^4	4.00	0.772
S-IRSA - x^2	4.00	0.575
IRSA - x^6	6.00	0.637
S-IRSA - x^3	6.00	0.568
IRSA - $0.5x^2 + 0.28x^3 + 0.22x^8$	3.60	0.939
S-IRSA - $0.86x^2 + 0.14x^8$	5.68	0.678

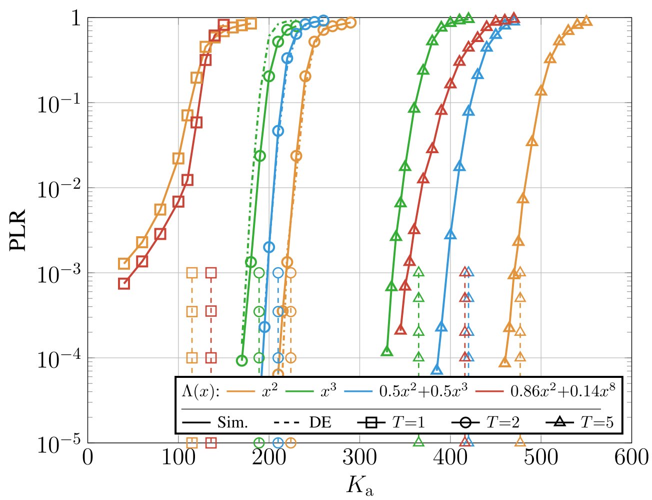

Fig. 1. Packet loss rate (PLR) versus number of active users K_{a} for $N_{\mathrm{P}} = 64$ and different MPR capabilities T. Solid lines: Monte Carlo simulations of S-IRSA. Dashed lines: analytical predictions via DE. Dash-dotted lines: realistic PHY-layer performance including coding, modulation, channel estimation, MUD, and practical SIC.

matching the total number of transmitted replicas per user $r \cdot f$. Specifically, we adopt concentrated IRSA degree distributions x^4 and x^6 with $f = 1$, which match the total replica counts of the S-IRSA schemes using degree distributions x^2 and x^3 with $f = 2$. For S-IRSA, $F = 4$, while for IRSA, $F = f = 1$. IRSA achieves higher asymptotic thresholds due to the flexibility of distributing replicas across the entire frame, but this requires end-of-frame decoding, increasing latency. S-IRSA, in contrast, exploits frequency-domain replicas within each slot, enabling intra-slot SIC and faster decoding. For example, with $\Lambda(x) = x^2$ for S-IRSA and $\Lambda(x) = x^4$ for IRSA, the normalized load values are $G^\star/F = 0.575$ and 0.772, respectively. Despite the slightly lower load, S-IRSA can recover multiple packets per slot and perform slot-level SIC.

3) Validation of Density Evolution Predictions: Fig. 1 compares the analytical predictions obtained via DE with Monte Carlo simulation results for S-IRSA under different degree distributions $\Lambda(x)$. The curves show the packet loss rate versus the number of active users K_{a}, computed as $K_{\mathrm{a}} = N_{\mathrm{s}} \cdot G$.

An excellent match is observed across the entire range of active users and for different MPR capabilities T, confirming the accuracy of the DE framework. For completeness, the figure also includes (dash-dotted lines, $T = 2$) results from realistic PHY-layer simulations. These curves exhibit only a slight degradation compared to the idealized setting, while maintaining the same trend, thus validating the reliability of the analytical model under practical conditions.

V. CONCLUSION

In this letter, we developed a DE framework for S-IRSA that captures joint time-frequency diversity under MPR-enabled decoding with orthogonal pilots. The framework was first validated against a simplified model, showing excellent agreement with theoretical predictions, and further confirmed through symbol-level PHY-layer simulations. Further research could explore joint optimization of time- and frequency-domain repetitions, at the cost of increased analytical complexity.

REFERENCES

[1] C. Bockelmann et al., "Massive machine-type communications in 5G: Physical and MAC-layer solutions," *IEEE Commun. Mag.*, vol. 54, no. 9, pp. 59–65, Sep. 2016.

[2] C. Bockelmann et al., "Towards massive connectivity support for scalable mMTC communications in 5G networks," *IEEE Access*, vol. 6, pp. 28969–28992, 2018.

[3] L. G. Roberts, "ALOHA packet system with and without slots and capture," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 5, no. 2, pp. 28–42, Apr. 1975.

[4] M. Berioli, G. Cocco, G. Liva, and A. Munari, "Modern random access protocols," *Found. Trends Netw.*, vol. 10, no. 4, pp. 317–446, Nov. 2016.

[5] E. Casini, R. De Gaudenzi, and O. R. Herrero, "Contention resolution diversity slotted ALOHA (CRDSA): An enhanced random access schemefor satellite access packet networks," *IEEE Trans. Wireless Commun.*, vol. 6, no. 4, pp. 1408–1419, Apr. 2007.

[6] G. Liva, "Graph-based analysis and optimization of contention resolution diversity slotted ALOHA," *IEEE Trans. Commun.*, vol. 59, no. 2, pp. 477–487, Feb. 2011.

[7] E. Paolini, G. Liva, and M. Chiani, "Coded slotted ALOHA: A graph-based method for uncoordinated multiple access," *IEEE Trans. Inf. Theory*, vol. 61, no. 12, pp. 6815–6832, Dec. 2015.

[8] T. J. Richardson, M. A. Shokrollahi, and R. L. Urbanke, "Design of capacity-approaching irregular low-density parity-check codes," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 619–637, Feb. 2001.

[9] C. Stefanovic, E. Paolini, and G. Liva, "Asymptotic performance of coded slotted ALOHA with multipacket reception," *IEEE Commun. Lett.*, vol. 22, no. 1, pp. 105–108, Jan. 2018.

[10] L. Valentini, M. Chiani, and E. Paolini, "A joint PHY and MAC layer design for coded random access with massive MIMO," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Rio de Janeiro, Brazil, Dec. 2022, pp. 2505–2510.

[11] L. Valentini, M. Chiani, and E. Paolini, "Interference cancellation algorithms for grant-free multiple access with massive MIMO," *IEEE Trans. Commun.*, vol. 71, no. 8, pp. 4665–4677, Aug. 2023.

[12] E. Björnson, J. Hoydis, and L. Sanguinetti, "Massive MIMO networks: Spectral, energy, and hardware efficiency," *Found. Trends Signal Process.*, vol. 11, nos. 3–4, pp. 154–655, 2017.

[13] R. Storn and K. Price, "Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces," *J. Global Optim.*, vol. 11, no. 4, pp. 341–359, Dec. 1997.

Open Access funding provided by 'Alma Mater Studiorum - Università di Bologna' within the CRUI CARE Agreement