

RASP: Reliability-Aware SINR Prediction for Realistic Industrial Subnetworks

Pramesh Gautam*, Christian Arendt[‡], Steffen Fricke[‡], Carsten Bockelmann*,
Armin Dekorsy*, and Christian Wietfeld[‡]

*Department of Communications Engineering, University of Bremen, Germany,
Email: {gautam, bockelmann, dekorsy}@ant.uni-bremen.de

[‡]Communication Networks Institute, TU Dortmund University, Germany,
Email: {christian.arendt, steffen.fricke, christian.wietfeld}@tu-dortmund.de

Abstract—Industrial networks often demand hyper-reliable, low-latency communication (HRLLC) to support closed-loop control and automation. However, performance is severely undermined by dynamic fluctuations caused by unknown interferers, making it difficult to uphold reliability requirements. Channel statistics and mobility of interferers lead to significant variations in perceived signal-to-interference-plus-noise ratio (SINR). To cope with such interference dynamics, predictive resource allocation is required, whereby resource selection must be determined before transmission, making SINR prediction essential. This paper introduces Reliability-Aware SINR Prediction (RASP), a probabilistic framework that predicts the SINR distribution under explicit block error rate (BLER) constraints. RASP employs a Mixture Density Network within a Conditional Value-at-Risk (CVaR)-based primal-dual optimization scheme to adapt streaming SINR samples while satisfying reliability constraints. Based on real-world industrial measurements with co-subband interferers, we show that RASP achieves the target BLER of 10^{-6} at the 95th percentile across three distinct scenarios, outperforming the baselines by a significant margin.

Index Terms—6G, SINR Prediction, Industrial networks, Resource Allocation, 3GPP.

I. INTRODUCTION

As wireless systems evolve beyond 5G, the vision of 6G extends not only to higher data rates and massive connectivity but also to unprecedented levels of reliability, low latency, and intelligence for emerging use cases such as the Industrial Internet of Things (IIoT), autonomous systems, and immersive applications [1]. Within this vision, the International Telecommunication Union (ITU) has proposed the evolution of ultra-reliable, low-latency communication (URLLC) to hyper-reliable, low-latency communication (HRLLC), which aims to meet even more stringent requirements and to enable scenarios such as industrial robots operating as sub-networks (SNs) with dependable closed-loop control communication [1]. These industrial SNs often collaborate to accomplish complex tasks, where multiple machines or robots exchange control information in real time. In such settings, achieving sub-millisecond latency is particularly challenging, as it requires

high bandwidth for closed-loop control communication, and this leads to significant co-subband interference and dynamic channel conditions that can severely degrade reliability [2]. To address this, we first measure such an industrial scenario in real time using a controlled test environment that captures the impact of co-subband interference and dynamic channel variations under realistic conditions. Using these measurements, we propose and evaluate a reliability-aware signal to interference plus noise ratio (SINR) prediction technique with target block error rate (BLER) in the HRLLC regime as a reliability constraint [3].

SINR prediction has been widely studied in the literature. Early work by Xu *et al.* employed channel quality indicator (CQI) prediction using Wiener filtering and spline extrapolation, demonstrating improved quality of service (QoS) in long-term evolution (LTE) systems by compensating for feedback delays. Kriging-based methods have also been proposed to account for correlated shadowing, capturing spatial and inter-cell dependencies for enhanced coverage and throughput estimation [4]. More recently, machine learning (ML)-based approaches have gained traction because wireless channels in dynamic environments often violate assumptions such as stationarity and fixed interference, rendering model-based methods inadequate to generalize across scenarios [5]. In contrast, ML can learn the underlying dynamics directly from data without relying on such assumptions. For instance, predictive models based on long short-term memory (LSTM), Bi-LSTM, and Random Forests have been applied to real-time 5G measurements in private campus networks, achieving reasonable accuracy in controlled environments [6]. Similarly, deep learning has been applied for signal to noise ratio (SNR) prediction using real-world datasets captured with software-defined radios, confirming the feasibility of data-driven prediction in realistic setups [7]. However, existing works focus primarily on point prediction of future SINR, rather than its probability distribution, which limits their ability to capture fast-fading effects that are critical for proactive interference management in the HRLLC regime. This limitation is further exacerbated in industrial subnetworks, where interference and channel statistics do not vary slowly or average out over time,

This work is supported by the German Federal Ministry of Research, Technology and Space (BMFTR) under the grants of 16KISK109 (6G-ANNA), 16KISK016 (Open6GHub) and 16KISK038 (6GEM).

but instead exhibit rapidly changing and heavy-tailed behavior driven by dynamic co-subband interference, thereby necessitating distributional, online, and reliability-aware prediction [3].

In this work, we propose a SINR probability density function (PDF) prediction technique for industrial SNs. We formulate the task as an online learning problem in which the predictive distribution adapts to streaming SINR samples under explicit BLER-based reliability constraints. To address the intractability of the reliability constraint, we employ a Conditional Value at Risk (CVaR)-based relaxation and embed an Mixture Density Network (MDN) within a reliability-aware Lagrangian framework. The solution is implemented via an online primal-dual algorithm and validated using real-time industrial measurement traces. To the best of our knowledge, this is the first work to formulate and solve SINR distribution prediction with explicit reliability guarantees for industrial SNs, which we refer to as the Reliability-Aware SINR Prediction (RASP) framework. The main contributions are summarized below:

- We formulate reliability-constrained SINR prediction for industrial SNs as an online learning problem with explicit BLER-based constraints.
- We propose a CVaR-based relaxation with an exponentially weighted estimator and embed an MDN into a reliability-aware Lagrangian framework.
- We develop an online primal-dual algorithm to jointly update the MDN parameters, auxiliary Value at Risk (VaR) variable, and dual variable.
- We validate the proposed approach on real-world industrial measurement traces, showing that it achieves the target BLER of 10^{-6} at the 95th percentile across three distinct scenarios.

The rest of the paper is structured as follows: the system model is presented in section II. The problem formulation and proposed RASP framework with learning algorithm are described in section III and section IV. In section V, the evaluation setup and performance results based on industrial measurements are discussed. Finally, the conclusion and directions for future work are given in section VI.

Notation: Bold lowercase and uppercase letters denote vectors and matrices, respectively, and $|\cdot|$ denotes the cardinality of a set.

II. SYSTEM MODEL AND SINR MEASUREMENT

SINR measurements constitute a fundamental radio resource management mechanism in 5G New Radio (NR) systems, providing a quantitative assessment of signal quality by evaluating the ratio between desired signal power and the combined interference and noise power. These measurements serve critical functions in mobility management procedures, including handover decisions, beam management, and link adaptation, and enable dynamic modulation and coding scheme selection to optimize spectral efficiency. In 5G NR, SINR measurements are performed using either synchronisation signal (SS) blocks to estimate SS-SINR, or channel state information-reference



Fig. 1: Transfer lab environment at TU Dortmund University.

Signal (CSI-RS) to estimate CSI-SINR [8], with the former utilizing Secondary Synchronization Signals and Physical Broadcast Channel Demodulation Reference Signals for cell-level quality assessment, and the latter employing dedicated CSI-RS resources for more granular beam-level evaluations. In this work, SS-SINR is used as periodically reported by a commercial off-the-shelf (COTS) modem. Mathematically, the measured SINR at t -th transmission time interval (TTI) can be written as

$$\Gamma(t) = \frac{S(t)}{\sum_{c \in \mathcal{C}} I_c(t) + w}, \quad (1)$$

where $S(t)$ denotes the received signal power between the access point (AP) and user equipment (UE) under continuous packet transmission, capturing path loss, small-scale fading, and shadowing. The interference term I_c originates from a defined set of mobile interferers \mathcal{C} . The $w \sim \mathcal{N}(0, \sigma^2)$ represents additive white Gaussian noise with zero mean and σ^2 variance.

The scenario in this work represents an industrial transfer environment at TU Dortmund University depicted in Fig. 1, where a fixed AP communicates with associated UE over line-of-sight (LOS) or non-line-of-sight (NLOS) paths to support closed-loop control communication. The AP-UE link is subject to interference from two mobile interferers, realized by robots moving along predefined trajectories. Each interferer is implemented with a Software Defined Radio (SDR) that generates signals in the same subband, operating at the same carrier frequency and bandwidth, thereby producing co-subband interference to the main AP-UE link. The measurement results published by the COTS modem are quantized into 128 discrete levels (0-127) representing SINR values from -23 dB to +40 dB with 0.5 dB resolution, subsequently reported to the AP (gNodeB) for informed radio resource management decisions.

III. PROBLEM FORMULATION

This section formulates the SINR PDF prediction problem, with particular emphasis on achieving the reliability demands

of HRLLC in industrial SNs. In industrial SN, conventional offline prediction is inadequate, as mobility, deployment conditions, and channel dynamics change over time, and the complete dataset is often unavailable a priori. Moreover, SINR measurements arrive sequentially and the distribution evolves over time. Therefore, an online learning approach is well-suited to incrementally refine the PDF as new data becomes available. Let $\Gamma(t)$ denote the measured SINR, and \mathcal{P}_t denote its predictive distribution with respect to its realized sample $\gamma(t)$, inferred from historical SINR samples. We predict \mathcal{P}_t using a PDF estimator $f_\Gamma(\gamma(t); \theta_t)$, with parameters θ_t learned and refined from streaming observations. In the HRLLC regime, SINR prediction alone is insufficient, since reliability is ultimately quantified by the BLER. In this work, reliability refers to the probability of successful packet decoding, which under finite blocklength theory is captured by the decoding error probability ε . Specifically, each predicted SINR $\gamma(t)$ maps to a BLER $\varepsilon(\gamma(t); R(t), n(t))$ that depends on the coding rate $R(t)$ and blocklength $n(t)$ (further details are included in section IV.C). Therefore, the prediction task must be formulated jointly: the model should learn the evolving SINR distribution while ensuring that the induced BLER satisfies the target reliability requirement. The reliability-constrained requirement in this joint framework is formally expressed as

$$\Pr_{\gamma(t) \sim f_\Gamma(\cdot; \theta_t)} \left\{ \varepsilon(\gamma(t); R(t), n(t)) \leq \varepsilon_{\text{target}} \right\} \geq 1 - \delta, \quad (2)$$

where $\varepsilon_{\text{target}}$ is the BLER target and $1 - \delta$ the required confidence level. The term $\varepsilon(\gamma(t); R(t), n(t))$ denotes the finite-blocklength theory BLER at SINR $\gamma(t)$ for the chosen $(R(t), n(t))$ [9]. The SINR prediction task is thus formulated as the following reliability-constrained learning problem:

$$\begin{aligned} \min_{\theta_t} \quad & -\log f_\Gamma(\gamma(t); \theta_t) \\ \text{s.t.} \quad & \Pr_{\gamma(t) \sim f_\Gamma(\cdot; \theta_t)} \left\{ \varepsilon(\gamma(t); R(t), n(t)) \leq \varepsilon_{\text{target}} \right\} \geq 1 - \delta, \end{aligned} \quad (3)$$

This formulation jointly addresses (i) prediction of the SINR distribution under sequential data, and (ii) enforcement of explicit reliability guarantees consistent with HRLLC requirements.

IV. PROBABILISTIC SINR PREDICTION

Based on the reliability-constrained formulation in section III, the proposed approach comprises two key components: (i) constraint relaxation via CVaR, (ii) non-parametric density modeling using a MDN, and a Lagrangian-based loss function derived accordingly.

A. CVaR-based Relaxation

The probabilistic constraint in (2) poses challenges due to its non-convexity and the intractability of computing gradients. To address this, we adopt the CVaR relaxation, which provides a convex surrogate by bounding the probability of violating the BLER constraint [10]. This relaxation ensures tractable

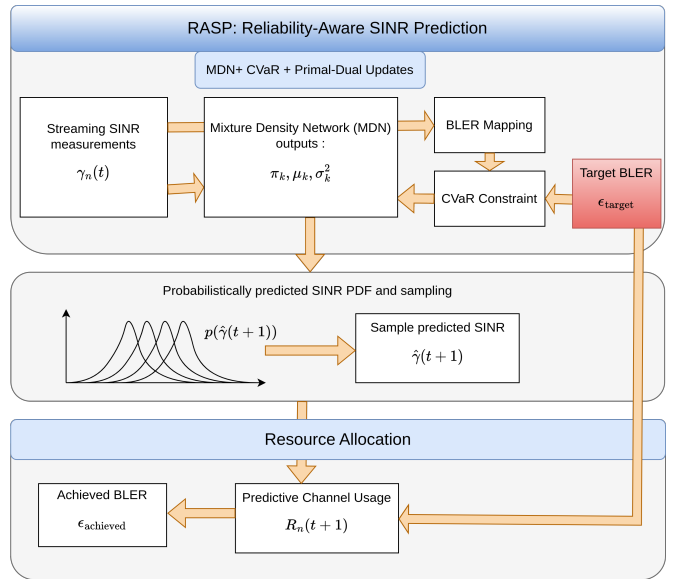


Fig. 2: Simplified schematics of the proposed framework.

optimization while preserving the reliability requirement. The violation variable is defined as:

$$v_t \triangleq \varepsilon(\gamma(t); R(t), n(t)) - \varepsilon_{\text{target}}, \quad (4)$$

and (2) is equivalently defined as a chance constraint as:

$$\Pr\{v_t \leq 0\} \geq 1 - \delta \iff \Pr\{v_t > 0\} \leq \delta, \quad (5)$$

with predicted SINR $\gamma(t) \sim f_\Gamma(\cdot; \theta_t)$. Consider a random variable Z and confidence level $\alpha \in (0, 1)$, CVaR is given by the Rockafellar–Uryasev representation as [10]:

$$\text{CVaR}_\alpha(Z) = \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{1 - \alpha} \mathbb{E}[(Z - \eta)_+] \right\}, \quad (6)$$

where $(x)_+ \triangleq \max\{x, 0\}$ and η is auxiliary variable. Here, $\text{VaR}_\alpha(Z)$ denotes the value-at-risk (VaR), i.e., the α -quantile of Z , meaning the largest threshold not exceeded with probability α . Intuitively, VaR identifies the boundary of the critical tail, whereas $\text{CVaR}_\alpha(Z)$ quantifies the expected value of Z within that tail region. Since $\text{CVaR}_\alpha(Z)$ is always greater than or equal to $\text{VaR}_\alpha(Z)$, it can be considered a conservative surrogate that guarantees the chance constraint [11]. Here, the random variable of interest is the violation v_t from (4), which measures how much the BLER exceeds the target. Accordingly, setting $Z = v_t$ and $\alpha = 1 - \delta$, the sufficient condition can be written as

$$\text{CVaR}_{1-\delta}(v_t) \leq 0 \implies \text{VaR}_{1-\delta}(v_t) \leq 0, \quad (7)$$

which shows that the CVaR condition is a conservative but tractable surrogate that guarantees the original chance constraint (5). From (6)–(7), the relaxed constraint is equivalently written via an auxiliary variable η_t as

$$\eta_t + \frac{1}{\delta} \mathbb{E}_{\gamma(t) \sim f_\Gamma(\cdot; \theta_t)} \left[(v_t - \eta_t)_+ \right] \leq 0. \quad (8)$$

As SINR measurement is considered as streaming data over successive TTI, the expectation in (8) is approximated using sequential samples. To capture temporal variation in industrial SNs, we employ an exponentially weighted estimator that assigns higher weight to recent observations while still incorporating historical information from past TTIs. The estimator is initialized with $\hat{\phi}_0(\cdot) = 0$ and updated recursively as $\hat{\phi}_t(\eta_t) = (1 - \beta)\hat{\phi}_{t-1}(\eta_t) + \beta(v_t - \eta_t)_+$. Specifically, the relaxed CVaR constraint with $v_t = \varepsilon(\gamma(t); R(t), n(t)) - \varepsilon_{\text{target}}$ is rewritten as

$$g_t(\eta_t) := \eta_t + \frac{1}{\delta}\hat{\phi}_t(\eta_t) \leq 0, \quad (9)$$

where $\hat{\phi}_{t-1}(\eta)$ is the recursively updated estimate from historical TTIs. The parameter $\beta = B_0(\omega \cdot \tau) \in [0, 1]$ is approximated using the zeroth-order Bessel function $B_0(\omega \cdot \tau)$, where $\omega = 2\pi f_d$ denotes the Doppler spread considering dynamics due to mobility of interferers [12]. This formulation ensures that the predictor remains adaptive to evolving SINR distributions over time.

B. Reliability-Aware Mixture Density Network

Building on the CVaR-based relaxation introduced in Subsection A, we embed this constraint into a predictive framework by employing a Mixture Density Network (MDN). The MDN predicts the evolving SINR distribution from streaming data while inherently enforcing the relaxed reliability constraint. Unlike conventional neural networks that output a single point estimate, an MDN produces the parameters of a mixture distribution, allowing it to model uncertainty and multimodality in the data [13] and illustrated in Fig. 2 highlighted in blue. Specifically, it combines a feedforward neural network with a Gaussian Mixture Model (GMM). The neural network takes the current input features and generates the parameters of the mixture, i.e., mixture weights, means, and variances. These parameters define the probability density function of the predicted SINR, thereby capturing both the local variability within the PDF, incorporating channel fading and industrial dynamics. At each time step t , the MDN outputs mixture weights $\pi_k(t; \theta_t)$, means $\mu_k(t; \theta_t)$, and variances $\sigma_k^2(t; \theta_t)$ for $k = 1, \dots, K$, with validity ensured through softmax, exponential, and linear activations, respectively. The predictive density is fitted to the observed SINR samples by minimizing the negative log likelihood (NLL). For an observation $\Gamma(t)$, the NLL $\mathcal{L}_{\text{NLL}}(\theta_t)$ is defined as follows:

$$\mathcal{L}_{\text{NLL}}(\theta_t) = -\log \left(\sum_{k=1}^K \pi_k(t; \theta_t) \cdot \mathcal{N}(\gamma(t) \mid \mu_k(t; \theta_t), \sigma_k^2(t; \theta_t)) \right). \quad (10)$$

As outlined earlier, the reliability constraint in (9) is enforced by computing the violation v_t , updating $\hat{\phi}_t(\eta_t)$, and constructing the surrogate. The complete reliability-aware Lagrangian loss is defined as

$$\mathcal{L}(\theta_t, \eta_t, \lambda_t) = \mathcal{L}_{\text{NLL}}(\theta_t) + \lambda_t g_t(\eta_t), \quad (11)$$

Algorithm 1 Online Training of RASP

Input: Streaming samples $(\gamma(t), R(t), n(t))$, stepsizes $(\alpha_\theta, \alpha_\eta, \alpha_\lambda)$, risk δ , smoothing β
Initialize:
 θ_0 (MDN weights), η_0 (VaR), $\lambda_0 \leftarrow 0$, $\hat{\phi}_0(\cdot) \leftarrow 0$
1: **for** $t = 1, 2, \dots$ **do**
2: Parameters $(\{\pi_k, \mu_k, \sigma_k^2\}_{k=1}^K) \leftarrow \text{MDN}(\theta_t) \triangleright \text{MDN forward}$
3: $\mathcal{L}_{\text{NLL}} \leftarrow -\log \left(\sum_{k=1}^K \pi_k \mathcal{N}(\gamma(t) \mid \mu_k, \sigma_k^2) \right) \triangleright \text{NLL loss}$
4: $v_t \leftarrow \varepsilon(\gamma(t); R(t), n(t)) - \varepsilon_{\text{target}} \triangleright \text{Violation}$
5: $\hat{\phi}_t(\eta_t) \leftarrow (1 - \beta)\hat{\phi}_{t-1}(\eta_t) + \beta(v_t - \eta_t)_+ \triangleright \text{EWMA update}$
6: $g_t(\eta_t) \leftarrow \eta_t + \frac{1}{\delta}\hat{\phi}_t(\eta_t) \triangleright \text{CVaR surrogate}$
7: $\mathcal{L} \leftarrow \mathcal{L}_{\text{NLL}} + \lambda_t g_t(\eta_t) \triangleright \text{Composite loss}$
8: $\theta_{t+1} \leftarrow \theta_t - \alpha_\theta \nabla_{\theta} \mathcal{L} \triangleright \text{Primal update on } \theta$
9: $\eta_{t+1} \leftarrow \eta_t - \alpha_\eta \nabla_{\eta} (\lambda_t g_t(\eta_t)) \triangleright \text{Primal update on } \eta$
10: $\lambda_{t+1} \leftarrow \max\{0, \lambda_t + \alpha_\lambda g_t(\eta_t)\} \triangleright \text{Projected dual update}$
11: **end for**
12: **Output:** Reliability-aware predictive density $f_{\Gamma}(\cdot; \theta_t)$

where η_t is the auxiliary VaR variable, $\lambda_t \geq 0$ the dual variable, $\beta \in (0, 1]$ the smoothing factor, and δ the risk level. The parameters are updated online via stochastic gradient descent, performing primal descent on (θ_t, η_t) and projected dual ascent on λ_t . This procedure ensures continual adaptation of the predictive distribution to streaming SINR samples while rigorously enforcing hyper-reliability. The detailed steps are outlined in Algorithm 1 and refer to the overall prediction technique as Reliability-Aware SINR Prediction (RASP). In addition, the proposed approach has bounded computational complexity per reporting interval. For an MDN with L fully connected layers of widths $\{h_\ell\}_{\ell=0}^L$ and a K -component Gaussian mixture output, each online update has time complexity $\mathcal{O}\left(\sum_{\ell=1}^L h_{\ell-1} h_\ell + K\right)$, dominated by the network forward and backward passes, while the reliability-related primal-dual updates are scalar and contribute only $\mathcal{O}(1)$ overhead.

C. Resource Allocation

We evaluate the accuracy of predicted SINR in terms of the achieved BLER under a target BLER and packet size D as an theoretical framework for resource allocation, using finite blocklength theory [9]. The required number of resource elements $R(t)$ for the link between AP-UE is calculated with the target BLER as a reliability constraint. According to [9], the required number of channel uses for transmitting D bits with decoding error probability $\varepsilon_{\text{target}}$ in an additive white Gaussian noise (AWGN) channel with predicted SINR $\hat{\gamma}(t)$

is approximated [14] as:

$$R(t) \approx \frac{D}{C(\hat{\gamma}(t))} + \frac{Q^{-1}(\varepsilon_{\text{target}})^2 V(\hat{\gamma}(t))}{2C(\hat{\gamma}(t))^2} \left[1 + \sqrt{1 + \frac{4DC(\hat{\gamma}(t))}{Q^{-1}(\varepsilon_{\text{target}})^2 V(\hat{\gamma}(t))}} \right], \quad (12)$$

where $C(\hat{\gamma}(t)) = \log_2(1 + \hat{\gamma}(t))$ denotes the Shannon capacity, $V(\cdot)$ the channel dispersion, and Q^{-1} the inverse Q-function. In order to compute the discrepancies between the target and achieved BLER, we use the predicted channel uses $R(t)$, derived from the predicted SINR, to determine the achieved BLER with genie-aided SINR $\tilde{\gamma}(t)$ as follows [9]:

$$\varepsilon_{\text{achieved}} := \Theta \left(\frac{R(t) \cdot C(\tilde{\gamma}(t)) - D}{\sqrt{R(t) \cdot V(\tilde{\gamma}(t))}} \right), \quad (13)$$

where $\Theta(\cdot)$ is the Gaussian Q-function. We then compare the achieved BLER, $\varepsilon_{\text{achieved}}$, with the target BLER. The perfect estimate of the interference, i.e., the genie-based predictor $\gamma(t)$, achieves the target BLER exactly with minimal channel uses. The overall simplified illustration of the proposed prediction technique and resource allocation framework is shown in Fig. 2.

V. NUMERICAL RESULTS

A. Baseline Methods

Following baseline methods are considered for comparison in this work:

- 1) *Genie Resource Allocation (RA)*: In this baseline, we consider the number of resource elements selection is based on the target BLER under the assumption that the prediction is perfect without any error. It is not a practical approach, but this reflects the optimal BLER with perfect prediction.
- 2) *Moving-average (MA) Predictor*: The SINR at every t is obtained as a weighted sum of SINRs estimated at TTIs $t-1$ and $t-2$. This method has been used in link adaptation for traditional enhanced mobile broadband (eMBB) services [14].
- 3) *LSTM Predictor*: To capture temporal dependencies in both directions, we employ a bidirectional LSTM (Bi-LSTM). The gating mechanisms of the LSTM units (input, forget, and output) facilitate stable learning and mitigate vanishing gradients. Such architectures are widely adopted in sequence prediction tasks, including interference and channel prediction [15].

B. Performance Analysis

1) *Measurement Setup*: Measurements were carried out in a $10\text{m} \times 10\text{m}$ experiment hall at TU Dortmund University (Fig. 1), resembling a scaled-down industrial environment.

The network under test is a commercial telco-grade private 5G system configured with 50 MHz bandwidth in the n78 band at 3.775 GHz. The Device Under Test (DUT) is a

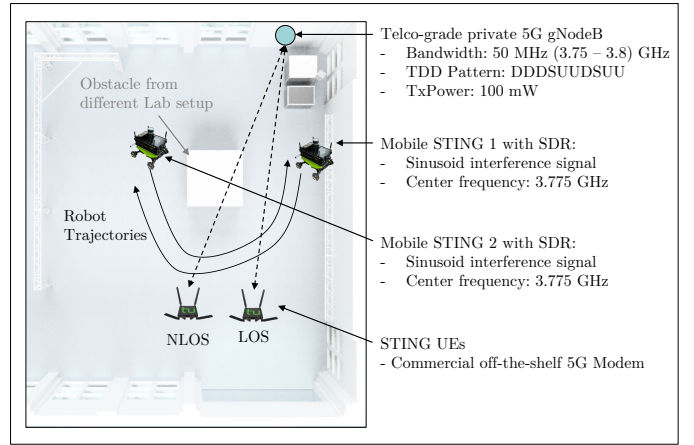


Fig. 3: Schematic setup of the experiment.

STING UE [16] equipped with a COTS modem, representative of industrial deployments. Interference was generated by a mobile STING unit based on a mobile robotic platform, with an SDR transmitting sinusoidal signals of varying bandwidths which was also used in [17] where an AI based image processing algorithm was used for anomaly detection. Here, this setup is used to generate co-subband interference where interferers operate within the same subband as the UE. The SINR traces used for the following prediction are recorded from a COTS 5G modem, which assures applicability for real-world deployments. The modem reports the last experienced SS-SINR and was queried every 150 ms. To ensure constant updates, a User Datagram Protocol (UDP) downlink transmission was active during measurements using iperf3. Traces have been recorded in three distinct scenarios:

- **Scenario 1**: DUT in **LOS** with the gNodeB; **One** robot with interference generation moving in a "U" shape around the obstacle in the middle of the scenario (cf. Fig. 3)
- **Scenario 2**: DUT in **LOS** with the gNodeB; **Two** robots with interference generation moving in a "U" shape in the respective opposite direction.
- **Scenario 3**: DUT in **NLOS** with the gNodeB; **Two** robots with interference generation

Based on these realistic SINR samples collected at the UE, we evaluate the performance of RASP against baselines in the following subsection.

C. Results

In this section, we evaluate the link reliability within the above-introduced scenarios based on section IV-C using a finite blocklength resource allocation framework [14]. In Scenario 1, with interferers moving at 0.4 m/s, we consider a packet size of $D = 160$ bits and a transmission rate of 5.33 bit/s/Hz, corresponding to the average interference-free SNR. For a target BLER of 10^{-6} , the violation probability across 1000 test samples is 12.4%. This high rate shows that methods relying only on historical SINR statistics cannot meet




Fig. 4: Variability of SINR and prediction from RASP.

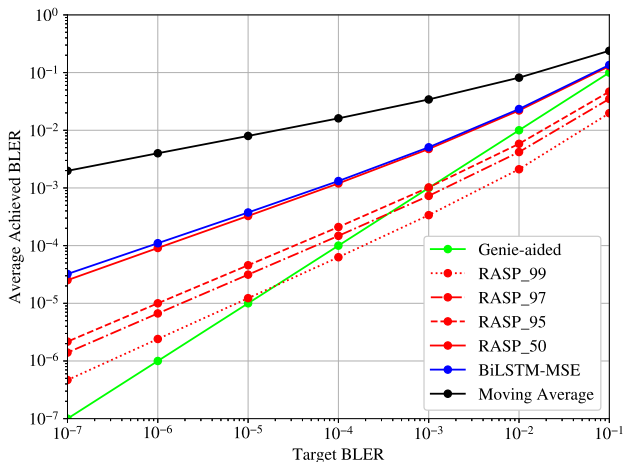


Fig. 5: Average achieved BLER with target BLER

HRLLC reliability demands, motivating the proposed RASP framework, which embeds explicit reliability constraints into prediction to ensure stringent BLER targets.

The MDN in RASP consists of three 64-node feedforward layers with $K = 5$ Gaussian mixture components, trained incrementally on streaming data. It estimates the PDF parameters of the SINR, from which 1000 samples are drawn to capture tail statistics. Fig. 4 illustrates the measured SINR at the UE together with RASP predictions at the 50th and 95th quantiles. Unlike LSTM and quantile regression baselines that yield only point estimates, RASP predicts the full distribution, offering flexible reliability assessment.

Using predicted SINR, we evaluate the achieved BLER for target values from 10^{-1} to 10^{-7} , as shown in Fig. 5. The moving average fails to track local interference, while BiLSTM reduces the gap to the target BLER. In contrast, the proposed RASP predicts the full PDF, allowing quantification at any chosen quantile. RASP-95 refers to the reliability-aware decision rule where the lower quantile of the predicted SINR distribution corresponding to a 95% reliability level is used for resource allocation. A clear trend emerges: moving toward conservative reliability levels (95th, 97th, 99th) better captures tail interference, thereby narrowing the gap to the target BLER and improving reliability. We next focus on the case of 10^{-6} using a full Empirical cumulative distribution function (ECDF) analysis. As illustrated in Fig. 6, incorporating and quantifying

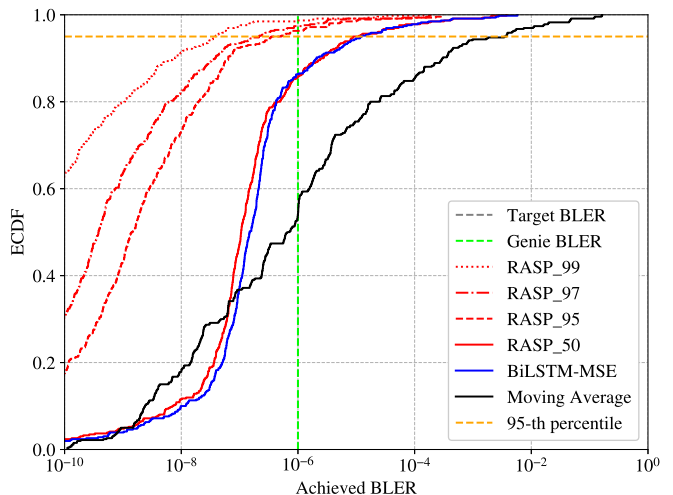


Fig. 6: ECDF of achieved BLER with target BLER of 10^{-6} .

Methods	Scenario-1	Scenario-2	Scenario-3
RASP_95	1.66×10^{-7}	6.57×10^{-7}	1.51×10^{-7}
RASP_50	2.03×10^{-5}	1.99×10^{-5}	6.54×10^{-5}
Bi_LSTM	1.33×10^{-5}	2.12×10^{-5}	5.69×10^{-5}

TABLE I: 95th percentile of achieved BLER across three scenarios with target BLER of 10^{-6} .

interference in the tail—such as at the 95th quantile—provides a more accurate characterization of reliability. This tail-aware approach reduces the gap to the target BLER and helps achieve stringent reliability levels more effectively. When comparing reporting intervals for Scenario-1, for a target BLER of 10^{-6} , the achieved BLER was 1.07×10^{-6} at 150 ms and 0.53×10^{-6} at 300 ms, with the latter requiring 3% more resources as a result of the longer reporting interval by RASP_95. Table I highlights the challenge of achieving hyper-reliability at the 95th percentile. While BiLSTM with mean square error (MSE) and RASP_50, which rely on mean or median prediction, fall short of the 10^{-6} target, RASP exploits the full predicted PDF. Although the scenarios are measured within the same test environment, it is shown that the results follow the same characteristics for different configurations of interference. By leveraging tail statistics, RASP_95 consistently meets the target BLER, confirming its effectiveness as a reliability-aware predictor for hyper-reliable communication.

VI. CONCLUSIONS

This paper introduced RASP, a reliability-aware SINR prediction framework that enforces BLER constraints through distributional prediction. Using real industrial measurements, RASP consistently achieved the 10^{-6} BLER target at the 95th percentile, outperforming mean-based predictors. Future work will extend RASP to multi-user scheduling and large-scale industrial deployments. Additionally, advancing MDN with a Transformer network or a recurrent neural network as the backbone will be evaluated, investigating robustness

under heterogeneous traffic patterns, and exploring integration with federated or distributed learning to support large-scale industrial deployments.

REFERENCES

- [1] ITU-R, "Framework and overall objectives of the future development of IMT for 2030 and beyond," *International Telecommunication Union (ITU) Recommendation (ITU-R)*, 2023.
- [2] R. Adeogun, G. Berardinelli, P. E. Mogensen, I. Rodriguez, and M. Razzaghpour, "Towards 6G in-X subnetworks with sub-millisecond communication cycles and extreme reliability," *IEEE Access*, vol. 8, pp. 110 172–110 188, 2020.
- [3] P. Gautam, S. Sapkota, C. Bockelmann, S. Raj Pandey, and A. Dekorsy, "Extreme Value Theory-Based Distributed Interference Prediction for 6G Industrial Sub-Networks," *IEEE Open Journal of the Communications Society*, vol. 6, pp. 9121–9143, 2025.
- [4] I. Hadj-Kacem, S. B. Jemaa, H. Braham, and A. M. Alam, "Sinr prediction in presence of correlated shadowing in cellular networks," *IEEE Transactions on Wireless Communications*, vol. 21, no. 10, pp. 8744–8756, 2022.
- [5] P. Gautam, C. Bockelmann, and A. Dekorsy, "Interference Prediction in Unconnected In-X Mobile 6G Subnetworks Using a Data-Driven Approach," in *2024 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2024, pp. 2046–2052.
- [6] S. B. Mallikarjun, S. C. Kusumapani, N. P. Kuruvatti, B. G. Bhat, and H. D. Schotten, "Machine learning based sinr prediction in private campus networks," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*. IEEE, 2023, pp. 1–6.
- [7] Q. Zhou, W. Jiang, D. Wang, and H. D. Schotten, "Deep learning-based signal-to-noise ratio prediction for realistic wireless communication," in *2022 IEEE 95th vehicular technology conference:(VTC2022-Spring)*. IEEE, 2022, pp. 1–5.
- [8] 3GPP, "5G; NR; Physical layer measurements," 3rd Generation Partnership Project, Technical Specification (TS) 38.215, Jan. 2025.
- [9] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.
- [10] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.
- [11] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, "Risk-sensitive and robust decision-making: a cvar optimization approach," *Advances in neural information processing systems*, vol. 28, 2015.
- [12] P. Gautam, R. S. B. A. G *et al.*, "Dynamic Interference Prediction for In-X 6G Sub-Networks," in *2025 14th International ITG Conference on Systems, Communications and Coding (SCC)*, 2025, pp. 1–6.
- [13] C. M. Bishop, "Mixture density networks," 1994.
- [14] N. H. Mahmood, O. A. López, H. Alves, and M. Latva-Aho, "A Predictive Interference Management Algorithm for URLLC in Beyond 5G Networks," *IEEE Communications Letters*, vol. 25, no. 3, pp. 995–999, 2021.
- [15] P. Gautam, M. Vakili, C. Bockelmann, and A. Dekorsy, "Cooperative interference estimation using lstm-based federated learning for in-x subnetworks," in *IEEE GLOBECOM 2023*.
- [16] C. Arendt, S. Böcker, C. Bektas, and C. Wietfeld, "Better safe than sorry: Distributed testbed for performance evaluation of private networks," in *IEEE Future Networks World Forum, FNWF'22*, Montreal, Canada, Oct. 2022.
- [17] K.-I. Šabanović, C. Arendt, S. Fricke, M. Geis, S. Böcker, and C. Wietfeld, "AI-Based Anomaly Detection for Industrial 5G Networks by Distributed SDR Measurements," in *2024 IEEE International Symposium on Measurements & Networking (M&N)*. Rome, Italy: IEEE, Jul. 2024, pp. 1–5.