




Beam-Slicing for Jammer Mitigation in mmWave Massive MU-MIMO

Conference Paper**Author(s):**

[Castañeda Fernández, Oscar](#) ; [Marti, Gian](#) ; [Studer, Christoph](#) 

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000499937>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/SiPS52927.2021.00039>

Beam-Slicing for Jammer Mitigation in mmWave Massive MU-MIMO

Oscar Castañeda*, Gian Marti*, and Christoph Studer

Department of Information Technology and Electrical Engineering, ETH Zürich, Switzerland

Abstract—Millimeter-wave (mmWave) massive multi-user multiple-input multiple-output (MU-MIMO) technology promises unprecedentedly high data rates for next-generation wireless systems. To be practically viable, mmWave massive MU-MIMO basestations (BS) must (i) rely on low-resolution data-conversion and (ii) be robust to jammer interference. This paper considers the problem of mitigating the impact of a permanently transmitting jammer during uplink transmission to a BS equipped with low-resolution analog-to-digital converters (ADCs). To this end, we propose SNIPS, short for Soft-Nulling of Interferers with Partitions in Space. SNIPS combines beam-slicing—a localized, analog spatial transform that focuses the jammer energy onto a subset of all ADCs—together with a soft-nulling data detector that exploits knowledge of which ADCs are contaminated by jammer interference. Our numerical results show that SNIPS is able to successfully serve 65% of the user equipments (UEs) for scenarios in which a conventional antenna-domain soft-nulling data detector is only able to serve 2% of the UEs.

I. INTRODUCTION

In order to meet the ever-growing demand for higher data rates, next-generation wireless communication systems are expected to rely on the vast amount of unused bandwidth available at millimeter wave (mmWave) frequencies [1]. Communication at mmWave frequencies is characterized by a high path loss that can be compensated for with massive multiple-input multiple-output (MIMO) technology. Besides providing the basestation (BS) with a high array gain, massive MIMO also enables multi-user (MU) communication [2].

The deployment of a BS equipped with a large number of antennas and corresponding radio-frequency (RF) chains poses implementation challenges in terms of system costs, power consumption, and circuit complexity. A potential solution is to use low-resolution data converters that (i) reduce the power consumption of data conversion and (ii) simplify the linearity and noise requirements of the RF chain, which in turn translates into power consumption and circuit complexity savings [3].

Unfortunately, the use of low-resolution analog-to-digital converters (ADCs) leaves the BS vulnerable to jammers that could be introduced, for example, by a rogue user equipment (UE) or a malicious transmitter. Previous works [4]–[13] have analyzed the impact of different types of jamming attacks on

massive MU-MIMO systems and proposed mitigation methods based on digital equalization. However, none of these works have taken into consideration the compounding challenge of low-resolution data conversion: A jammer can either saturate low-resolution ADCs or (if gain-control is used) widen their quantization range to an extent that drowns the signal in quantization noise, and thereby introduce distortions that are difficult to remove with subsequent digital equalization.

A. Contributions

In this work, we propose novel means to limit the impact of a single, permanently transmitting jammer on mmWave massive MU-MIMO systems while taking into account the nonlinear distortions caused by ADCs. Our numerical results indicate that, as long as the ADCs have a resolution that is sufficiently high (e.g., 8 bits for a 25 dB jammer), the jammer can be effectively mitigated by using a linear equalizer in the digital domain that utilizes an estimate of the jammer covariance matrix. However, practical deployments of massive MU-MIMO are likely to rely on low-resolution ADCs. In such situations, the jammer will force the ADCs' quantization range to drown the UE signals in quantization noise. In order to enable jammer-robust communication with low-resolution ADCs, we propose a novel method called *Soft-Nulling of Interferers with Partitions in Space (SNIPS)*. SNIPS partitions the space using a non-adaptive, localized spatial transform prior to the data-conversion step. We refer to this transform as *beam-slicing*. Thanks to the strong directionality of mmWave signals, beam-slicing focuses the jammer's energy onto a subset of all ADCs. An estimate of how strong each ADC is affected by the jammer is then used by the data detector to generate improved estimates of the transmitted data.

B. Related Prior Work

Several works have studied means that improve the resiliency of MIMO systems against jamming attacks. These works have considered different attacks, such as constant jamming attacks [4], [5], in which the jammer is permanently transmitting, as well as other types of attacks in which the jammer transmits only at specific time instances, such as when the UEs transmit [4] or during pilot transmission [6]. Furthermore, due to the complexity of the jammer problem, some works have devoted themselves only to detecting the presence of a jammer [6], [7], while other works have proposed methods to suppress jamming [4], [5], [8]–[13].

*OC and GM contributed equally to this work.

The work of OC and CS was supported in part by ComSenTer, one of six centers in JUMP, a SRC program sponsored by DARPA. The work of CS was also supported by an ETH Research Grant and by the US National Science Foundation (NSF) under grants CNS-1717559 and ECCS-1824379.

The authors thank D. Nonaca for helpful comments on this manuscript. Contact authors: OC and GM; e-mail: caoscar@ethz.ch, gimarti@ethz.ch

In our work, we focus on mitigating the interference of a permanently transmitting jammer. We now describe existing approaches that are related to handling the jammer’s interference. In [4], a method which uses the angle of arrival of the interfering signal for projecting the receive vector onto the subspace orthogonal to the interference is proposed for small-scale MIMO systems. Also in the context of small-scale MIMO, reference [5] proposes a method that uses differential encoding and exploits the ratio between channel coefficients. Turning to massive MIMO systems, reference [8] uses random matrix theory to estimate the UEs’ eigensubspace to then project the received signals onto that subspace, while [9] proposes methods which require perfect channel state information and cooperation between UEs and BS. Moreover, references [10], [11] utilize an estimate of the jammer channel to implement different versions of a jammer-robust zero-forcing detector.

Similarly to our work, references [12], [13] propose to exploit spatially correlated channels in order to suppress jammer interference. In particular, the work in [13] applies a beamspace transform to separate the jammer from the UEs in the angular domain. The beamspace transform [14] is related to the concept of beam-slicing used by SNIPS. As a matter of fact, the conventional beamspace transform (and even the absence of any spatial transform, i.e., the antenna domain) can be formulated as a special case of beam-slicing, and one can think in general of beam-slicing as a beamspace transform with a coarser angular resolution. The key advantage of beam-slicing over the beamspace transform is that the former is composed of *localized* transforms which only take inputs from a few adjacent antennas, making such transforms more amenable for analog circuit implementation [15]. Another key difference of SNIPS is that it takes into account the effects of low-resolution quantization. Most of the existing results do not consider the effects of hardware impairments—with one exception: Reference [12] models the effects of hardware impairments (including quantization errors) as additive Gaussian noise, hence failing to accurately model the signal- and jammer-dependent distortions introduced by coarse ADCs.

C. Notation

Matrices and column vectors are represented by boldface uppercase and lowercase letters, respectively. For a matrix \mathbf{A} , the conjugate transpose is \mathbf{A}^H , the k th column is \mathbf{a}_k , and the Frobenius norm is $\|\mathbf{A}\|_F$. The $N \times N$ identity matrix and discrete Fourier transform (DFT) are \mathbf{I}_N and \mathbf{F}_N , respectively, where $\mathbf{F}_N^H \mathbf{F}_N = \mathbf{I}_N$. For a vector \mathbf{a} , the k th entry is a_k , the ℓ_2 -norm is $\|\mathbf{a}\|_2$, the real part is $\Re\{\mathbf{a}\}$, and the imaginary part is $\Im\{\mathbf{a}\}$. Furthermore, $\text{diag}(\mathbf{a})$ is a diagonal matrix whose diagonal is formed by \mathbf{a} . Expectation with respect to the random vector \mathbf{x} is denoted by $\mathbb{E}_{\mathbf{x}}[\cdot]$. The floor function $\lfloor x \rfloor$ returns the greatest integer less than or equal to x . We define $i^2 = -1$.

II. PROPAGATION MODEL

We consider the uplink of a mmWave massive MU-MIMO system in which U single-antenna UEs transmit data to a B antenna BS, while a permanently transmitting, single-antenna

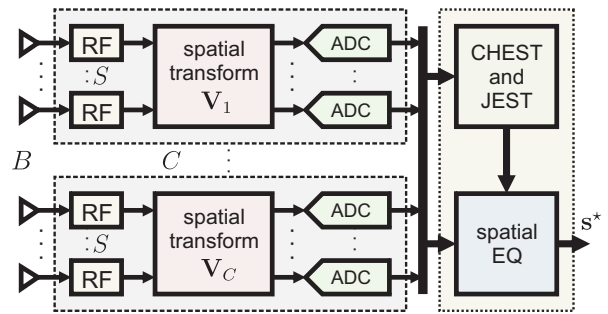


Fig. 1. System overview of SNIPS: The B RF chains are clustered into C clusters of size S . The analog baseband signals are transformed cluster-wise to the beam-slice domain before being converted to the digital domain. SNIPS then performs jammer interference estimation, channel estimation, and data detection in digital beam-slice domain.

jammer interferes with the BS receive signal. For this scenario, we consider the following frequency-flat input-output relation:

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{j}w + \mathbf{n}. \quad (1)$$

Here, $\mathbf{y} \in \mathbb{C}^B$ is the (unquantized) vector received by the BS antennas, $\mathbf{H} \in \mathbb{C}^{B \times U}$ models the MIMO uplink channel matrix, $\mathbf{s} \in \mathcal{S}^U$ is the transmit vector whose (independent) entries correspond to the per-UE transmit symbols which take value in a constellation set \mathcal{S} (e.g., 16-QAM), $\mathbf{j} \in \mathbb{C}^B$ is the channel from the jammer to the BS, $w \in \mathbb{C}$ is the jamming signal, and $\mathbf{n} \in \mathbb{C}^B$ is i.i.d. circularly-symmetric complex Gaussian noise with a per-entry variance of N_0 . In what follows, we assume that the UE transmit symbols s_u , $u = 1, \dots, U$, are independent and zero mean with variance E_s so that $\mathbb{E}_{\mathbf{s}}[\mathbf{s}\mathbf{s}^H] = E_s \mathbf{I}_U$. We model the jamming signal w as circularly-symmetric complex Gaussian with variance E_w . All probabilistic quantities are assumed to be mutually independent.

III. SNIPS: SOFT-NULLING OF INTERFERERS WITH PARTITIONS IN SPACE

Our proposed method aims to protect most of the ADCs from the jammer by exploiting the strong spatial directivity of mmWave signals. Prior to analog-to-digital (A/D)-conversion, we apply a spatial transform that resolves the incident waves so that only a few ADCs are strongly affected by the jammer. One may then discount the outputs of these jammer-distorted ADCs during equalization and detect the data symbols mainly based on the outputs of the distortion-free ADCs.

A naïve approach would be to transform the (unquantized) receive vector \mathbf{y} into the beamspace (or angular) domain [14] using a DFT according to $\mathbf{y}_B = \mathbf{F}_B \mathbf{y}$, and set the entries of the beamspace vector \mathbf{y}_B dominated by jammer interference to zero. Accordingly, the resulting vector $\mathbf{y}_{B,\text{mask}}$ will have entries equal to zero if they belong to jammer-contaminated beams, and otherwise equal to the corresponding entry in \mathbf{y}_B . One could then equalize $\mathbf{y}_{B,\text{mask}}$ as if neither interference nor interference-cancellation had occurred, for instance with linear minimum mean-square error (LMMSE) estimation,




Fig. 2. Effect of $S=4$ beam-slicing for a $B=8$ antenna array considering one UE and one jammer (both in far-field). The spatial transform \mathbf{V}_1 illustrates the 4 beams (a), (b), (c), and (d) of \mathbf{F}_4 , of which (b) is perfectly aligned to capture the UE's transmitted power, while beams (c) and (d) partially capture the jammer's power. Such partial capturing would lead to unsatisfactory jammer mitigation. To increase angular diversity, the spatial transform \mathbf{V}_2 first rotates the received signal so that the 4 beams (a), (b), (c), and (d) of the resulting transform are also shifted. As a result, the shifted beam (d) is now able to capture the jammer's power, which allows for improved jammer mitigation.

$\mathbf{s}^* = (\mathbf{H}^H \mathbf{H} + N_0/E_s \mathbf{I}_B)^{-1} \mathbf{H}^H \mathbf{y}_{B,\text{mask}}$. While such an approach would be effective in suppressing jammers, implementing large spatial transforms in the analog domain is nontrivial, especially when considering hundreds of BS antennas [15].

SNIPS is a jammer-mitigation method that relies on a distributed, localized, and hence small analog transform. SNIPS does not discard the outputs of jammer-affected ADCs completely, but instead takes into account each ADC-output's fidelity by estimating the amount of jammer interference at the individual ADCs. Moreover, SNIPS utilizes Busgang's decomposition in order to take into account the effects of practical, low-resolution ADCs. Figure 1 illustrates the complete SNIPS pipeline, which we detail hereafter.

A. Beam-Slicing

Beam-slicing transforms partitions of the (unquantized) receive vector \mathbf{y} into shifted angular domains. Beam-slicing is fully analog, non-adaptive, and operates in decentralized fashion. Specifically, we partition the BS antenna array into C equisized clusters, each consisting of $S = B/C$ adjacent BS antennas. The corresponding partitioning of the receive vector is denoted $\mathbf{y} = [\mathbf{y}_1^T, \dots, \mathbf{y}_C^T]^T$, where $\mathbf{y}_c \in \mathbb{C}^S, c = 1, \dots, C$. Beam-slicing then transforms each cluster into the beam-slice domain as follows:

$$\hat{\mathbf{y}}_c = \mathbf{V}_c \mathbf{y}_c, \quad c = 1, \dots, C. \quad (2)$$

Here, the c th cluster matrix \mathbf{V}_c is given as a progressively phase-shifted S -point DFT matrix \mathbf{F}_S according to

$$\mathbf{V}_c = \mathbf{F}_S \text{diag}\left(1, \dots, e^{-i\frac{2\pi}{S}(c-1)(s-1)}, \dots, e^{-i\frac{2\pi}{S}(c-1)(S-1)}\right). \quad (3)$$

Such phase-rotated DFTs increase the ‘‘angular diversity’’ of beam-slicing to better capture the possible directions of

jammers—see Figure 2 for a graphical explanation with two clusters. The action of beam-slicing is summarized as

$$\hat{\mathbf{y}} = \mathbf{V} \mathbf{y} = \begin{bmatrix} \mathbf{V}_1 \mathbf{y}_1 \\ \vdots \\ \mathbf{V}_C \mathbf{y}_C \end{bmatrix}, \quad (4)$$

where $\mathbf{V} = \text{diag}(\mathbf{V}_1, \dots, \mathbf{V}_C)$ and $\mathbf{V}^H \mathbf{V} = \mathbf{I}_B$. We also point out that for an (impractical) cluster size $S = B$, beam-slicing corresponds to performing a conventional beamspace transform. In what follows, it will be convenient to define the *beam-sliced channel matrix* $\hat{\mathbf{H}} = \mathbf{V} \mathbf{H}$, and the *beam-sliced jammer channel* $\hat{\mathbf{j}} = \mathbf{V} \mathbf{j}$, which allows us to rewrite (1) as

$$\hat{\mathbf{y}} = \hat{\mathbf{H}} \mathbf{s} + \hat{\mathbf{j}} w + \hat{\mathbf{n}}, \quad (5)$$

where $\hat{\mathbf{n}} = \mathbf{V} \mathbf{n}$ has the same distribution as \mathbf{n} .

B. Data Conversion

The beam-sliced signal is then converted to the digital domain. To take into account the quantization errors of low-resolution ADCs, we assume that the beam-sliced vector $\hat{\mathbf{y}}$ is quantized as

$$\mathbf{r} = \mathbf{G}^{-1} (Q(\Re\{\mathbf{G}\hat{\mathbf{y}}\}) + iQ(\Im\{\mathbf{G}\hat{\mathbf{y}}\})). \quad (6)$$

Here, $\mathbf{G} = \text{diag}(g_1, \dots, g_B)$ is a diagonal matrix that represents beam-wise gain-control. The quantization function $Q(\cdot)$ is applied entry-wise to its input and represents a q -bit uniform midrise quantizer with step size Δ defined as

$$Q(x) \triangleq \begin{cases} \Delta \lfloor \frac{x}{\Delta} \rfloor + \frac{\Delta}{2}, & \text{if } |x| < \Delta 2^{q-1} \\ \frac{\Delta}{2} (2^q - 1) \frac{x}{|x|}, & \text{if } |x| > \Delta 2^{q-1}. \end{cases} \quad (7)$$

For the quantizer's step size Δ , we use the value which minimizes the mean-square error (MSE) between the quantizer's output $Q(x)$ and its input x under the assumption that x is standard normal [16]. For convenience, we will denote (6) as

$$\mathbf{r} = \mathcal{Q}(\hat{\mathbf{y}}). \quad (8)$$

The per-beam gains g_b aim to ensure that the values entering the quantizers have unit variance per real dimension and are obtained from a set of T training vectors $\tilde{\mathbf{Y}} = [\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_T]$ as

$$g_b = \sqrt{\frac{2T}{\|\tilde{\mathbf{y}}_{(b)}\|_2^2}}, \quad b = 1, \dots, B, \quad (9)$$

where $\tilde{\mathbf{y}}_{(b)}$ is the b th row of $\tilde{\mathbf{Y}}$.

C. Jammer Interference Estimation

Our soft-nulling data detection scheme (see below) treats the jamming term $\hat{\mathbf{j}} w$ in (5) as spatially correlated noise. For this, we need to know its covariance matrix $\mathbf{C}_j = \mathbb{E}_w [\hat{\mathbf{j}} w (\hat{\mathbf{j}} w)^H]$. We suggest to estimate \mathbf{C}_j from a number of channel uses during which the UEs do not transmit any symbols, and where the jammer transmits i.i.d. jamming symbols $[w_1, \dots, w_N]$, so the quantized, beam-sliced receive matrix is

$$\mathbf{R}_J = \mathcal{Q}(\hat{\mathbf{Y}}_J) \text{ with } \hat{\mathbf{Y}}_J = \hat{\mathbf{j}} [w_1, \dots, w_N] + \hat{\mathbf{N}}. \quad (10)$$

In order to learn the jammer channel, we propose to estimate the gain matrix \mathbf{G} with (9) directly from the received signals, $\mathbf{G} = \mathbf{G}(\hat{\mathbf{Y}}_J)$. Our estimate $\mathbf{\Lambda}$ of the covariance matrix \mathbf{C}_j is

$$\mathbf{\Lambda} = \frac{1}{N} \mathbf{R}_J \mathbf{R}_J^H. \quad (11)$$

D. Channel Estimation

We estimate the UEs' channel matrix using a pilot-based least squares (LS) estimator from U orthogonal pilot sequences $\mathbf{S}_P = [\mathbf{s}_1, \dots, \mathbf{s}_U]$. The channel estimation pipeline passes through the beam-slicer and the quantizer. The beam-sliced receive matrix and its corresponding quantization output are

$$\hat{\mathbf{Y}}_P = \hat{\mathbf{H}} \mathbf{S}_P + \hat{\mathbf{j}}[w_1, \dots, w_U] + \hat{\mathbf{N}} \quad (12)$$

$$\mathbf{R}_P = \mathcal{Q}(\hat{\mathbf{Y}}_P), \quad (13)$$

where we estimate the gain matrix \mathbf{G} with (9) from the pilot sequence itself, $\hat{\mathbf{Y}} = \hat{\mathbf{Y}}_P$. (We fix this choice of $\mathbf{G} = \mathbf{G}(\hat{\mathbf{Y}}_P)$ also for the data detection phase described below.) We then estimate the channel directly in the beam-slice domain with an LS channel estimate:

$$\hat{\mathbf{H}}_{\text{est}} = \mathbf{R}_P \mathbf{S}_P^H (\mathbf{S}_P \mathbf{S}_P^H)^{-1} \quad (14)$$

$$\stackrel{(a)}{=} \frac{1}{UE_s} \mathbf{R}_P \mathbf{S}_P^H, \quad (15)$$

where (a) holds because the pilot sequence is orthogonal.

E. Soft-Nulling Data Detection

Our soft-nulling data detector aims to detect the UE signal by means of linear equalization. In the detector's derivation, we will make certain idealizing assumptions under which it would be the LMMSE estimator.

So far (i.e., for jammer covariance estimation and channel estimation), we have neglected the distortion introduced by the quantization step in (6)–(8). We do not, however, neglect this distortion for the data detection step as quantization introduces distortions which are correlated with the quantizer inputs. We assume that the components of the quantizer inputs are real-valued zero-mean unit-variance Gaussian which allows us to perform a component-wise Bussgang decomposition [17] of the quantization signal as follows:

$$Q(x) = \gamma x + d. \quad (16)$$

Here, γ is the quantizer's *Bussgang gain*, and the distortion d has zero mean and is uncorrelated with x . The Bussgang gain and the second moment of the distortion are given by

$$\gamma = \frac{\mathbb{E}[Q(x)x]}{\mathbb{E}[x^2]} \quad (17)$$

and

$$D = \mathbb{E}[d^2] = \mathbb{E}[Q(x)^2] - \gamma^2 \mathbb{E}[x^2], \quad (18)$$

respectively. The Bussgang decomposition allows us to rewrite (8) as

$$\mathbf{r} = \mathcal{Q}(\hat{\mathbf{y}}) \quad (19)$$

$$= \mathbf{G}^{-1} (Q(\Re\{\mathbf{G}\hat{\mathbf{y}}\}) + iQ(\Im\{\mathbf{G}\hat{\mathbf{y}}\})) \quad (20)$$

$$= \mathbf{G}^{-1} (\gamma \Re\{\mathbf{G}\hat{\mathbf{y}}\} + \mathbf{d}_r + i(\gamma \Im\{\mathbf{G}\hat{\mathbf{y}}\} + \mathbf{d}_i)) \quad (21)$$

$$= \gamma \mathbf{G}^{-1} (\Re\{\mathbf{G}\hat{\mathbf{y}}\} + i\Im\{\mathbf{G}\hat{\mathbf{y}}\}) + \mathbf{G}^{-1}(\mathbf{d}_r + i\mathbf{d}_i) \quad (22)$$

$$= \gamma \hat{\mathbf{y}} + \mathbf{G}^{-1} \mathbf{d}, \quad (23)$$

where we define $\mathbf{d} = \mathbf{d}_r + i\mathbf{d}_i$. Based on (18), we make the idealized assumption that the covariance matrix of \mathbf{d} is

$$\mathbf{C}_d = \mathbb{E}_d[\mathbf{d}\mathbf{d}^H] \approx 2D \mathbf{I}_B. \quad (24)$$

We now combine (23) with (5), which yields the input-output relation

$$\mathbf{r} = \gamma \hat{\mathbf{H}} \mathbf{s} + \gamma \hat{\mathbf{j}} w + \gamma \hat{\mathbf{n}} + \mathbf{G}^{-1} \mathbf{d}. \quad (25)$$

Our linear equalizer is

$$\mathbf{s}^* = \mathbf{W} \mathbf{r}, \quad (26)$$

where the matrix \mathbf{W} is given as

$$\mathbf{W} = \gamma E_s \hat{\mathbf{H}}_{\text{est}}^H (\gamma^2 E_s \hat{\mathbf{H}}_{\text{est}} \hat{\mathbf{H}}_{\text{est}}^H + \gamma^2 \mathbf{\Lambda} + \gamma^2 N_0 \mathbf{I}_B + 2D \mathbf{G}^{-2})^{-1} \quad (27)$$

$$= \frac{1}{\gamma} \hat{\mathbf{H}}_{\text{est}}^H \left(\hat{\mathbf{H}}_{\text{est}} \hat{\mathbf{H}}_{\text{est}}^H + \frac{1}{E_s} (\mathbf{\Lambda} + N_0 \mathbf{I}_B + 2D \gamma^{-2} \mathbf{G}^{-2}) \right)^{-1}. \quad (28)$$

Here, we use again the gain control matrix acquired during the pilot phase, $\mathbf{G} = \mathbf{G}(\hat{\mathbf{Y}}_P)$. If the diagonal approximation (24) and the approximations $\hat{\mathbf{H}}_{\text{est}} \approx \hat{\mathbf{H}}$ and $\mathbf{\Lambda} \approx \mathbf{C}_j$ were exact, then equation (26) would implement the LMMSE estimator.

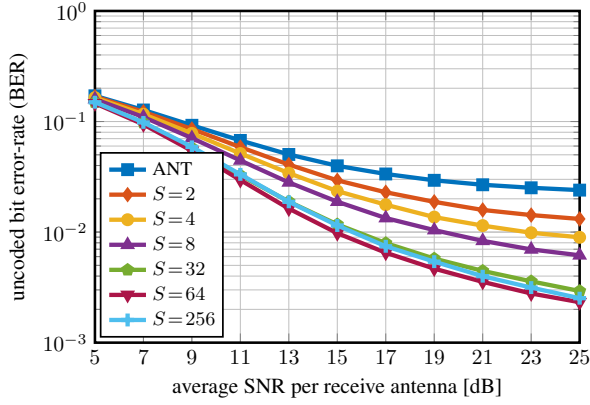
IV. RESULTS

We now demonstrate the efficacy of SNIPS by comparing it with a baseline which differs from SNIPS only in lacking the analog beam-slicing stage. We note that the operation of this baseline corresponds to SNIPS with cluster size $S = 1$, which implies $\mathbf{V} = \mathbf{I}_B$: This baseline performs A/D-conversion and soft-nulling directly in the antenna domain. We will show that in the presence of a strong jammer, beam-slicing with a cluster size of $S = 2$ already yields significant improvements over this baseline. We will also analyze in which cases a beam-slicer is required considering different ADC resolutions and levels of jamming power.

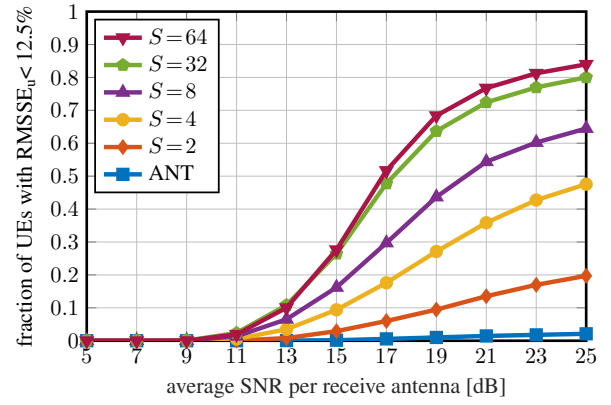
A. Simulation Setup and Performance Metrics

We simulate a mmWave massive MIMO system in which $U = 32$ single-antenna UEs transmit to a $B = 256$ antenna BS under line-of-sight (LoS) conditions. The channel matrices are generated using the QuaDRiGa mmMAGIC UMi model [18] for a 60 GHz carrier and a uniform linear array (ULA) with $\lambda/2$ spacing. We let the U UEs and the jammer be randomly placed at distances from 10 m to 100 m within a 120° angular sector in front of the BS. The minimum angular separation between two UEs, as well as between the jammer and any UE, is 1° . We assume ± 3 dB per-UE power control, so that the ratio between maximum and minimum per-UE-receive power is 4. The transmit constellation is 16-QAM. In our simulations, we define the average receive signal-to-noise ratio (SNR) as

$$\text{SNR} \triangleq \frac{\mathbb{E}_s[\|\mathbf{H}\mathbf{s}\|_2^2]}{\mathbb{E}_n[\|\mathbf{n}\|_2^2]}. \quad (29)$$



(a) BER vs. SNR



(b) Fraction of successfully served UEs vs. SNR

Fig. 3. Comparison between soft-nulling in antenna domain (ANT) and SNIPS with different cluster sizes S , in terms of (a) uncoded bit error-rate (BER) and (b) fraction of successfully served UEs. The relative jammer power and ADC resolution are $\rho = 25$ dB and $q = 4$ bits per real dimension, respectively.

To quantify the jammer's power in comparison to a single UE, we define the relative jammer power ρ as

$$\rho \triangleq \frac{U \mathbb{E}_w [\|\mathbf{j}w\|_2^2]}{\mathbb{E}_s [\|\mathbf{H}_s\|_2^2]} = \frac{U E_w \|\mathbf{j}\|_2^2}{E_s \|\mathbf{H}\|_F^2}. \quad (30)$$

We will consider two performance metrics: Uncoded bit error-rate (BER) and a metric introduced in [19] called the per-UE root mean-square symbol error (RMSSE). The RMSSE for the u th UE over n data symbol slots is defined as:

$$RMSSE_u \triangleq \sqrt{\frac{\sum_{k=1}^n |s_{u,k}^* - \hat{s}_{u,k}|^2}{\sum_{k=1}^n |s_{u,k}|^2}}. \quad (31)$$

Here, $s_{u,k}$ and $s_{u,k}^*$ are the transmitted and estimated data symbol of the u th UE at time slot k , respectively. To understand the relevance of $RMSSE_u$ as a performance metric, it is helpful to compare it to the error vector magnitude (EVM) requirements in the 3GPP 5G NR technical specification [20]. The EVM is loosely speaking the square root of the sum of $RMSSE_u$ -squared over all U UEs. Vice versa, the $RMSSE_u$ is loosely speaking a single-UE proxy for EVM. We will therefore interpret $RMSSE_u$ as a random variable and analyze its distribution by means of Monte-Carlo simulations. For 16-QAM transmission, the 3GPP 5G NR technical specification requires an EVM below 12.5% [20, Tbl. 6.5.2.2-1]. We therefore consider as a performance metric the fraction of UEs (averaged over UE placements/channel realizations, noise and jammer realizations, and data transmissions) for which the $RMSSE_u$ is below 12.5%.

B. The Efficacy of Beam-Slicing

Figure 3 evaluates the performance of SNIPS for different antenna cluster sizes S . We compare the baseline, which performs soft-nulling in antenna domain, against SNIPS with cluster sizes $S = \{2, 4, 8, 16, 32, 64, 256\}$, hence considering analog beam-slicing that only operates on a pair of adjacent antennas up to a single cluster consisting of the whole antenna array. We note that with a cluster size $S = B = 256$, beam-slicing corresponds to performing a beamspace transform, which is difficult to implement in the analog domain.

The evaluation is done for a strong relative jammer power $\rho = 25$ dB and a BS with ADCs with a resolution of $q = 4$ bits (per real dimension).

In Figure 3(a), we see that beam-slicing with two-antenna clusters ($S = 2$) already gives noticeable BER improvements over the antenna-domain baseline. Figure 3(a) also shows that large clusters are superior to small ones. However, the performance of a full beamspace transform ($S = 256$) is inferior to SNIPS with $S = 64$, which exhibits the best performance. The reason for this has nothing to do with how the beamspace transform distributes the jammer interference to the ADCs. Instead, the performance decrease can be explained with the fact that, after a full beamspace transform, the UE signals are concentrated to only a few ADCs, so that low-resolution ($q = 4$) ADCs can no longer represent them appropriately. This observation suggests that the fully-centralized beamspace transform, which in any case is impractical, may not necessarily be optimal for achieving the full potential of SNIPS. For this reason, we do not consider it in the following experiments.

The behavior of the fraction of UEs whose $RMSSE_u$ is below 12.5% at a given SNR is shown in Figure 3(b). In terms of this criterion, the antenna-domain baseline is unable to serve a significant fraction of UEs regardless of SNR. In contrast, beam-slicing with two-antenna clusters ($S = 2$) serves almost 20% of UEs at high SNR. This fraction increases with the cluster size S , with SNIPS being able to serve more than 80% of UEs at high SNR for $S = 32$ and $S = 64$.

C. When is Beam-Slicing Needed?

The experiments in Section IV-B indicate that, for a strong jammer, SNIPS outperforms the antenna-domain baseline and that the best performance is achieved with a large cluster size S . However, analog transforms spanning a large number of antennas are difficult to implement in practice [15], with the fully-centralized beamspace transform being likely infeasible for massive MU-MIMO. We therefore consider a moderately-sized $S = 8$ beam-slicer for our subsequent evaluations.

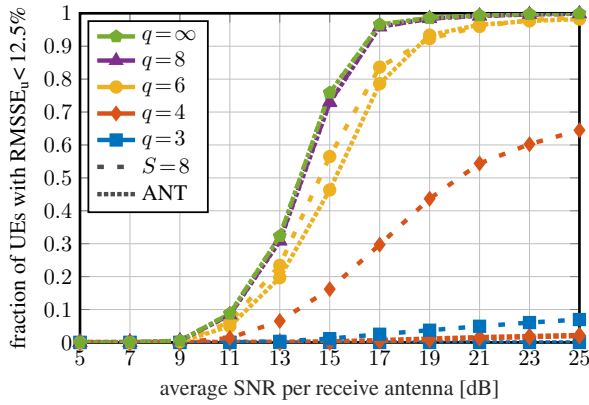


Fig. 4. Fraction of successfully served UEs vs. SNR, compared between antenna domain (ANT; dotted curves) and SNIPS with $S = 8$ (dashed curves) for different ADC resolutions q [bits per real dimension]. The relative jammer power is $\rho = 25$ dB.




Fig. 5. Fraction of successfully served UEs vs. SNR, compared between antenna domain (ANT; dotted curves) and SNIPS with $S = 8$ (dashed curves) for different relative jammer powers ρ [dB]. The ADC resolution is $q = 4$ bits per real dimension.

In Figure 4, we analyze the impact of ADC resolution for a relative jammer power $\rho = 25$ dB. We consider the fraction of UEs successfully served in terms of the criterion $RMSSE_u < 12.5\%$. For infinite- or high-resolution ($q = 8$) ADCs, SNIPS and the antenna-domain baseline have identical performance. However, already for 6-bit ADCs, SNIPS outperforms the baseline. For low-resolution ADCs with $q \leq 4$, the antenna-domain baseline is unable to serve a significant fraction of UEs (less than 2%). In contrast, SNIPS can at least serve some UEs at high SNR even for $q = 3$, and it can serve up to 65% of UEs for $q = 4$.

In Figure 5, we consider the impact of the relative jammer power ρ for 4-bit ADCs. We again compare SNIPS with clusters of size $S = 8$ against the antenna-domain baseline. We see that their performance is virtually identical for weak jammers that are as strong as the average UE ($\rho = 0$ dB). However, for a jammer with $\rho = 15$ dB, SNIPS already significantly outperforms the antenna-domain baseline in terms of successfully served UEs, and this gap continues to widen as jamming power

increases further.

These experiments confirm that strong jammers pose a serious problem for classical all-digital jamming suppression methods when combined with low-resolution ADCs. Our results also demonstrate that SNIPS is able to successfully mitigate jammer interference in a practical manner.

REFERENCES

- [1] T. S. Rappaport, R. W. Heath Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Prentice Hall, 2015.
- [2] E. G. Larsson, F. Tufvesson, O. Edfors, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [3] O. Castañeda, S. Jacobsson, G. Durisi, M. Coldrey, T. Goldstein, and C. Studer, "1-bit massive MU-MIMO precoding in VLSI," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 7, no. 4, pp. 508–522, Dec. 2017.
- [4] Q. Yan, H. Zeng, T. Jiang, M. Li, W. Lou, and Y. T. Hou, "MIMO-based jamming resilient communication in wireless networks," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2014, pp. 2697–2705.
- [5] W. Shen, P. Ning, X. He, H. Dai, and Y. Liu, "MCR decoding: A MIMO approach for defending against wireless jamming attacks," in *Proc. IEEE Conf. Commun. Netw. Security (CNS)*, Oct. 2014, pp. 133–138.
- [6] D. Kapetanović, G. Zheng, K. Wong, and B. Ottersten, "Detection of pilot contamination attack using random training and massive MIMO," in *Proc. IEEE Int. Symp. Personal, Indoor, Mobile Radio Commun. (PIMRC)*, Sep. 2013, pp. 13–18.
- [7] H. Akhlaghpasand, S. M. Razavizadeh, E. Björnsson, and T. T. Do, "Jamming detection in massive MIMO systems," *IEEE Wireless Commun. Lett.*, vol. 7, no. 2, pp. 242–245, Apr. 2018.
- [8] J. Vinogradova, E. Björnsson, and E. G. Larsson, "Detection and mitigation of jamming attacks in massive MIMO systems using random matrix theory," in *Proc. IEEE Int. Workshop Signal Process. Advances Wireless Commun. (SPAWC)*, Jul. 2016.
- [9] N. Zhao, J. Guo, F. R. Yu, and V. C. M. Leung, "Antijamming schemes for interference-alignment-based wireless networks," *IEEE Trans. Veh. Technol.*, vol. 66, no. 2, pp. 1271–1283, Feb. 2017.
- [10] T. T. Do, E. Björnsson, E. G. Larsson, and S. M. Razavizadeh, "Jamming-resistant receivers for the massive MIMO uplink," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 1, pp. 210–223, Jan. 2018.
- [11] H. Akhlaghpasand, E. Björnsson, and S. M. Razavizadeh, "Jamming suppression in massive MIMO systems," *IEEE Trans. Circuits Syst. II*, vol. 68, no. 1, pp. 182–186, Jan. 2020.
- [12] —, "Jamming-robust uplink transmission for spatially correlated massive MIMO systems," *IEEE Trans. Commun.*, vol. 68, no. 6, pp. 3495–3504, Mar. 2020.
- [13] S. Bagherinejad and S. M. Razavizadeh, "Direction-based jamming detection and suppression in mmwave massive MIMO networks," *IET Commun.*, Mar. 2021.
- [14] J. Brady, N. Behdad, and A. M. Sayeed, "Beamspace MIMO for millimeter-wave communications: System architecture, modeling, analysis, and measurements," *IEEE Trans. Antennas Propag.*, vol. 61, no. 7, pp. 3814–3827, Jul. 2013.
- [15] S. Yoshi, "High-fidelity spatial signal processing in low-power mixed-signal VLSI arrays," Ph.D. dissertation, University of California, San Diego (UCSD), San Diego, CA, 2017.
- [16] J. Max, "Quantizing for minimum distortion," *IRE Trans. Inf. Theory*, vol. 6, no. 1, pp. 7–12, Mar. 1960.
- [17] J. J. Bussgang, "Crosscorrelation functions of amplitude-distorted Gaussian signals," Res. Lab. Elec., Cambridge, MA, USA, Tech. Rep. 216, Mar. 1952.
- [18] S. Jaeckel, L. Raschkowski, K. Börner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3242–3256, Jun. 2014.
- [19] H. Song, T. Goldstein, X. You, C. Zhang, O. Tirkkonen, and C. Studer, "Joint channel estimation and data detection in cell-free massive MU-MIMO systems," *submitted*, 2021.
- [20] 3GPP, "5G; NR; base station (BS) radio transmission and reception," Mar. 2021, TS 38.104 version 17.1.0 Rel. 17.