

Mitigating Smart Jammers in MU-MIMO via Joint Channel Estimation and Data Detection

Gian Marti and Christoph Studer

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

e-mail: gimarti@ethz.ch and studer@ethz.ch

Abstract—Wireless systems must be resilient to jamming attacks. Existing mitigation methods require knowledge of the jammer's transmit characteristics. However, this knowledge may be difficult to acquire, especially for smart jammers that attack only specific instants during transmission in order to evade mitigation. We propose a novel method that mitigates attacks by smart jammers on massive multi-user multiple-input multiple-output (MU-MIMO) basestations (BSs). Our approach builds on recent progress in joint channel estimation and data detection (JED) and exploits the fact that a jammer cannot change its subspace within a coherence interval. Our method, called MAED (short for MitigAtion, Estimation, and Detection), uses a novel problem formulation that combines jammer estimation and mitigation, channel estimation, and data detection, instead of separating these tasks. We solve the problem approximately with an efficient iterative algorithm. Our simulation results show that MAED effectively mitigates a wide range of smart jamming attacks without having any *a priori* knowledge about the attack type.

I. INTRODUCTION

Jamming attacks pose a serious threat to the continuous operability of wireless communication systems [1]. Effective methods to mitigate such attacks are necessary as wireless systems become increasingly critical to modern infrastructure [2]. In the uplink of massive multi-user multiple-input multiple-output (MU-MIMO) systems, effective jammer mitigation is rendered possible by the asymmetry in the number of antennas between the basestation (BS), which has many antennas, and a mobile jamming device, which has one or few antennas. One possibility, for instance, is to project the receive signals on the subspace orthogonal to the jammer's channel [3], [4]. But such methods require accurate knowledge of the jammer's channel. If a jammer transmits permanently and with a static signature (often called barrage jamming), the BS can estimate the required quantities, for instance during a dedicated period in which the user equipments (UEs) do not transmit [3] or in which they transmit predefined symbols [4]. Instead of barrage jamming, however, a smart jammer might jam the system only at specific time instants. Such attacks may prevent the BS from estimating a jammer's channel with simple methods.

A. State of the Art

Multi-antenna wireless systems have the unique potential to effectively mitigate jamming attacks, and a variety of multi-


antenna methods have been proposed for the mitigation of jamming attacks in MIMO systems [3]–[11]. Common to all of them is the assumption—in one way or other—that information about the jammer's transmit characteristics (e.g., the jammer's channel, or the covariance matrix between the UE transmit signals and the jammed receive signals) can be estimated based on some specific subset of the receive samples.¹ Fig. 1(a) illustrates the approach taken by these methods: The data phase is preceded by an augmented training phase in which the jammer's transmit characteristics are estimated in addition to the channel matrix. This augmented training phase can either complement a traditional pilot phase with a period during which the UEs do not transmit in order to enable jammer estimation (e.g., [3], [5]), or it can consist of an extended pilot phase so that there exist pilot sequences which are unused by the UEs and on whose span the receive signals can be projected to estimate the jammer's subspace (e.g., [8]–[10]). The estimated jammer characteristics are then used to perform jammer-mitigating data detection. Such an approach succeeds in the case of barrage jammers, but is unreliable for estimating the transmit characteristics of smart jammers: A smart jammer can evade estimation and thus circumvent mitigation by not transmitting in those samples, for instance because it is aware of the defense mechanism, or simply because it jams in brief bursts only. For this reason, our proposed method MAED unifies jammer estimation, channel estimation and data detection, see Fig. 1(b).

Many studies have already shown how smart jammers can disrupt wireless communication systems by targeting only specific parts of the wireless transmission [12]–[20] instead of using barrage jamming. Jammers that jam only the pilot phase have received considerable attention [12]–[16], as such attacks can be more energy-efficient than barrage jamming in disrupting communication systems that do not defend themselves against jammers [14]–[16]. However, if a jammer is active during the pilot phase, then a BS that *does* defend itself against jamming attacks can estimate the jammer's channel by exploiting knowledge of the UE transmit symbols, for instance with the aid of unused pilot sequences [8]–[10]. To disable such jammer-mitigating communication systems, a smart jammer might therefore refrain from jamming the pilot phase and only target the data phase, even if such jamming attacks have not received


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

The authors thank O. Castañeda and S. Taner for comments and suggestions.

¹The method of [7] is to some extent an exception as it estimates the UEs' subspace and projects the receive signals thereon. However, this method distinguishes the UEs' from the jammer's subspace based on the receive power, thereby presuming that the UEs and the jammer transmit with different power.



(a) Existing jammer-mitigation methods separate jammer estimation (JEST) and channel estimation (CHEST) from the jammer-resilient data detection (DET). Such methods are ineffective against jammers that attack only the data phase and do not transmit in the training phase.



(b) MAED unifies jammer estimation and mitigation, channel estimation, and data detection to mitigate jamming attacks regardless of their activity pattern.

Fig. 1. The approach to jammer mitigation taken by existing methods (a) compared to the approach taken by MAED (b). In the figure, $\{y_1, \dots, y_K\}$ are the receive signals, and $\hat{\mathbf{j}}$, $\hat{\mathbf{H}}$ and $\hat{\mathbf{S}}_D$ are the estimates of the jammer channel, the UE channel matrix, and the UE transmit symbols, respectively.

much attention so far [18], [19]. Other threat models that have been analyzed include jammers that attack specific control channels [16]–[18] or the time synchronization phase [20].

B. Contributions

We propose MAED (short for MitigAtion, Estimation, and Detection), a novel method for jammer mitigation that does not depend on the jammer being active during *any* specific period. Instead, our approach leverages the fact that a jammer cannot change its subspace instantaneously. To this end, MAED utilizes a problem formulation which unifies jammer estimation and cancellation, channel estimation, and data detection, instead of dealing with these tasks independently, see Fig. 1(b). In doing this, the method builds upon techniques for joint channel estimation and data detection (JED) [21]–[27]. The proposed problem formulation is approximately solved using an efficient iterative algorithm. Extensive simulation results demonstrate that MAED is able to effectively mitigate a wide variety of naïve and smart jamming attacks, without requiring any *a priori* knowledge about the attack type.

C. Notation

Matrices and column vectors are represented by boldface uppercase and lowercase letters, respectively. For a matrix $\mathbf{A}$, the transpose is $\mathbf{A}^T$, the conjugate transpose is $\mathbf{A}^H$, the entry in the ℓ th row and k th column is $[\mathbf{A}]_{\ell,k}$, the submatrix consisting of the columns from n through m is $\mathbf{A}_{[n:m]}$, and the Frobenius norm is $\|\mathbf{A}\|_F$. The $N \times N$ identity matrix is $\mathbf{I}_N$. For a vector $\mathbf{a}$, the ℓ_2 -norm is $\|\mathbf{a}\|_2$, the real part is $\Re\{\mathbf{a}\}$, and the imaginary part is $\Im\{\mathbf{a}\}$. The expectation with respect to the random vector $\mathbf{x}$ is denoted by $\mathbb{E}_{\mathbf{x}}[\cdot]$. We define $i^2 = -1$.

The complex n -hypersphere of radius r is denoted by $\mathbb{S}_r^n$, and the span of a vector $\mathbf{a}$ is denoted by $\text{span}(\mathbf{a})$.

II. SYSTEM SETUP

We consider the uplink of a massive MU-MIMO system in which U single-antenna UEs transmit data to a B antenna BS in the presence of a single-antenna jammer. The channels are assumed to be frequency flat and block-fading with coherence time $K = T + D$. The first T time slots are used to transmit pilot symbols; the remaining D time slots are used to transmit payload data. The UE transmit matrix is $\mathbf{S} = [\mathbf{S}_T, \mathbf{S}_D]$, where $\mathbf{S}_T \in \mathbb{C}^{U \times T}$ and $\mathbf{S}_D \in \mathbb{C}^{U \times D}$ contain the pilots and the transmit symbols, respectively, and $\mathbf{S}$ is the transmit constellation, which is taken to be QPSK with symbol energy E_s throughout this paper.² We assume that the jammer does not prevent the UEs and the BS from establishing time synchronization, which allows us to use the discrete-time input-output relation

$$\mathbf{Y} = \mathbf{H}\mathbf{S} + \mathbf{j}\mathbf{w}^T + \mathbf{N}. \quad (1)$$

Here, $\mathbf{Y} \in \mathbb{C}^{B \times K}$ is the BS-side receive matrix that contains the B -dimensional receive vectors over all K time slots, $\mathbf{H} \in \mathbb{C}^{B \times U}$ models the MIMO uplink channel, $\mathbf{j} \in \mathbb{C}^B$ models the channel between the jammer and the BS, $\mathbf{w} \in \mathbb{C}^K$ contains the jammer's transmit symbols over all K time slots, and $\mathbf{N} \in \mathbb{C}^{B \times K}$ models thermal noise and contains independently and identically distributed (i.i.d.) circularly-symmetric complex Gaussian entries with variance N_0 . In what follows, we usually assume that the jammer's transmit symbols $\mathbf{w}$ are independent of the UE transmit symbols $\mathbf{S}$. However, we will also consider a scenario in which we drop this assumption (see Section IV-F). We make no other assumptions about the distribution of $\mathbf{w}$. In particular, we do not assume that its entries are distributed i.i.d.

III. JOINT JAMMER ESTIMATION AND MITIGATION, CHANNEL ESTIMATION, AND DATA DETECTION

Our goal is threefold: (i) mitigating the jammer's interference by locating its one-dimensional subspace $\text{span}(\mathbf{j})$ and projecting the receive matrix $\mathbf{Y}$ onto the orthogonal complement of that subspace, (ii) estimating the channel matrix $\mathbf{H}$, and (iii) recovering the UE data in $\mathbf{S}_D$. To this end, we develop a novel problem formulation that combines all three of these aspects. We then propose an iterative algorithm to approximately solve this optimization problem.

A. The MAED Optimization Problem

In the absence of a jammer, the maximum-likelihood problem for joint channel estimation and data detection is [21]

$$\{\hat{\mathbf{H}}, \hat{\mathbf{S}}_D\} = \arg \min_{\substack{\hat{\mathbf{H}} \in \mathbb{C}^{B \times U} \\ \hat{\mathbf{S}}_D \in \mathbb{S}^{U \times D}}} \|\mathbf{Y} - \hat{\mathbf{H}}\hat{\mathbf{S}}_D\|_F^2, \quad (2)$$

where, for brevity, we define $\tilde{\mathbf{S}} \triangleq [\mathbf{S}_T, \tilde{\mathbf{S}}_D]$ and leave the dependence on $\tilde{\mathbf{S}}_D$ implicit. This objective already integrates

²Our method can be extended to higher-order modulation schemes.

the goals of estimating the channel matrix and detecting the transmit symbols. The jammer, however, will lead to a residual

$$\|\mathbf{Y} - \mathbf{H}\mathbf{S}\|_F^2 = \|\mathbf{j}\mathbf{w}^T + \mathbf{N}\|_F^2 \approx \|\mathbf{j}\mathbf{w}^T\|_F^2 \quad (3)$$

when plugging the true channel and data matrices into (2), and assuming that the contribution of the noise $\mathbf{N}$ is negligible.

Consider now the projection onto the orthogonal subspace (POS) for jammer mitigation [3]: POS nulls a jammer by orthogonally projecting the receive signals onto the orthogonal complement of $\text{span}(\mathbf{j})$ using the matrix $\mathbf{P}(\mathbf{j}) \triangleq \mathbf{I}_B - \mathbf{j}\mathbf{j}^H / \|\mathbf{j}\|^2$:

$$\mathbf{P}(\mathbf{j})\mathbf{Y} = \mathbf{P}(\mathbf{j})\mathbf{H}\mathbf{S} + \mathbf{P}(\mathbf{j})\mathbf{j}\mathbf{w}^T + \mathbf{P}(\mathbf{j})\mathbf{N} \quad (4)$$

$$= \mathbf{P}(\mathbf{j})\mathbf{H}\mathbf{S} + \mathbf{P}(\mathbf{j})\mathbf{N}. \quad (5)$$

One can then define $\mathbf{Y}_{\mathbf{P}(\mathbf{j})} \triangleq \mathbf{P}(\mathbf{j})\mathbf{Y}$, $\mathbf{H}_{\mathbf{P}(\mathbf{j})} \triangleq \mathbf{P}(\mathbf{j})\mathbf{H}$, and perform channel estimation and data detection using the resulting jammer-free system. The difficulty is, of course, that the projection matrix $\mathbf{P}(\mathbf{j})$ depends on the (unknown) direction $\mathbf{j}/\|\mathbf{j}\|$ of the jammer's channel.

Now, consider what happens when we take the matrix³ $\tilde{\mathbf{P}} \triangleq \mathbf{I} - \tilde{\mathbf{p}}\tilde{\mathbf{p}}^H$, $\tilde{\mathbf{p}} \in \mathbb{S}_1^B$ which orthogonally projects a signal onto the orthogonal complement of some arbitrary one-dimensional subspace $\text{span}(\tilde{\mathbf{p}})$, and then apply that projection to the objective of (2) as follows:

$$\|\tilde{\mathbf{P}}(\mathbf{Y} - \tilde{\mathbf{H}}\tilde{\mathbf{S}})\|_F^2. \quad (6)$$

If we now plug the true channel and data matrices into (6) and assume that the noise $\mathbf{N}$ is negligible, then we obtain

$$\|\tilde{\mathbf{P}}(\mathbf{Y} - \mathbf{H}\mathbf{S})\|_F^2 = \|\tilde{\mathbf{P}}\mathbf{j}\mathbf{w}^T + \tilde{\mathbf{P}}\mathbf{N}\|_F^2 \quad (7)$$

$$\approx \|\tilde{\mathbf{P}}\mathbf{j}\mathbf{w}^T\|_F^2 \quad (8)$$

$$\geq 0, \quad (9)$$

with equality if and only if $\tilde{\mathbf{p}}$ is collinear with $\mathbf{j}$. In other words, the unit vector $\tilde{\mathbf{p}}$ which—in combination with the true channel and data matrices—minimizes (6) is collinear with the jammer's channel, and hence $\tilde{\mathbf{P}}$ is the POS matrix.

So, assuming the noise $\mathbf{N}$ is negligible, then $\{\tilde{\mathbf{p}}, \tilde{\mathbf{H}}, \tilde{\mathbf{S}}\}$ minimizes (6) if (i) $\tilde{\mathbf{P}}$ is the orthogonal projection onto the orthogonal complement of $\text{span}(\mathbf{j})$, (ii) $\tilde{\mathbf{H}}$ is the true channel matrix, and (iii) $\tilde{\mathbf{S}}$ contains the true data matrix. These are, of course, exactly the goals that we wanted to attain. Following this insight, we now formulate the MAED joint jammer estimation and mitigation, channel estimation, and data detection problem:

$$\{\hat{\mathbf{p}}, \hat{\mathbf{H}}_{\mathbf{P}}, \hat{\mathbf{S}}_D\} = \arg \min_{\substack{\tilde{\mathbf{p}} \in \mathbb{S}_1^B \\ \hat{\mathbf{H}}_{\mathbf{P}} \in \mathbb{C}^{B \times U} \\ \hat{\mathbf{S}}_D \in \mathbb{S}^{U \times D}} \|\tilde{\mathbf{P}}\mathbf{Y} - \tilde{\mathbf{H}}_{\mathbf{P}}\tilde{\mathbf{S}}\|_F^2. \quad (10)$$

Note that, compared to (6), we have absorbed the projection matrix $\tilde{\mathbf{P}}$ directly into the unknown channel matrix $\tilde{\mathbf{H}}_{\mathbf{P}}$. Otherwise the columns of $\tilde{\mathbf{H}}_{\mathbf{P}}$ would be ill-defined with respect to the length of their components in the direction of $\tilde{\mathbf{p}} \approx \mathbf{j}$, meaning that one could not distinguish between channel estimates $\tilde{\mathbf{H}} + \alpha\mathbf{j}\mathbf{w}^T$ with different $\alpha, \tilde{\mathbf{w}}$.

³The dependence of $\tilde{\mathbf{P}}$ on $\tilde{\mathbf{p}}$ is left implicit here and throughout the paper.

B. Solving the MAED Optimization Problem

The objective (10) is quadratic in $\tilde{\mathbf{H}}_{\mathbf{P}}$, so we can derive the optimal value of $\tilde{\mathbf{H}}_{\mathbf{P}}$ as a function of $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{S}}$, as

$$\hat{\mathbf{H}}_{\mathbf{P}} = \tilde{\mathbf{P}}\mathbf{Y}\tilde{\mathbf{S}}^\dagger, \quad (11)$$

where $\tilde{\mathbf{S}}^\dagger = \tilde{\mathbf{S}}^H(\tilde{\mathbf{S}}\tilde{\mathbf{S}}^H)^{-1}$ is the Moore-Penrose pseudo-inverse of $\tilde{\mathbf{S}}$. Substituting $\hat{\mathbf{H}}_{\mathbf{P}}$ back into (10) yields

$$\{\hat{\mathbf{p}}, \hat{\mathbf{S}}_D\} = \arg \min_{\substack{\tilde{\mathbf{p}} \in \mathbb{S}_1^B \\ \hat{\mathbf{S}}_D \in \mathbb{S}^{D \times U}}} \|\tilde{\mathbf{P}}\mathbf{Y}(\mathbf{I}_K - \tilde{\mathbf{S}}^\dagger\tilde{\mathbf{S}})\|_F^2. \quad (12)$$

Solving (12) is difficult due to its combinatorial nature, so we resort to solving it approximately. First, we relax the constraint set $\mathcal{S}$ to its convex hull $\mathcal{C} \triangleq \text{conv}(\mathcal{S})$ as in [24]. We then solve this relaxed problem formulation approximately by alternately performing a forward-backward splitting step in $\tilde{\mathbf{S}}$ and a minimization step in $\tilde{\mathbf{P}}$.

C. Forward-Backward Splitting Step in $\tilde{\mathbf{S}}$

Forward-backward splitting (FBS) [28], also called proximal gradient descent, is an iterative method that solves convex optimization problems of the form

$$\arg \min_{\tilde{\mathbf{s}}} f(\tilde{\mathbf{s}}) + g(\tilde{\mathbf{s}}), \quad (13)$$

where f is convex and differentiable, and g is convex but not necessarily differentiable, smooth, or bounded. Starting from an initialization vector $\tilde{\mathbf{s}}^{(0)}$, FBS solves the problem in (13) iteratively by computing

$$\tilde{\mathbf{s}}^{(t+1)} = \text{prox}_g(\tilde{\mathbf{s}}^{(t)} - \tau^{(t)}\nabla f(\tilde{\mathbf{s}}^{(t)}); \tau^{(t)}). \quad (14)$$

Here, $\nabla f(\tilde{\mathbf{s}})$ is the gradient of $f(\tilde{\mathbf{s}})$, $\tau^{(t)}$ is the stepsize at iteration t , and prox_g is the proximal operator of g [29]. For a suitably chosen sequence of stepsizes $\{\tau^{(t)}\}$, FBS solves convex optimization problems exactly (provided that the number of iterations is sufficiently large). FBS can also be utilized to approximately and efficiently solve non-convex problems, even though there are typically no guarantees for optimality or even convergence [28].

For the optimization problem in (12), we define f and g as

$$f(\tilde{\mathbf{S}}) = \|\tilde{\mathbf{P}}\mathbf{Y}(\mathbf{I}_K - \tilde{\mathbf{S}}^\dagger\tilde{\mathbf{S}})\|_F^2, \quad (15)$$

and

$$g(\tilde{\mathbf{S}}) = \begin{cases} 0 & \text{if } \tilde{\mathbf{S}}_{[1:T]} = \mathbf{S}_T \text{ and } \tilde{\mathbf{S}}_{[T+1:K]} \in \mathcal{C}^{U \times D} \\ \infty & \text{else.} \end{cases} \quad (16)$$

The gradient of f in $\tilde{\mathbf{S}}$ is given by

$$\nabla f(\tilde{\mathbf{S}}) = (\tilde{\mathbf{S}}^\dagger)^H \mathbf{Y}^H \tilde{\mathbf{P}}\mathbf{Y}(\mathbf{I}_K - \tilde{\mathbf{S}}^\dagger\tilde{\mathbf{S}}) \quad (17)$$

and the proximal operator for g is simply the orthogonal projection onto the constraint set, which is

$$[\text{prox}_g(\tilde{\mathbf{S}})]_{u,k} = \begin{cases} [\mathbf{S}_T]_{u,k} & \text{if } k \in [1 : T] \\ \text{proj}_{\mathcal{C}}([\tilde{\mathbf{S}}_{u,k}]) & \text{else,} \end{cases} \quad (18)$$

where (for QPSK) $\text{proj}_{\mathcal{C}}$ acts entry-wise on $[\tilde{\mathbf{S}}]_{u,k}$ as

$$\text{proj}_{\mathcal{C}}(x) = \min\{\max\{\Re(x), -\sqrt{E_s/2}\}, \sqrt{E_s/2}\} + i \min\{\max\{\Im(x), -\sqrt{E_s/2}\}, \sqrt{E_s/2}\}. \quad (19)$$

For the selection of the stepsizes $\{\tau^{(t)}\}$, we use the Barzilai-Borwein method [30] as detailed in [28], [31].

D. Minimization Step in $\tilde{\mathbf{P}}$

After the FBS step in $\tilde{\mathbf{S}}$, we minimize (12) with respect to the vector $\tilde{\mathbf{p}}$. Defining $\tilde{\mathbf{E}} \triangleq \mathbf{Y}(\mathbf{I}_K - \tilde{\mathbf{S}}^\dagger \tilde{\mathbf{S}})$ and performing standard algebraic manipulations yields

$$\hat{\mathbf{p}} = \arg \min_{\tilde{\mathbf{p}} \in \mathbb{S}_1^B} \|\tilde{\mathbf{P}} \tilde{\mathbf{E}}\|_F^2 \quad (20)$$

$$= \arg \max_{\tilde{\mathbf{p}} \in \mathbb{S}_1^B} \tilde{\mathbf{p}}^H \tilde{\mathbf{E}} \tilde{\mathbf{E}}^H \tilde{\mathbf{p}}. \quad (21)$$

This implies that the minimizer $\hat{\mathbf{p}}$ is the unit vector which maximizes the Rayleigh quotient of $\tilde{\mathbf{E}} \tilde{\mathbf{E}}^H$. The solution to this problem is the eigenvector $\mathbf{v}_1(\tilde{\mathbf{E}} \tilde{\mathbf{E}}^H)$ belonging to the largest eigenvalue of $\tilde{\mathbf{E}} \tilde{\mathbf{E}}^H$, normalized to unit length [32, Thm. 4.2.2],

$$\hat{\mathbf{p}} = \frac{\mathbf{v}_1(\tilde{\mathbf{E}} \tilde{\mathbf{E}}^H)}{\|\mathbf{v}_1(\tilde{\mathbf{E}} \tilde{\mathbf{E}}^H)\|_2}. \quad (22)$$

Calculating this eigenvector for every iteration is computationally expensive, so we only do it for the very first iteration. In all subsequent iterations, we then approximate its value with a single power method step [33, Sec. 8.2.1], i.e., we estimate

$$\hat{\mathbf{p}}^{(t+1)} = \frac{\tilde{\mathbf{E}}^{(t+1)}(\tilde{\mathbf{E}}^{(t+1)})^H \hat{\mathbf{p}}^{(t)}}{\|\tilde{\mathbf{E}}^{(t+1)}(\tilde{\mathbf{E}}^{(t+1)})^H \hat{\mathbf{p}}^{(t)}\|_2^2}, \quad (23)$$

where we initialize the power method with the subspace estimate $\hat{\mathbf{p}}^{(t)}$ from the previous iteration.

E. The MAED Algorithm

We now have all the building blocks for MAED, which is summarized in Algorithm 1. Its only input is the receive matrix $\mathbf{Y}$. MAED is initialized with $\tilde{\mathbf{S}}_D^{(0)} = \mathbf{0}_{U \times D}$, $\tilde{\mathbf{P}}^{(0)} = \mathbf{I}_B$, and $\tau^{(0)} = \tau_0 = 0.1$, and runs for a fixed number of $t_{\max}$ iterations.

IV. SIMULATION RESULTS

A. Simulation Setup

We simulate a massive MU-MIMO system as described in Section IV-A with $B = 128$ BS antennas, $U = 32$ single-antenna UEs, and with one single-antenna jammer. The UEs transmit for $K = 96$ time slots. The first $T = 32$ time slots are used to transmit orthogonal pilots $\mathbf{S}_T$ in the form of a 32×32 Hadamard matrix (scaled to symbol energy E_s). The remaining $D = 64$ time slots are used to transmit QPSK payload data (also with symbol energy E_s). The channels of the UEs and the jammer are modelled as i.i.d. Rayleigh fading. We define the average receive signal-to-noise ratio (SNR) as follows:

$$\text{SNR} \triangleq \frac{\mathbb{E}_{\mathbf{S}}[\|\mathbf{H}\mathbf{S}\|_F^2]}{\mathbb{E}_{\mathbf{N}}[\|\mathbf{N}\|_F^2]}. \quad (24)$$

Algorithm 1 MAED

```

1: input:  $\mathbf{Y}$ 
2: initialize:  $\tilde{\mathbf{S}}^{(0)} = [\mathbf{S}_T, \mathbf{0}_{U \times D}]$ ,  $\tilde{\mathbf{P}}^{(0)} = \mathbf{I}_B$ ,  $\tau^{(0)} = \tau_0$ 
3: for  $t = 0$  to  $t_{\max} - 1$  do
4:    $\nabla f(\tilde{\mathbf{S}}^{(t)}) = (\tilde{\mathbf{S}}^{(t)\dagger})^H \mathbf{Y}^H \tilde{\mathbf{P}}^{(t)} \mathbf{Y} (\mathbf{I}_K - \tilde{\mathbf{S}}^{(t)} \tilde{\mathbf{S}}^{(t)\dagger})$ 
5:    $\tilde{\mathbf{S}}^{(t+1)} = \text{prox}_g(\tilde{\mathbf{S}}^{(t)} + \tau^{(t)} \nabla f(\tilde{\mathbf{S}}^{(t)}))$  (cf. (18))
6:    $\tilde{\mathbf{E}}^{(t+1)} = \mathbf{Y}(\mathbf{I}_K - \tilde{\mathbf{S}}^{(t+1)\dagger} \tilde{\mathbf{S}}^{(t+1)})$ 
7:   if  $t = 0$  then
8:      $\hat{\mathbf{p}}^{(t+1)} = \mathbf{v}_1(\tilde{\mathbf{E}}^{(t+1)} \tilde{\mathbf{E}}^{(t+1)H}) / \|\mathbf{v}_1(\tilde{\mathbf{E}}^{(t+1)} \tilde{\mathbf{E}}^{(t+1)H})\|_2$ 
9:   else
10:     $\hat{\mathbf{p}}^{(t+1)} = \tilde{\mathbf{E}}^{(t+1)} \tilde{\mathbf{E}}^{(t+1)H} \hat{\mathbf{p}}^{(t)} / \|\tilde{\mathbf{E}}^{(t+1)} \tilde{\mathbf{E}}^{(t+1)H} \hat{\mathbf{p}}^{(t)}\|_2$ 
11:   end if
12:    $\tilde{\mathbf{P}}^{(t+1)} = \mathbf{I}_B - \hat{\mathbf{p}}^{(t+1)} \hat{\mathbf{p}}^{(t+1)H}$ 
13:    $\tau^{(t+1)} = \text{Barzilai-Borwein}(\tau^{(t)}, \tilde{\mathbf{S}}^{(t)}, \tilde{\mathbf{S}}^{(t+1)}, \dots, \nabla f(\tilde{\mathbf{S}}^{(t)}), \nabla f(\tilde{\mathbf{S}}^{(t+1)}))$  [28]
14: end for
15: output:  $\tilde{\mathbf{S}}_{[T+1:K]}^{(t_{\max})}$ 

```

In our evaluation, we consider four different types of jammers: (J1) barrage jammers that transmit i.i.d. jamming symbols during the entire coherence interval, (J2) pilot jammers that transmit i.i.d. jamming symbols during the pilot phase but do not jam the data phase, (J3) data jammers that transmit i.i.d. jamming symbols during the data phase but do not jam the pilot phase, (J4) sparse jammers that transmit i.i.d. jamming symbols during some fraction α of randomly selected bursts of unit length (i.e., one time slot), but do not jam the remaining time slots. The jamming symbols are either circularly symmetric complex Gaussian or selected uniformly at random from the QPSK constellation. Unless stated otherwise, the jamming symbols are also independent of the UE transmit symbols $\mathbf{S}$. We quantify the strength of the jammer's interference relative to the strength of the average UE, either as the ratio between total receive *energy*

$$\rho_E \triangleq \frac{\mathbb{E}_{\mathbf{w}}[\|\mathbf{j}\mathbf{w}\|_2^2]}{\frac{1}{U} \mathbb{E}_{\mathbf{S}}[\|\mathbf{H}\mathbf{S}\|_2^2]}, \quad (25)$$

or as the ratio between receive *power during those phases that the jammer is jamming*

$$\rho_P = \frac{\rho_E}{\lambda}, \quad (26)$$

where λ is the jammer's duty cycle and equals 1 , $\frac{T}{K}$, $\frac{D}{K}$, or α for barrage, pilot, data, or sparse jammers, respectively.

B. Performance Baselines

Unless stated otherwise, we run MAED with $t_{\max} = 30$ iterations and compare it with the following baseline methods: The first baseline called "LMMSE" does not mitigate the jammer in any way and separately performs least-squares channel estimation and LMMSE data detection. The second baseline called "geniePOS" represents a jammer-robust variant of LMMSE that is furnished with ground-truth knowledge of the jammer channel $\mathbf{j}$ and projects the receive signals $\mathbf{Y}$ onto




Fig. 2. Uncoded bit error-rate (BER) for different detectors in the presence of a *strong* jammer. The jammer transmits Gaussian symbols either (a) during the entire coherence interval, (b) during the pilot phase only, (c) during the data phase only, or (d) in random unit-symbol bursts with a duty cycle of $\alpha = 20\%$. The jammer receive energy over the whole coherence interval is $\rho_E = 25$ dB higher than that of the average UE.




Fig. 3. Uncoded bit error-rate (BER) for different detectors in the presence of a *weak* jammer. The jammer transmits QPSK symbols either (a) during the entire coherence interval, (b) during the pilot phase only, (c) during the data phase only, or (d) in random unit-symbol bursts with a duty cycle of $\alpha = 20\%$. The jammer receive power during the jamming phase is as high as that of the average UE ($\rho_P = 0$ dB).

the orthogonal complement of $\text{span}(\mathbf{j})$ as in (4). The method then separately performs least-squares channel estimation and LMMSE data detection in this projected subspace. The third baseline called “JL-JED” serves as a performance upper bound and operates in a jammerless but otherwise equivalent scenario. JL-JED performs joint channel estimation and data detection by approximately solving (2) (with $\mathcal{S}$ relaxed to its convex hull $\mathcal{C}$) using the same FBS procedure as MAED (cf. Section III-C), except that it misses the projection $\tilde{\mathbf{P}}$.

C. Mitigation of Strong Gaussian Jammers

We first investigate the ability of MAED to mitigate strong jamming attacks. For this, we simulated Gaussian jammers with $\rho_E = 25$ dB of all four jamming types introduced in Section IV-A and evaluated the performance of MAED compared to the baselines of Section IV-B (Fig. 2). We point out that the performance of geniePOS and JL-JED is independent of the considered jammer: geniePOS uses the genie-provided jammer channel to null the jammer perfectly, and JL-JED operates on a jammerless system from the beginning. Unsurprisingly, the jammer-oblivious LMMSE baseline performs significantly worse than the jammer-robust geniePOS baseline under all attack scenarios. MAED succeeds in mitigating all four jamming attacks with very high effectiveness, even outperforming the

genie-assisted geniePOS method by a considerable margin.⁴ The efficacy of MAED is further reflected in the fact that its BER approaches the BER of the jammerless reference baseline JL-JED to within 1 dB in all considered scenarios.

D. Mitigation of Weak QPSK Jammers

We now turn to the analysis of more restrained jamming attacks in which the jammer transmits QPSK symbols with relative power $\rho_P = 0$ dB during its on-phase (to pass itself off as just another UE, for instance [7]). For now, we still make the assumption that the jamming symbols are independent of the UE transmit matrix $\mathbf{S}$. (We will consider an alternative scenario in Section IV-F.) Simulation results for all four jammer types are shown in Fig. 3. The baseline performance of geniePOS and JL-JED are again independent of the jammer and mirror the curves of Fig. 2. Because of the weaker jamming attacks, the jammer-oblivious LMMSE baseline performs much closer to the jammer-resistant geniePOS baseline. MAED again mitigates all attack types successfully, outperforming geniePOS and approaching the JL-JED baseline to within 1 dB.

Comparing Fig. 3 with Fig. 2 reveals an interesting phenomenon. MAED achieves better absolute performance under

⁴The potential for MAED to outperform geniePOS is a consequence of the superiority of JED over separating channel estimation from data detection.

strong jamming attacks than under weak ones, even if the difference is subtle. The reason for this behavior is the following: MAED searches for the jamming subspace by looking for a dominant dimension of the iterative residual error $\tilde{\mathbf{E}}^{(t)}$, see (21). If the received jamming energy is small compared to the received signal energy, then it becomes harder to distinguish the residual errors due to the jammer's impact from those that are caused by the errors in the channel and transmit matrix estimates $\tilde{\mathbf{H}}_P^{(t)}$ and $\tilde{\mathbf{S}}^{(t)}$.

E. What if No Jammer Is Present?

This observation leads to the question of how MAED performs if no jammer is present, or—equivalently—if a jammer does not transmit for a given coherence interval. Fig. 4(a) shows simulation results for this scenario. MAED still outperforms the LMMSE baseline at low SNR, but shows an error floor at high SNR. This error floor is caused by the slower convergence of MAED with (infinitely) weak jammers. Fig. 4(b) shows the jammerless performance of MAED for different numbers $t_{\max}$ of algorithm iterations. For $t_{\max} = 100$ iterations, MAED essentially achieves the excellent performance that it has in combination with strong jammers. For $t_{\max} = 10$ iterations, however, MAED exhibits an error floor as high as 0.2%. In contrast, in the presence of a $\rho_E = 25$ dB strong barrage Gaussian jammer, MAED requires no more than $t_{\max} = 10$ iterations for optimal performance.

The slow convergence in the absence of a jammer can be explained by the fact that, in every iteration, the strongest dimension of the residual error matrix $\tilde{\mathbf{E}}^{(t)}$ is mistakenly attributed to a hypothesized jammer instead of to the residual errors in the channel and transmit matrix estimates $\tilde{\mathbf{H}}_P^{(t)}$ and $\tilde{\mathbf{S}}^{(t)}$. This recurring misattribution prevents fast convergence. Nonetheless, while MAED was conceived for jammer mitigation, it shows robust performance even in the absence of jamming.

F. What Happens with a Truly Smart Jammer?

Finally, we turn to a scenario in which the jammer knows the UE pilot sequences and attacks a specific UE by transmitting that UE's pilot sequence during the pilot phase (at $\rho_P = 25$ dB higher power). The jammer does not transmit during the data phase. Fig. 5 shows simulation results for this scenario. The geniePOS baseline nulls the jammer perfectly using its ground-truth knowledge. Thus, its performance remains unaffected regardless of the jammer. In contrast, MAED exhibits an error floor as high as 1%, only marginally outperforming the LMMSE baseline. Excluding the attacked UE and evaluating the BER among the remaining 31 UEs (labeled $\bar{\mathbf{U}}_j$ in Fig. 5(a)) reveals that the decoding errors are focused entirely on the attacked UE, and that the BER among the remaining UEs appears to be unaffected by the jammer. This experiment shows that MAED cannot identify the jammer's subspace if the jammer passes itself off as a UE by transmitting that UE's pilot sequence. It is not clear, however, whether such a jammer could be distinguished from a legitimate UE, even in principle. One way to prevent smart jammers from utilizing such impersonation attacks would be to use encrypted pilot sequences [34].




Fig. 4. Uncoded bit error-rate (a) for different detectors in absence of a jammer and (b) for MAED with different number $t_{\max}$ of iterations in absence of a jammer (JL), as well as in the presence of a barrage jammer that transmits Gaussian symbols at $\rho_E = 25$ dB higher energy than the average UE ($J_{25\text{dB}}$).




Fig. 5. Uncoded bit error rate (BER) for a smart jammer that (a) transmits the pilot sequence of $\mathbf{U}_{E_j}$ or (b) transmits the average of multiple UE pilot sequences. The jammer transmits at $\rho_P = 25$ dB higher power than the average UE. The $\bar{\mathbf{U}}_j$ curve depicts the BER averaged over all UEs except $\mathbf{U}_{E_j}$.

Finally, Fig. 5(b) shows the performance of MAED (over all 32 UEs) when the jammer transmits the average of multiple pilot sequences during the pilot phase (and refrains from transmitting during the data phase). Evidently, a jammer that targets multiple UEs quickly enables MAED to locate the jammer's subspace and mitigate the jammer effectively.

V. CONCLUSIONS

We have proposed MAED in order to mitigate smart jamming attacks on the uplink of massive MU-MIMO systems. In contrast to existing mitigation methods, MAED does not rely on jamming activity during any particular epoch for successful jammer mitigation. Instead, our method exploits the fact that the jammer's subspace remains constant within a coherence interval. To this end, MAED uses a novel problem formulation that combines jammer estimation and mitigation, channel estimation, and data detection. The resulting optimization problem is approximately solved using an efficient iterative algorithm. Without requiring any *a priori* knowledge, MAED is able to effectively mitigate a wide range of jamming attacks. In particular, MAED succeeds in mitigating attack types like data jamming and sparse jamming, for which—to the best of our knowledge—no mitigation methods have existed so far.

REFERENCES

- [1] "Satellite-navigation systems such as GPS are at risk of jamming," *The Economist*. [Online]. Available: <https://www.economist.com/science-and-technology/2021/05/06/satellite-navigation-systems-such-as-gps-are-at-risk-of-jamming>
- [2] P. Popovski, "Ultra-reliable communication in 5G wireless systems," in *Proc. Int. Conf. 5G Ubiquitous Connectivity*, Nov. 2014, pp. 146–151.
- [3] G. Marti, O. Castañeda, and C. Studer, "Jammer mitigation via beam-slicing for low-resolution mmWave massive MU-MIMO," *IEEE Open J. Circuits Syst.*, vol. 2, 2021.
- [4] Q. Yan, H. Zeng, T. Jiang, M. Li, W. Lou, and Y. T. Hou, "Jamming resilient communication using MIMO interference cancellation," *IEEE Trans. Inf. Forensics Security*, vol. 11, no. 7, pp. 1486–1499, Jul. 2016.
- [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] H. Zeng, C. Cao, H. Li, and Q. Yan, "Enabling jamming-resistant communications in wireless MIMO networks," in *Proc. IEEE Conf. Commun. Netw. Security (CNS)*, Oct. 2017, pp. 1–9.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] H. Akhlaghpasand, E. Björnsson, and S. M. Razavizadeh, "Jamming-robust uplink transmission for spatially correlated massive MIMO systems," *IEEE Trans. Commun.*, vol. 68, no. 6, pp. 3495–3504, Mar. 2020.
- [11] G. Marti, O. Castañeda, S. Jacobsson, G. Durisi, T. Goldstein, and C. Studer, "Hybrid jammer mitigation for all-digital mmWave massive MU-MIMO," in *Asilomar Conf. Signals, Syst., Comput.*, Nov. 2021, pp. 1–7.
- [12] R. Miller and W. Trappe, "Subverting MIMO wireless systems by jamming the channel estimation procedure," in *Proc. ACM Conf. Wireless Netw. Security*, Mar. 2010, pp. 19–24.
- [13] R. Miller and W. Trappe, "On the vulnerabilities of CSI in MIMO wireless communication systems," *IEEE Trans. Mobile Comput.*, vol. 11, no. 8, pp. 1386–1398, Aug. 2011.
- [14] T. C. Clancy, "Efficient OFDM denial: Pilot jamming and pilot nulling," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2011, pp. 1–5.
- [15] S. Sodagari and T. C. Clancy, "Efficient jamming attacks on MIMO channels," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2012, pp. 852–856.
- [16] M. Lichtman, J. H. Reed, T. C. Clancy, and M. Norton, "Vulnerability of LTE to hostile interference," in *Proc. IEEE Global Conf. Signal Inf. Process.*, Dec. 2013, pp. 285–288.
- [17] M. Lichtman, R. Rao, V. Marojevic, J. Reed, and R. P. Jover, "5G NR jamming, spoofing, and sniffing: Threat assessment and mitigation," in *IEEE Int. Conf. Commun. Workshop (ICCW)*, May 2018, pp. 1–6.
- [18] M. Lichtman, R. P. Jover, M. Labib, R. Rao, V. Marojevic, and J. H. Reed, "LTE/LTE-A jamming, spoofing, and sniffing: threat assessment and mitigation," *IEEE Commun. Mag.*, vol. 54, no. 4, pp. 54–61, Apr. 2016.
- [19] F. Girke, F. Kurtz, N. Dorsch, and C. Wietfeld, "Towards resilient 5G: Lessons learned from experimental evaluations of LTE uplink jamming," in *IEEE Int. Conf. Commun. Workshop (ICCW)*, May 2019, pp. 1–6.
- [20] M. J. La Pan, T. C. Clancy, and R. W. McGwier, "Jamming attacks against OFDM timing synchronization and signal acquisition," in *IEEE Mil. Commun. Conf. (MILCOM)*, Oct. 2012, pp. 1–7.
- [21] H. Vikalo, B. Hassibi, and P. Stoica, "Efficient joint maximum-likelihood channel estimation and signal detection," *IEEE Trans. Wireless Commun.*, vol. 5, no. 7, pp. 1838–1845, Jul. 2006.
- [22] W. Xu, M. Stojnic, and B. Hassibi, "On exact maximum-likelihood detection for non-coherent MIMO wireless systems: a branch-estimate-bound optimization framework," in *IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2008, pp. 2017–2021.
- [23] E. Kofidis, C. Chatzichristos, and A. L. de Almeida, "Joint channel estimation/data detection in MIMO-FBMC/OQAM systems—a tensor-based approach," in *Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2017, pp. 420–424.
- [24] O. Castaneda, T. Goldstein, and C. Studer, "VLSI designs for joint channel estimation and data detection in large SIMO wireless systems," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 65, no. 3, pp. 1120–1132, Mar. 2018.
- [25] B. B. Yilmaz and A. T. Erdogan, "Channel estimation for massive MIMO: A semiblind algorithm exploiting QAM structure," in *Asilomar Conf. Signals, Syst., Comput.*, IEEE, Nov. 2019, pp. 2077–2081.
- [26] H. He, C.-K. Wen, S. Jin, and G. Y. Li, "Model-driven deep learning for MIMO detection," *IEEE Trans. Signal Process.*, vol. 68, pp. 1702–1715, Feb. 2020.
- [27] H. Song, X. You, C. Zhang, and C. Studer, "Soft-output joint channel estimation and data detection using deep unfolding," in *IEEE Inf. Theory Workshop (ITW)*, Oct. 2021, pp. 1–5.
- [28] T. Goldstein, C. Studer, and R. G. Baraniuk, "A field guide to forward-backward splitting with a FASTA implementation," Feb. 2016. [Online]. Available: <https://arxiv.org/abs/1411.3406>
- [29] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, Jan. 2014.
- [30] J. Barzilai and J. M. Borwein, "Two-point step size gradient methods," *IMA J. Numer. Anal.*, vol. 8, no. 1, pp. 141–148, 1988.
- [31] B. Zhou, L. Gao, and Y.-H. Dai, "Gradient methods with adaptive step-sizes," *Comp. Optimization Applicat.*, vol. 35, no. 1, pp. 69–86, 2006.
- [32] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge Univ. Press, 2013.
- [33] G. H. Golub and C. F. van Loan, *Matrix Computations*, 3rd ed. The Johns Hopkins Univ. Press, 1996.
- [34] Y. O. Basciftci, C. E. Koksul, and A. Ashikhmin, "Securing massive MIMO at the physical layer," in *IEEE Conf. Commun. Netw. Security (CNS)*, Sep. 2015, pp. 272–280.