

# Detection by Sampling: Massive MIMO Detector based on Langevin Dynamics

Nicolas Zilberstein\*, Chris Dick†, Rahman Doost-Mohammady\*, Ashutosh Sabharwal\*, Santiago Segarra\*  
\*Rice University, USA †Nvidia, USA

**Abstract**—Optimal symbol detection in multiple-input multiple-output (MIMO) systems is known to be an NP-hard problem. Hence, the objective of any detector of practical relevance is to get reasonably close to the optimal solution while keeping the computational complexity in check. In this work, we propose a MIMO detector based on an annealed version of Langevin (stochastic) dynamics. More precisely, we define a stochastic dynamical process whose stationary distribution coincides with the posterior distribution of the symbols given our observations. In essence, this allows us to approximate the maximum a posteriori estimator of the transmitted symbols by sampling from the proposed Langevin dynamic. Furthermore, we carefully craft this stochastic dynamic by gradually adding a sequence of noise with decreasing variance to the trajectories, which ensures that the estimated symbols belong to a pre-specified discrete constellation. Through numerical experiments, we show that our proposed detector yields state-of-the-art symbol error rate performance.

**Index Terms**—Massive MIMO detection, Langevin dynamics, Markov chain Monte Carlo

## I. INTRODUCTION

Massive multiple-input multiple-output (MIMO) systems are crucial for modern and future communications [1], [2]. They are expected to play a key role in moving from the fifth to the sixth generation of cellular communications by achieving high data rates and spectral efficiency [3]. In massive MIMO systems, base stations are equipped with a large number of antennas, enabling them to handle several users simultaneously. However, these systems entail many challenges such as designing low complexity MIMO detection schemes, which is the focus of our paper.

Exact MIMO detection is an NP-hard problem [4]. Given  $N_u$  users and a modulation of  $K$  symbols, the exact maximum likelihood (ML) estimator has an exponential decoding complexity  $\mathcal{O}(K^{N_u})$ . Thus, obtaining this ML estimate is computationally infeasible and becomes intractable even for moderately-sized systems. Many approximate solutions for symbol detection have been proposed in the classical literature including zero forcing (ZF) and minimum mean squared error (MMSE) [5]. Although both (linear) detectors have low complexity and good performance for small systems, their performance degrades severely for larger systems [6]. Another classical detector is approximate message passing (AMP), which is asymptotically optimal for large MIMO systems with Gaussian channels but degrades significantly for other (more practical) channel distributions [7]. In the past few years, several massive MIMO symbol detectors based on machine

learning – and, in particular, deep learning – have been derived. One can roughly categorize them into channel-specific methods, like MMNet [8], and channel-agnostic methods like RE-MIMO [9], OAMPNet [10], and hyperMIMO [11], [12].

An alternative family of detectors is based on Markov chain Monte Carlo (MCMC) methods [6, Chapter 8]. Given that the ML estimator is prohibitive for large systems, these methods seek for a solution by generating candidate samples from the search space. In particular, in [13] a detector based on the Gibbs sampler was presented. Recently, in [14], the authors proposed a detector based on Metropolis-Hasting, in which they propose to make a random walk along the gradient descent direction of the least-square surface defined by the continuous-relaxed version of the ML. In the past few years, in the context of image processing, sampling algorithms based on the Langevin dynamic have been proposed as generative models or to solve inverse problems. This iterative technique enables sampling from a given distribution by leveraging the availability of the score function (the gradient of the log-probability density function) without the necessity of computing the classical acceptance/rejection step in MCMC methods. In [15], an *annealed* Langevin dynamic is used in the context of generative modeling for images. Assuming an unknown distribution of the images, they parameterize the score function as a neural network and use the annealed Langevin dynamic to sample from the underlying probability distribution. In [16], the authors proposed to solve noisy image inverse problems by sampling from the posterior.

In this work, we propose the first method that uses annealed Langevin dynamics for MIMO detection. Given that the transmitted symbols come from a discrete constellation, we leverage the annealed process to include information of the prior in the dynamic. As we have access to the prior distribution, we can define a closed-form expression for the score of the prior through the MMSE estimator. This allows us to avoid the training process required in state-of-the-art learning-based detectors. Hence, our detector can be applied to any observed channel and can handle a different number of users and mixed modulation schemes.

**Contribution.** The contributions of this paper are twofold:

- 1) We propose a novel detector based on *annealed* Langevin dynamics, allowing us to include information of our discrete prior in the exploration of the posterior distribution.
- 2) Through numerical experiments, we analyze the behavior of our method for different hyperparameter settings and demonstrate that the proposed detector achieves lower symbol error rate (SER) than baseline methods for massive MIMO systems.

This work was partially supported by Nvidia. Email: {nzilberstein, doost, ashu, segarra}@rice.edu, cdick@nvidia.com.

## II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a communication channel with  $N_u$  single-antenna transmitters or users and a receiving base station with  $N_r$  antennas. The forward model for this MIMO system is given by

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{z}, \quad (1)$$

where  $\mathbf{H} \in \mathbb{C}^{N_r \times N_u}$  is the channel matrix,  $\mathbf{z} \sim \mathcal{CN}(\mathbf{0}, \sigma_0^2 \mathbf{I}_{N_r})$  is a vector of complex circular Gaussian noise,  $\mathbf{x} \in \mathcal{X}^{N_u}$  is the vector of transmitted symbols,  $\mathcal{X}$  is a finite set of constellation points, and  $\mathbf{y} \in \mathbb{C}^{N_r}$  is the received vector. In this work, a quadrature amplitude modulation (QAM) is used and symbols are normalized to attain unit average power. It is assumed that the constellation is the same for all transmitters and each symbol has the same probability of being chosen by the users  $N_u$ . Moreover, perfect channel state information (CSI) is assumed, which means that  $\mathbf{H}$  and  $\sigma_0^2$  are known at the receiver.<sup>1</sup> Under this setting, the MIMO detection problem can be defined as follows.

**Problem 1:** *Given perfect CSI and an observed  $\mathbf{y}$  following (1), find an estimate of  $\mathbf{x}$ .*

Given the stochastic nature of  $\mathbf{z}$  in (1), a natural way of solving Problem 1 is to search for the  $\mathbf{x}$  that maximizes its *posterior* probability given the observations  $\mathbf{y}$ . Hence, the optimal decision rule can be written as

$$\begin{aligned} \hat{\mathbf{x}}_{\text{MAP}} &= \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}^{N_u}} p(\mathbf{x}|\mathbf{y}, \mathbf{H}) \\ &= \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}^{N_u}} p_{\mathbf{z}}(\mathbf{y} - \mathbf{H}\mathbf{x})p(\mathbf{x}), \end{aligned} \quad (2)$$

where we have applied Bayes' rule. As we assume that the symbols' prior distribution is uniform among the constellation elements and the measurement noise  $\mathbf{z}$  is Gaussian, the maximum a posteriori (MAP) detector reduces to an ML detector. Specifically, (2) boils down to solving the following optimization problem

$$\hat{\mathbf{x}}_{\text{ML}} = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}^{N_u}} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2, \quad (3)$$

which is NP-hard due to the finite constellation constraint  $\mathbf{x} \in \mathcal{X}^{N_u}$ , rendering  $\hat{\mathbf{x}}_{\text{ML}}$  intractable in practical applications. Consequently, several schemes have been proposed in the last decades to provide efficient approximate solutions to Problem 1, as mentioned in Section I. In this paper, we propose to solve Problem 1 by (approximately) sampling from the posterior distribution in (2) using an annealed Langevin dynamic.

## III. LANGEVIN FOR MIMO DETECTION

In Section III-A we briefly introduce the Langevin dynamic while in Section III-B we explain how we propose to use it for MIMO detection. In particular, we detail the expressions of the score functions involved in the sampling process to solve our Problem 1.

<sup>1</sup>To avoid notation overload, we adopt the convention that whenever we assume  $\mathbf{H}$  to be known,  $\sigma_0^2$  is also known.

### A. Langevin dynamics

The Langevin dynamic algorithm is an MCMC algorithm [17], [18], described by the following equation<sup>2</sup>

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \epsilon \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \sqrt{2\epsilon} \mathbf{w}_t, \quad (4)$$

where  $p(\mathbf{x})$  is the target distribution from which we want to generate samples  $\mathbf{x} \in \mathbb{R}^N$  and  $\mathbf{w}_t \sim \mathcal{N}(0, \mathbf{I}_N)$ . The dynamic in (4) explores the target distribution by moving in the direction of the gradient of the logarithm of the target density  $\nabla_{\mathbf{x}} \log p(\mathbf{x})$ , known as *score function*. In essence, it is a combination of stochastic gradient ascent in the direction of the score function and injected noise, which allows the method to avoid collapsing to local maxima. Under some regularity conditions [20], the distribution of  $\mathbf{x}_T$  is equal to  $p(\mathbf{x})$  when  $\epsilon \rightarrow 0$  and  $T \rightarrow \infty$ , in which case  $\mathbf{x}_T$  becomes an exact sample from  $p(\mathbf{x})$ . In practice, neither  $\epsilon \rightarrow 0$  nor  $T \rightarrow \infty$ , so a Metropolis-Hastings acceptance/rejection step is used to ensure convergence, leading to the so-called Metropolis-adjusted Langevin algorithm (MALA) [18]. An alternative, proposed in [20], is to use a time-inhomogeneous variant of (4), i.e., defining a variable step size  $\epsilon_t$ . They demonstrate that when  $\epsilon_t$  decreases to zero for large  $t$ , the error becomes negligible and the acceptance/rejection step can be omitted. It should be noted that the *only* requirement for sampling from  $p(\mathbf{x})$  using this procedure is knowing the score function.

### B. Detection by sampling from the posterior distribution

Given the intractability of Problem 1 due to the finite constellation constraint, we propose to *generate a set of samples that approximately come from the posterior distribution  $p(\mathbf{x}|\mathbf{y}, \mathbf{H})$  using (4) and then select the one that minimizes the objective in (3)*. The key ingredient in the Langevin dynamic is the score function, which for our case is given by  $\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y}, \mathbf{H})$ . After applying Bayes' rule, this score function can be rewritten as

$$\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y}, \mathbf{H}) = \nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}, \mathbf{H}) + \nabla_{\mathbf{x}} \log p(\mathbf{x}), \quad (5)$$

where the term  $\nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}, \mathbf{H})$  corresponds to the score function of the likelihood and  $\nabla_{\mathbf{x}} \log p(\mathbf{x})$  to the score function of the prior. Notice that this latter term is not well defined due to the discrete nature of the symbols.

To circumvent this obstacle, inspired by [16], we approximate the prior by using an *annealed* version of the Langevin dynamic. Specifically, instead of working with the discrete symbols  $\mathbf{x}$ , we define a perturbed version of the symbols  $\tilde{\mathbf{x}} = \mathbf{x} + \mathbf{n}$  with  $\mathbf{n} \sim \mathcal{CN}(0, \sigma^2 \mathbf{I})$ , for different values of  $\sigma^2$ . Note that  $\tilde{\mathbf{x}}$  now has a continuous prior allowing us to run a Langevin dynamic in  $\tilde{\mathbf{x}}$  instead of  $\mathbf{x}$ . Moreover, if we make  $\sigma^2 \rightarrow 0$  then  $\tilde{\mathbf{x}}$  concentrates around  $\mathbf{x}$ , allowing us to effectively sample from  $p(\mathbf{x}|\mathbf{y}, \mathbf{H})$ , as wanted.

In a nutshell, the algorithm works as follows. First, we initialize  $\tilde{\mathbf{x}}_0$  uniformly at random in  $[-1, 1] \times [-1, 1]$ , since the symbols are assumed to be normalized. Then, we follow

<sup>2</sup>This is technically known as unadjusted Langevin algorithm (ULA), which is obtained from the Euler-Maruyama discretization of the overdamped Langevin stochastic differential equation [19].




Fig. 1: Scheme of the annealed process. We consider a QPSK constellation and at each level we add Gaussian noise. The variance of the noise decreases for higher levels. In the last level  $L$ , the Gaussian is very sharp around each symbol, mimicking our true discrete prior over the constellation.

the direction of the score function of the log-posterior density of the perturbed symbol  $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}|\mathbf{y}, \mathbf{H})$ , starting with a high  $\sigma$  and gradually decreasing its value until  $\tilde{\mathbf{x}} \approx \mathbf{x}$ . Apart from enabling the approximation of the score function of the prior distribution, the annealing process also improves the mixing time of the Langevin dynamic [15]. This is particularly important in multimodal distributions, as is the case of MIMO detection. Having introduced the high-level idea of our method, we now provide more details on the *annealing process*, exact expressions for the terms in the *score function*, and a step-by-step description of the *algorithm*.

**Annealing process.** We define a sequence of noise levels  $\{\sigma_l\}_{l=1}^L$  such that  $\sigma_1 > \sigma_2 > \dots > \sigma_L > 0$ . Then, at each level we define a perturbed version of the true symbols  $\mathbf{x}$

$$\tilde{\mathbf{x}}_l = \mathbf{x} + \mathbf{n}_l, \quad (6)$$

where  $\mathbf{n}_l \sim \mathcal{CN}(0, \sigma_l^2 \mathbf{I})$ . A representation (for a QPSK modulation) of this process is shown in Fig. 1. Since the variance of the noise injected at each level is a predefined parameter, we will design the sequence in such way that the noise injected in the last levels is very small, approximating the true discrete distribution given by a set of delta functions at each symbol with uniform weight.

**Score function.** Given the perturbed symbols in (6), the forward model in (1) can be rewritten as

$$\mathbf{y} = \mathbf{H}\tilde{\mathbf{x}}_l + (\mathbf{z} - \mathbf{H}\mathbf{n}_l). \quad (7)$$

In this new forward model, the likelihood is given by  $p(\mathbf{y}|\tilde{\mathbf{x}}_l, \mathbf{H}) = p(\mathbf{z} - \mathbf{H}\mathbf{n}_l|\tilde{\mathbf{x}}_l)$ , which is not Gaussian: although  $\mathbf{n}_l$  is a Gaussian random variable, when conditioning on  $\tilde{\mathbf{x}}_l$  the conditional distribution is no longer Gaussian due to (6). However, an analytical expression for the score of the likelihood can still be obtained by following the approach in [16], where a synthetic annealed noise carved from the measurement noise  $\mathbf{z}$  is constructed in a gradual fashion. Furthermore, in order to get a tractable expression, we have to rely on the singular value decomposition (SVD) of the channel matrix given by  $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  as well as in the spectral representation of  $\tilde{\mathbf{x}}_l$  and  $\mathbf{y}$  defined as  $\tilde{\chi}_l = \mathbf{V}^\top \tilde{\mathbf{x}}_l$  and  $\boldsymbol{\eta} = \mathbf{U}^\top \mathbf{y}$ . In essence, the gradual noise addition is constructed in such a way that the noise  $\mathbf{z} - \mathbf{H}\mathbf{n}_l$  is uncorrelated and independent of  $\tilde{\mathbf{x}}_l$ , given the singular values  $s_j = [\mathbf{\Sigma}]_{jj}$ . More precisely, the distribution of  $\mathbf{z} - \mathbf{H}\mathbf{n}_l$  is given by a multivariate

Gaussian distribution, where each component is distributed as  $[\mathbf{z} - \mathbf{H}\mathbf{n}_l]_j \sim \mathcal{N}(0, |\sigma_0^2 - \sigma_l^2 s_j^2|)$  for  $j = 1, \dots, N_u$ .

With this spectral representation in mind, our goal is to run a Langevin dynamic whose score function for every noise level  $l$  is given by [cf. (5)]

$$\nabla_{\tilde{\chi}_l} \log p(\tilde{\chi}_l|\boldsymbol{\eta}, \mathbf{H}) = \nabla_{\tilde{\chi}_l} \log p(\boldsymbol{\eta}|\tilde{\chi}_l, \mathbf{H}) + \nabla_{\tilde{\chi}_l} \log p(\tilde{\chi}_l). \quad (8)$$

We now provide closed-form expression for both constituent terms in this score function.

*i) Score of the likelihood:* Given the above discussion, the final expression for the score of the likelihood in the spectral domain is given by

$$\nabla_{\tilde{\chi}_l} \log p(\boldsymbol{\eta}|\tilde{\chi}_l, \mathbf{H}) = \mathbf{\Sigma}^\top |\sigma_0^2 \mathbf{I} - \sigma_l^2 \mathbf{\Sigma} \mathbf{\Sigma}^\top|^\dagger (\boldsymbol{\eta} - \mathbf{\Sigma} \tilde{\chi}_l). \quad (9)$$

To give some intuition, the score function of the likelihood is given by the gradient of a multivariate Gaussian distribution: the residual error  $(\boldsymbol{\eta} - \mathbf{\Sigma} \tilde{\chi}_l) = (\mathbf{U}^\top \mathbf{y} - \mathbf{\Sigma} \mathbf{V}^\top \tilde{\mathbf{x}}_l)$  is multiplied by the (pseudo-)inverse of the covariance matrix, which is diagonal with entries given by  $|\sigma_0^2 - \sigma_l^2 s_j^2|$ . For details about the derivation of this expression see [16].

*ii) Score of the annealed prior:* We first notice that  $\nabla_{\tilde{\chi}_l} \log p(\tilde{\chi}_l) = \mathbf{V}^\top \nabla_{\tilde{\mathbf{x}}_l} \log p(\tilde{\mathbf{x}}_l)$  due to the orthogonality of  $\mathbf{V}$ . Moreover, based on the Tweedie's identity [21], we can relate the score function  $\nabla_{\tilde{\mathbf{x}}_l} \log p(\tilde{\mathbf{x}}_l)$  and the MMSE denoiser as follows

$$\nabla_{\tilde{\mathbf{x}}_l} \log p(\tilde{\mathbf{x}}_l) = \frac{\mathbb{E}_{\sigma_l}[\mathbf{x}|\tilde{\mathbf{x}}_l] - \tilde{\mathbf{x}}_l}{\sigma_l^2}. \quad (10)$$

In particular, the conditional expectation can be calculated elementwise as

$$\mathbb{E}_{\sigma_l}[x_j|\tilde{\chi}_l] = \frac{1}{Z} \sum_{x_k \in \mathcal{X}} x_k \exp\left(\frac{-||[\tilde{\chi}_l]_j - x_k||^2}{2\sigma_l^2}\right), \quad (11)$$

where  $Z = \sum_{x_k \in \mathcal{X}} \exp\left(\frac{-||[\tilde{\chi}_l]_j - x_k||^2}{2\sigma_l^2}\right)$  and  $j = 1, \dots, N_u$ .

**Algorithm.** The algorithm to generate samples  $\hat{\mathbf{x}}$  from the (approximate) posterior  $p(\mathbf{x}|\mathbf{y}, \mathbf{H})$  is shown in Alg. 1. As discussed in [16], when computing the entries of the score function in (8) using the expressions in (9) and (10), one of these terms might be negligible with respect to the other depending on the noise level. Thus, the elementwise score of the posterior will be given by

$$[\nabla_{\tilde{\chi}_l} \log(\boldsymbol{\eta}|\tilde{\chi}_l, \mathbf{H})]_j = \begin{cases} [\nabla_{\tilde{\chi}_l} \log p(\boldsymbol{\eta}|\tilde{\chi}_l, \mathbf{H}) + \mathbf{V}^\top \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})]_j, & \sigma_0 \geq \sigma_l s_j \\ [\nabla_{\tilde{\chi}_l} \log p(\boldsymbol{\eta}|\tilde{\chi}_l, \mathbf{H})]_j, & \sigma_0 < \sigma_l s_j \\ [\mathbf{V}^\top \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})]_j, & s_j = 0. \end{cases} \quad (12)$$

Intuitively, when the injected noise at level  $l$  is such that  $\sigma_l s_j > \sigma_0$ , then the contribution of the score of the prior is negligible and can be ignored. Similarly, whenever  $s_j = 0$ , the corresponding entry  $\eta_j$  is uninformative and the score of the likelihood can be ignored. Furthermore, a refinement that we incorporate in the algorithm is the use of position-dependent step sizes. Instead of using a constant *scalar* step size  $\epsilon$  as in (4) or even time-varying versions of it, in Alg. 1

---

**Algorithm 1** Annealed Langevin for MIMO detection

---

**Require:**  $T, \{\sigma_l\}_{l=1}^L, \epsilon, \sigma_0, \mathbf{H}, \mathbf{y}$   
Compute SVD of  $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$   
Initialize  $\tilde{\mathbf{x}}_{t=0, l=1}$  with random noise  $\mathcal{U}[-1, 1]$   
**for**  $l = 1$  to  $L$  **do**  
     $[\mathbf{\Lambda}_l]_{jj} = \begin{cases} \frac{\epsilon\sigma_l^2}{\sigma_l^2} (1 - \frac{\sigma_l^2}{\sigma_0^2} s_j^2) & \text{if } \sigma_l s_j \leq \sigma_0 \\ \frac{\epsilon}{\sigma_l^2} (\sigma_l^2 - \frac{\sigma_0^2}{s_j^2}) & \text{if } \sigma_l s_j > \sigma_0 \end{cases}$   
    **for**  $t = 0$  to  $T - 1$  **do**  
        Draw  $\mathbf{w}_t \sim \mathcal{N}(0, \mathbf{I})$   
        Compute  $\nabla_{\tilde{\mathbf{x}}_{t,l}} \log p(\boldsymbol{\eta} | \tilde{\mathbf{x}}_{t,l}, \mathbf{H})$  as in (9)  
        Compute  $\nabla_{\tilde{\mathbf{x}}_{t,l}} \log p(\tilde{\mathbf{x}}_{t,l})$  as in (10)  
        Compute  $\nabla_{\tilde{\mathbf{x}}_{t,l}} \log p(\tilde{\mathbf{x}}_{t,l} | \boldsymbol{\eta}, \mathbf{H})$  as in (12)  
         $\tilde{\mathbf{x}}_{t+1,l} = \tilde{\mathbf{x}}_{t,l} + \mathbf{\Lambda}_l \nabla_{\tilde{\mathbf{x}}_{t,l}} \log p(\tilde{\mathbf{x}}_{t,l} | \boldsymbol{\eta}, \mathbf{H}) + \sqrt{2\mathbf{\Lambda}_l} \mathbf{w}_t$   
    **end for**  
     $\tilde{\mathbf{x}}_{0, l+1} = \tilde{\mathbf{x}}_{T, l}$   
**end for**  
**return**  $\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}^{N_u}} \|\mathbf{x} - \mathbf{V}\tilde{\mathbf{x}}_{T,L}\|_2^2$

---

we employ level-dependent *diagonal matrices*  $\mathbf{\Lambda}_l$ . In this way, different entries of our vector-valued Langevin dynamic can be updated at different rates depending on the singular values of the channel under consideration. Finally, given that  $\sigma_L \neq 0$ , the sample will be very close to the constellation but not exactly. Hence, we take  $\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}^{N_u}} \|\mathbf{x} - \mathbf{V}\tilde{\mathbf{x}}_{T,L}\|_2^2$ .

Given that Alg. 1 is stochastic, one can generate several samples  $\hat{\mathbf{x}}$  from the same (approximate) posterior distribution by running the algorithm multiple times. Therefore, as we want to approximate the MAP estimate – equivalently for this case, the ML estimate – we run  $M$  different Langevin trajectories for each pair  $\{\mathbf{y}, \mathbf{H}\}$  and keep the sample that minimizes (3). Formally, given  $M$  samples  $\{\hat{\mathbf{x}}_m\}_{m=1}^M$  obtained from Alg. 1, our final estimate is given by

$$\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x} \in \{\hat{\mathbf{x}}_m\}_{m=1}^M} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2. \quad (13)$$

Notice that these  $M$  trajectories can be run in parallel, as they are independent of each other.

**Computational complexity.** The first step in Alg. 1 is to compute the SVD of the channel  $\mathbf{H}$ , whose complexity is  $\mathcal{O}(N_u N_r \min\{N_u, N_r\})$ , and is done only once per channel. Moreover, notice that the matrices involved in each iteration are diagonal so the complexity of multiplying them is  $\mathcal{O}(N_u^2)$ , while the matrix inversion is  $\mathcal{O}(N_u)$ . Finally, given a modulation of  $K$  symbols, the complexity of (11) is  $\mathcal{O}(KN_u)$ . Hence, one iteration has a complexity of  $\mathcal{O}(N_u^2 + KN_u)$ . The overall complexity, including the SVD computation and all the iterations, is  $\mathcal{O}(N_u N_r \min\{N_u, N_r\} + LT(N_u^2 + KN_u))$ . Regarding the  $M$  trajectories, observe that these are independent of each other, so they can be computed in parallel. Therefore, the bottlenecks are twofold: the SVD computation and the number of iterations  $LT$ . While the former is inevitable, the latter is a parameter that we control and represents a trade-off between SER performance and computational complexity. In Section IV, we present some numerical experiments that analyze this trade-off and the impact on the SER performance.

## IV. RESULTS

In this section we present the results of our proposed method.<sup>3</sup> We start by presenting the channel model and the simulation setup. Then, we analyze the SER performance of the proposed method when considering different noise levels  $L$  and different numbers of trajectories  $M$ . Finally, we compare our method with both classical and learning-based baseline detectors.

**Channel model and simulation settings.** The channel model is generated following the Kronecker correlated model

$$\mathbf{H} = \mathbf{R}_r^{1/2} \mathbf{H}_e \mathbf{R}_u^{1/2}, \quad (14)$$

where  $\mathbf{H}_e$  is a Rayleigh fading channel matrix and  $\mathbf{R}_r$  and  $\mathbf{R}_u$  are the spatial correlation matrices at the receiver and transmitters, respectively, generated according to the exponential correlation matrix model with a correlation coefficient  $\rho$ ; see [22] for details. The signal-to-noise ratio (SNR) is given by

$$\text{SNR} = \frac{\mathbb{E}[\|\mathbf{H}\mathbf{x}\|^2]}{\mathbb{E}[\|\mathbf{z}\|^2]} = \frac{N_u}{\sigma_0^2 N_r}. \quad (15)$$

The simulation environment includes a base station with  $N_r = 64$  receiver antennas and  $N_u = 32$  single-antenna users. We consider a 16-QAM modulation and  $\rho = 0.6$ . The value of  $\epsilon$  is fixed at  $3 \times 10^{-5}$ , while the number of samples per noise level at  $T = 70$ . The batch size for testing is 5000.

**Varying the number of noise levels.** In the first experiment, given the sequence of noise levels with variance  $\{\sigma_l\}_{l=1}^L$ , we fix  $\sigma_1 = 1$ ,  $\sigma_L = 0.01$  and  $M = 40$  trajectories and change the number of noise levels between them. We consider four cases where  $L \in \{5, 10, 16, 20\}$ ; see Fig. 2a. First, notice that the performance when  $L = 5$  is much worse than the other three cases. This implies that the algorithm is not able to sufficiently explore the search space. On the other hand, the gap between the other three cases is much smaller. Therefore, a trade-off between computational burden and performance exists: when considering more levels, the SER performance improves at the cost of increasing the running time. From this experiment, we conclude that at least  $L = 10$  levels are needed in order to perform as well as the existing state-of-the-art detectors.

**Varying the number of trajectories.** The performance of the detector as a function of the number  $M$  of different Langevin trajectories [cf. (13)] is shown in Fig. 2b. We analyze five cases where  $M \in \{1, 5, 10, 20, 40\}$ . For all the cases, we consider  $L = 20$  noise levels. From the results, we see that  $M$  is a hyperparameter that has a high impact on the overall performance of the detector. In particular, if we consider only  $M = 1$ , then the performance degrades severely, with a SER in the order of the classical MMSE detector (not shown in the figure). However, if we consider  $M = 40$  trajectories, the proposed method outperforms state-of-the-art detectors, as we illustrate in our next experiment.

**Performance comparison with baseline methods.** Based on our previous experiments, we set  $L = 20$  noise levels between  $\sigma_1 = 1$  and  $\sigma_{20} = 0.01$ , and we run  $M = 40$  trajectories. We compare our method with the following detectors:

<sup>3</sup>Code to replicate the numerical experiments can be found at <https://github.com/nzilberstein/Langevin-MIMO-detector>




Fig. 2: Performance analysis of our proposed method. (a) SER as a function of SNR for our Langevin method for noise levels  $L \in \{5, 10, 16, 20\}$ . (b) SER as a function of SNR for our Langevin method for  $M \in \{1, 5, 10, 20, 40\}$  numbers of trajectories. (c) SER as a function of SNR for different detection methods evaluated in a Kronecker correlated channel model as in (14).

MMSE detector, V-BLAST detector [23], and two learning-based, RE-MIMO [9] and OAMPNet [10], which were trained as explained in the respective papers with channels drawn from (14). The comparison is shown in Fig. 2c. The figure reveals that our proposed method markedly outperforms the other detectors. It is particularly interesting to notice that our proposed method outperforms the learning-based detectors, which have been trained with channel instances drawn from the same model as those in the testing set. Moreover, since our method does not require training, it is very flexible and presents two main advantages over the learning-based baselines: it can handle channels drawn from any distribution (promoting its application to real-world channel instances) and a varying number of users without the need of any retraining as required in, e.g., OAMPNet [10]. This is key in MIMO communications, as the number of users connected to the network is constantly changing.

## V. CONCLUSIONS

We proposed a massive MIMO detector based on an annealed version of Langevin dynamics that achieves state-of-the-art SER performance on correlated channels in large-scale systems. To include the prior information in the sampling process, we approximated the discrete prior distribution with a sequence of annealed noises that tend to concentrate around the discrete constellation symbols. Future work includes running experiments for scenarios where users are transmitting with multiple modulation schemes simultaneously, extending our method to scenarios with imperfect CSI, and leveraging the rich Langevin theory to derive theoretical guarantees.

## REFERENCES

- [1] Shaoshi Yang and Lajos Hanzo, "Fifty years of MIMO detection: The road to large-scale MIMOs," *IEEE Commun. Surveys Tut.*, vol. 17, no. 4, pp. 1941–1988, 2015.
- [2] Arogyaswami J. Paulraj, Dhananjay A. Gore, Rohit U. Nabar, and Helmut Bolcskei, "An overview of MIMO communications - a key to gigabit wireless," *Proc. IEEE*, vol. 92, no. 2, pp. 198–218, 2004.
- [3] Khaled B. Letaief, Wei Chen, Yuanming Shi, Jun Zhang, and Ying-Jun Angela Zhang, "The roadmap to 6G: AI empowered wireless networks," *IEEE Commun. Mag.*, vol. 57, no. 8, pp. 84–90, 2019.
- [4] Alberto Del Pia, Santanu S. Dey, and M. Molinaro, "Mixed-integer quadratic programming is in NP," *Mathematical Programming*, vol. 162, pp. 225–240, 2017.
- [5] John G. Proakis, *Digital Communications 5th Edition*, McGraw Hill, 2007.
- [6] Ananthanarayanan Chockalingam and Balaji Sundar Rajan, *Large MIMO Systems*, Cambridge University Press, 2014.
- [7] Charles Jeon, Ramina Ghods, Arian Maleki, and Christoph Studer, "Optimality of large MIMO detection via approximate message passing," in *IEEE Intl. Symp. on Info. Theory (ISIT)*, 2015, pp. 1227–1231.
- [8] Mehrdad Khani, Mohammad Alizadeh, Jakob Hoydis, and Phil Fleming, "Adaptive neural signal detection for massive MIMO," *IEEE Trans. Wireless Commun.*, vol. 19, no. 8, pp. 5635–5648, 2020.
- [9] Kumar Pratik, Bhaskar D. Rao, and Max Welling, "RE-MIMO: Recurrent and permutation equivariant neural MIMO detection," *IEEE Trans. Signal Process.*, vol. 69, pp. 459–473, 2021.
- [10] Hengtao He, Chao-Kai Wen, Shi Jin, and Geoffrey Ye Li, "A model-driven deep learning network for MIMO detection," in *IEEE Global Conf. Signal and Info. Process. (GlobalSIP)*, 2018, pp. 584–588.
- [11] Nicolas Zilberstein, Chris Dick, Rahman Doost-Mohammady, Ashutosh Sabharwal, and Santiago Segarra, "Robust MIMO detection using hypernetworks with learned regularizers," *arXiv preprint arXiv:2110.07053*, 2021.
- [12] Mathieu Goutay, Fayçal Ait Aoudia, and Jakob Hoydis, "Deep hypernetwork-based MIMO detection," in *IEEE Intl. Workshop on Signal Process. Adv. in Wireless Commun. (SPAWC)*, 2020, pp. 1–5.
- [13] Tanumay Datta, Nagendra A. Kumar, Ananthanarayanan Chockalingam, and Balaji S. Rajan, "A novel Monte-Carlo-sampling-based receiver for large-scale uplink multiuser MIMO systems," *IEEE Trans. Veh. Technol.*, vol. 62, no. 7, pp. 3019–3038, 2013.
- [14] Niranjan M Gowda, Sundar Krishnamurthy, and Andrey Belogolov, "Metropolis-hastings random walk along the gradient descent direction for MIMO detection," in *IEEE Intl. Conf. on Comms.*, 2021, pp. 1–7.
- [15] Yang Song and Stefano Ermon, "Generative modeling by estimating gradients of the data distribution," in *Advances in Neural Inf. Process. Syst. (NIPS)*, 2019, p. 11918–11930.
- [16] Bahjat Kwar, Gregory Vaksman, and Michael Elad, "SNIPS: Solving noisy inverse problems stochastically," *arXiv preprint arXiv:2105.14951*, 2021.
- [17] Christian Robert and George Casella, *Monte Carlo Statistical Method*, Springer, 1999.
- [18] Gareth O. Roberts and Richard L. Tweedie, "Exponential convergence of Langevin distributions and their discrete approximations," *Bernoulli*, vol. 2, pp. 341–363, 1996.
- [19] Alain Durmus and Éric Moulines, "Non-asymptotic convergence analysis for the Unadjusted Langevin Algorithm," *Ann. Appl. Probab.*, vol. 27, no. 3, pp. 1551–1587, 2017.
- [20] Max Welling and Yee Whye Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Intl. Conf. on Machine Learning (ICML)*, 2011, p. 681–688.
- [21] Bradley Efron, "Tweedie's formula and selection bias," *Journal of the American Stat. Association*, vol. 106, no. 496, pp. 1602–1614, 2011.
- [22] Sergey L. Loyka, "Channel capacity of MIMO architecture using the exponential correlation matrix," *IEEE Commun. Lett.*, vol. 5, no. 9, pp. 369–371, 2001.
- [23] Peter W. Wolniansky, Gerard J. Foschini, Glenn D. Golden, and Reinaldo A. Valenzuela, "V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel," in *Proc. URSI Intl. Symp. Signals, Syst., Electron. Conf. Proc.*, 1998, pp. 295–300.