

Achieving Near MAP Performance With an Excited Markov Chain Monte Carlo MIMO Detector

Jonathan C. Hedstrom, *Member, IEEE*, Chung Him Yuen, *Member, IEEE*, Rong-Rong Chen, and Behrouz Farhang-Boroujeny, *Senior Member, IEEE*

Abstract—We introduce a revised derivation of the bitwise Markov Chain Monte Carlo (MCMC) multiple-input multiple-output (MIMO) detector. The new approach resolves the previously reported high SNR stalling problem of MCMC without the need for hybridization with another detector method or adding heuristic temperature scaling factors. Another common problem with MCMC algorithms is the unknown convergence time making predictable fixed-length implementations problematic. When an insufficient number of iterations are used on a slowly converging example, the output log likelihood ratios can be unstable and overconfident. Therefore, we develop a method to identify rare slowly converging runs and mitigate their degrading effects on the soft-output information. This improves forward-error-correcting code performance and removes a symptomatic error floor in bit error rate plots. Next, pseudo-convergence is identified with a novel way to visualize the internal behavior of the Gibbs sampler. An effective and efficient pseudo-convergence detection and escape strategy is suggested. Finally, the new excited MCMC (X-MCMC) detector is shown to have near maximum-a-posteriori performance even with challenging, realistic, and highly-correlated channels at the maximum MIMO sizes and modulation rates supported by the 802.11ac WiFi specification, 8 × 8 MIMO 256 quadrature amplitude modulation.

Index Terms—Markov Chain Monte Carlo (MCMC), X-MCMC, pseudo-convergence, stalling, multiple-input multiple-output (MIMO) detector, wireless communication.

I. INTRODUCTION

THE use of spatial-multiplexing multiple-input multiple-output (MIMO) is increasingly being adopted in wireless protocols as higher spectral efficiency is needed to meet the capacity requirements of modern wireless networks [1], [2]. It has the potential to linearly increase spectral reuse and capacity as the number of streams increases. Compared to massive-MIMO, more moderately sized spatial-multiplexing MIMO systems have many potential benefits including higher single-device throughput in scattering environments, small scale system capabilities appropriate for home use and small cells, and immediate availability in several common wireless standards. Up to 8-stream MIMO is already defined in the

Manuscript received February 27, 2017; revised June 28, 2017 and August 25, 2017; accepted September 4, 2017. Date of publication September 19, 2017; date of current version December 8, 2017. The work of J. C. Hedstrom was supported in part by the National Science Foundation under Grant 1449033 and Grant 1632569 and in part by collaboration with Farhang Wireless Inc. The associate editor coordinating the review of this paper and approving it for publication was D. Lopez-Perez. (*Corresponding author: Jonathan C. Hedstrom*.)

The authors are with the Electrical and Computer Engineering Department, University of Utah, Salt Lake City, UT 84112 USA (e-mail: jchedstrom@gmail.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TWC.2017.2750667

802.11ac WiFi and LTE Advanced Release 10 protocols. With these large sizes, the performance and complexity scaling of the MIMO detector, which separates the mixed together transmitted streams at the receiver, is extremely important. The maximum-a-posteriori (MAP) detector has a complexity that increases exponentially with modulation order and the number of transmit streams which makes it unimplementable for all but the most trivial of cases. Zero forcing (ZF) and minimum mean square error (MMSE) detectors have low complexity but suffer from noise enhancement which degrades their performance sufficiently to negate the potential spectral efficiency benefits of MIMO.

The sphere-decoding (SD) class of MIMO detectors are well known to have near MAP performance [3]. Specifically the K-Best variations of sphere-decoding have been demonstrated in effective VLSI designs but their complexity increases quickly with the number of antennas, the number of transmitted bits per channel use, and the list size κ [4]. Therefore, the search for alternative, lower complexity methods is ongoing.

Markov Chain Monte Carlo (MCMC) has been shown to have near optimal performance at low signal-to-noise-ratio (SNR) and to have efficient hardware implementations [5], [6]. It uses a random walk through the permutations of the transmitted bit sequence to estimate the posterior probability distribution and accordingly generate soft-output information. Unlike other detectors, the MCMC detector thus far has had the undesirable behavior that its performance degrades at higher values of SNR [7]. This issue is caused by stalling at high SNR. Most of the attempted solutions to this problem can be grouped into either hybridization or temperature scaling approaches. The hybridization schemes combine a method with good, low complexity performance at high SNR with MCMC to combine the best traits of both. Examples are ZF, MMSE, and sphere-decoding initialized MCMC methods [8], [9]. The temperature scaling approaches are so called because they recognize that the probabilities generated by the Gibbs sampler are too cold, resulting in slow convergence. A linear scaling coefficient is therefore applied to artificially increase the noise temperature. This coefficient must be heuristically optimized depending on system parameters and SNR [10]–[12], though a recent derivation of a near optimal value is showing promise under some testing conditions [13].

There are three main contributions presented in this paper. First, we present an improved derivation of the MCMC detector's Gibbs sampler which resolves the high SNR stalling problems that have been previously reported. Our new derivation

uses a new and more accurate system model when a Gibbs sampler is far from the correct solution, accounting for both noise and bits-in-error. This results in what is effectively a dynamic temperature scaling coefficient. Then, using the ideas developed in this derivation, the problem of the unpredictable convergence time of the MCMC detector is explored. This leads to development of soft-output conditioning based on the quality of the list of samples generated by the random walk. This allows the detection of poor convergence and optionally extending the search with more iterations or terminating with a fixed number of iterations and decreasing the output confidence. Moderating the overconfident outputs is found to improve the performance of the forward-error-correcting (FEC) decoder. We also note that since the output conditioning is derived using the sample list, this method is applicable to other list based algorithms such as list sphere-decoding and K-Best.

The third contribution of this paper is identification and remediation of pseudo-convergence of the Gibbs sampler [14]. A novel way to visualize the Gibbs sampler behavior is used which provides insight into both pseudo-convergence and stalling. To mitigate the resulting correlated sampling, we propose a simple motion based detection strategy and a 1-bit forced change procedure which sufficiently excites the Gibbs sampler to allow new uncorrelated samples to be made.

The excited MCMC (X-MCMC) detector that is proposed in this paper is comprised of the excited Gibbs sampler, output LLR conditioning, and pseudo-convergence remediation through additional Gibbs excitation. The new X-MCMC detector achieves near Max-MAP performance even at the largest 802.11ac WiFi rates with 8×8 MIMO and 256 QAM modulation. This is demonstrated on the challenging, highly correlated WiFi TGn Model-D channel [15] with line-of-sight (LOS) conditions, which is much more difficult to attain near MAP performance than with uncorrelated i.i.d. Gaussian channels [16].

This paper is organized as follows. In Section II we present the system model and some specialized notation used throughout the paper. A review of a typical MCMC derivation along with background information is provided in Section III. The first two contributions of this paper are explained in Section IV; in Subsection IV-A the excited Gibbs sampler is presented and our proposed method of output LLR conditioning is discussed in Subsection IV-B. Our third contribution, the pseudo-convergence of Gibbs samplers and remedy to it, is discussed in Section V. Simulation results and a complexity analysis are presented in Section VI, and the conclusions of the paper are made in Section VII.

II. SYSTEM MODEL

The goal of a spatial-multiplexing MIMO transceiver is to exploit the spatial multipath of the environment to support overlapping data streams. This concept allows reuse of the spectrum and therefore has the potential to dramatically increase the data rates and capacity of wireless networks.

To perform data transmission with spatial-multiplexing, the transmitter sends independent data streams simultaneously on a set of N_t transmit antennas which are then received

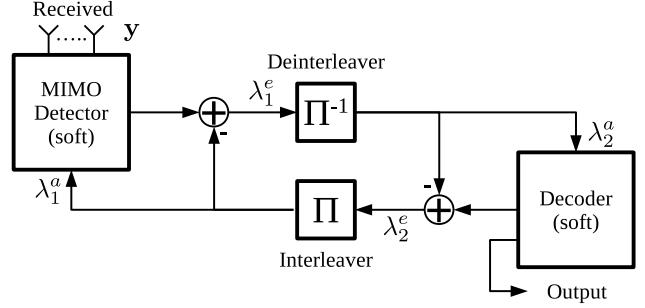


Fig. 1. Turbo loop structure with a MAP, MCMC, X-MCMC, or K-Best MIMO detector iteratively exchanging soft information with a forward error correction decoder.

at a set of N_r receive antennas. In an orthogonal frequency division multiplexing (OFDM) system, for each subcarrier, this can be represented with a flat-fading frequency-domain system model as

$$\mathbf{y} = \mathbf{Hs} + \mathbf{n}. \quad (1)$$

Here, \mathbf{y} is the received signal vector, \mathbf{H} is a slow flat fading complex channel matrix containing the pairwise gain and phase between antennas, \mathbf{s} is the vector of transmitted complex constellation symbols mapped from the bit sequence \mathbf{x} , and \mathbf{n} is a noise vector. Assuming the transmit and receive side have the same number of antennas $N_t = N_r = N$, the dimensions of the corresponding vectors and matrices are $N \times 1$ and $N \times N$, respectively. Note that often it is easier in hardware designs to use a real-valued system model which this paper is fully compatible with, see [6]. The channels for this paper are produced with the method described by the WiFi TGn Model-D LOS specification [15] which creates a correlated channel matrix \mathbf{H} . The TGn-D channel is much more challenging for the detector compared to the independent and identically distributed (i.i.d.) complex Gaussian channel typically used in the literature [3]–[6], [13]. The bit vector \mathbf{x} may be described as comprising of 1's and 0's or equivalently +1's and -1's, depending on context. In principle, \mathbf{s} can be a vector of any complex modulated symbols, but here we use the quadrature amplitude modulated (QAM) symbols defined in the IEEE 802.11ac specification. The elements of \mathbf{n} are assumed to be i.i.d. complex Gaussian random variables with variance of σ_n^2 per each real and imaginary dimension.

The primary challenge in a MIMO receiver, and the focus of this paper, is to accurately and efficiently estimate the $K = N \log_2(N_{\text{qam}})$ simultaneously transmitted bits in \mathbf{x} per realization, where N_{qam} is the QAM constellation size.

The MIMO detectors discussed in this paper are all compatible with turbo iterations, the exchange of soft-information between the detector and decoder as in Fig. 1. This allows for the iterative joint detection of the signal for enhanced performance. The soft-information is in the form of log likelihood ratios (LLRs) and represented by λ^a and λ^e for their *a priori* input and extrinsic output versions.

Notation: In the equations that follow, some specialized notation is used for compactness and clarity. Vectors and matrices are expressed with bold fonts and the latter are capitalized. The removal of the k^{th} element of a vector is shown with set

notation as $\{\cdot\}^{\setminus k}$. A variable or vector derived from the bit sequence \mathbf{x} with the k^{th} bit forced to a one or zero is shown with $\{\cdot\}^{k+}$ and $\{\cdot\}^{k-}$, respectively. When two nearly identical equations are needed differing only in use of $k+$ or $k-$, $k\pm$ is used to represent both versions. If the k^{th} bit is forced to the correct transmitted value, either one or zero, it is shown with $\{\cdot\}^{k*}$.

III. MCMC DETECTOR

The MCMC MIMO detector estimates the output log likelihood ratio (LLR) by means of Monte Carlo sampling [5]. This can be thought of as a method to identify a list of important bit permutation samples to approximate the full permutation list in the MAP detector. The challenge is to make the process computationally efficient by using an easy to calculate short list that accurately captures the statistics of the full permutation list. We select the bitwise method described in [6] as our foundational algorithm because it has been shown to be efficiently implementable in hardware.

The two main components of the MCMC detector is the Gibbs sampler and the LLR output calculation. The Gibbs sampler starts with an initial estimate of the transmitted bit sequence, either randomly selected or initialized with prior information. It then cycles through the bits, calculating the probability of a bit being a one or zero and uses it to weight a random decision to change the bit. Each cycle through the bits is an iteration. The list of permutations visited by the Gibbs sampler is used to calculate the output LLR.

To make the differences between our improved excited MCMC detector and the original bitwise MCMC detector clear, here, we will show the main steps in the original MCMC's Gibbs sampler and output LLR derivation. After a brief summary of the algorithm, we will provide some detail of the stalling problem and the MMSE initialized MCMC detector variation used in comparisons of Section VI.

A. Gibbs Sampler

The Gibbs sampler is at the core of the MCMC algorithm. It is used in difficult multi-variate posterior probability estimation problems when sampling probabilities jointly across all variables is too complex. It cycles across the variables, calculating probabilities conditioned on all other variables being fixed to the current state. In the bitwise MCMC MIMO detector, this means that the Gibbs sampler needs to calculate the probability P_{gibbs} of a specific bit x_k being a $+1$ or -1 conditioned on the current state of $\mathbf{x}^{\setminus k}$.

To derive P_{gibbs} we begin with the definition of the LLR for the k^{th} bit at the current Gibbs sampler state \mathbf{x}

$$\gamma_k = \ln \frac{P(x_k = +1|\mathbf{y}, \mathbf{x}^{\setminus k}, \lambda^a)}{P(x_k = -1|\mathbf{y}, \mathbf{x}^{\setminus k}, \lambda^a)}. \quad (2)$$

Then, by noting that

$$P(x_k = -1|\mathbf{y}, \mathbf{x}^{\setminus k}, \lambda^a) = 1 - P(x_k = +1|\mathbf{y}, \mathbf{x}^{\setminus k}, \lambda^a), \quad (3)$$

we can rearrange to have the definition of the Gibbs probability as

$$P_{\text{gibbs}} = P(x_k = +1|\mathbf{y}, \mathbf{x}^{\setminus k}, \lambda^a) = \frac{1}{1 + e^{-\gamma_k}}. \quad (4)$$

Next, Bayes' Theorem is applied to γ_k to separate the contribution of the prior λ^a . The result is

$$\begin{aligned} \gamma_k &= \ln \frac{p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k}, \lambda^a)P(x_k = +1|\mathbf{x}^{\setminus k}, \lambda^a)p(\mathbf{y})}{p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k}, \lambda^a)P(x_k = -1|\mathbf{x}^{\setminus k}, \lambda^a)p(\mathbf{y})} \\ &= \ln \frac{p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k})P(x_k = +1|\lambda_k^a)}{p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k})P(x_k = -1|\lambda_k^a)} \\ &= \ln \frac{p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k})}{p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k})} + \lambda_k^a \end{aligned} \quad (5)$$

where in the second line λ^a , $\mathbf{x}^{\setminus k}$, and $\lambda^{a, \setminus k}$ are removed from their respective conditionals because of the independence among the \mathbf{x} bits created by the interleaving effect in the turbo loop. The final line follows from the definition of the *a priori* LLR λ^a .

The system model is now used to provide the probability of the received sequence \mathbf{y} given \mathbf{x} and noise with variance σ_n^2 per dimension. This leads to

$$p(\mathbf{y}|\mathbf{x}) = (2\pi\sigma_n^2)^{-N} \exp\left(-\frac{\|\mathbf{y} - \mathbf{Hs}\|^2}{2\sigma_n^2}\right). \quad (6)$$

Substituting (6) into (5) and simplifying the result, we get

$$\gamma_k = \frac{1}{2\sigma_n^2} \left(\|\mathbf{y} - \mathbf{Hs}^{k-}\|^2 - \|\mathbf{y} - \mathbf{Hs}^{k+}\|^2 \right) + \lambda_k^a. \quad (7)$$

This may be used with (4) to calculate the needed P_{gibbs} probability for the Gibbs sampler. Note that this will only be an accurate calculation if AWGN noise with variance σ_n^2 is the only contributor to the error residual in $\mathbf{y} - \mathbf{Hs}$. This assumption is revisited in the derivation of the X-MCMC detector in Section IV where we find that error in the Gibbs sampler's current state of \mathbf{s} introduces additional error that should be accounted for.

B. Output LLR

To calculate the extrinsic output LLR λ^e we first derive the MAP output LLR. The definition of LLR for the k^{th} bit is the log ratio of probabilities of x_k being a $+1$ or -1 . This is expressed as

$$\begin{aligned} \lambda_{k,\text{MAP}}^e &= \ln \frac{P(x_k = +1|\mathbf{y}, \lambda^{a, \setminus k})}{P(x_k = -1|\mathbf{y}, \lambda^{a, \setminus k})} \\ &= \ln \frac{\sum_{\mathbf{x} \in \mathbb{X}^{k+}} \left(p(\mathbf{y}|\mathbf{x}) \prod_{j \neq k} P(x_j|\lambda_j^a) \right)}{\sum_{\mathbf{x} \in \mathbb{X}^{k-}} \left(p(\mathbf{y}|\mathbf{x}) \prod_{j \neq k} P(x_j|\lambda_j^a) \right)} \end{aligned} \quad (8)$$

where \mathbb{X}^{k+} and \mathbb{X}^{k-} are the sets of all permutations of \mathbf{x} with the k^{th} bit forced to a $+1$ and -1 , respectively. To simplify (8), we recall from [17] that

$$P(x_j|\lambda_j^a) = \frac{e^{-\lambda_j^a/2}}{1 + e^{-\lambda_j^a}} e^{x_j \lambda_j^a/2} = A(\lambda_j^a) e^{x_j \lambda_j^a/2}. \quad (9)$$

Since the coefficient $A(\lambda_j^a)$ is independent of x_j , it may be separated from the summation and canceled. This leads to

$$\lambda_{k,\text{MAP}}^e = \ln \frac{\sum_{\mathbf{x} \in \mathbb{X}^{k+}} \left(p(\mathbf{y}|\mathbf{x}) \exp \sum_{j \neq k} \frac{1}{2} x_j \lambda_j^a \right)}{\sum_{\mathbf{x} \in \mathbb{X}^k} \left(p(\mathbf{y}|\mathbf{x}) \exp \sum_{j \neq k} \frac{1}{2} x_j \lambda_j^a \right)}. \quad (10)$$

Finally, the max-log approximation to the Jacobian logarithm [18] is used to simplify (10) to

$$\begin{aligned} \lambda_{k,\text{Max-MAP}}^e &\approx \frac{1}{2} \max_{\mathbf{x} \in \mathbb{X}^{k+}} \left(-\frac{1}{\sigma_n^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}^{k+}\|^2 + \mathbf{x}^k \cdot \lambda^{a,\setminus k} \right) \\ &\quad - \frac{1}{2} \max_{\mathbf{x} \in \mathbb{X}^k} \left(-\frac{1}{\sigma_n^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}^k\|^2 + \mathbf{x}^k \cdot \lambda^{a,\setminus k} \right) \end{aligned} \quad (11)$$

where $\mathbf{x}^k \cdot \lambda^{a,\setminus k}$ is a vector dot product. This approximation is commonly used by approximate-MAP MIMO detectors as it results in a minor loss of performance while significantly decreasing algorithm complexity [3].

To calculate the output LLR with MCMC, the max-log MAP calculation of (11) is used but with the set \mathbb{X} replaced with the list \mathbb{Z} of the sampled permutations. This, with some minor rearrangements leads to

$$\begin{aligned} \lambda_{k,\text{MCMC}}^e &\approx \frac{1}{2} \min_{\mathbf{x} \in \mathbb{Z}^k} \left(\frac{1}{\sigma_n^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}^k\|^2 - \mathbf{x}^k \cdot \lambda^{a,\setminus k} \right) \\ &\quad - \frac{1}{2} \min_{\mathbf{x} \in \mathbb{Z}^{k+}} \left(\frac{1}{\sigma_n^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}^{k+}\|^2 - \mathbf{x}^{k+} \cdot \lambda^{a,\setminus k} \right) \end{aligned} \quad (12)$$

where the $\max()$ operation has been changed to an equivalent $\min()$ operation to put the equation in a more intuitive cost function minimization form.

C. Summary

The MCMC MIMO detector algorithm starts with an initial sequence of bits \mathbf{x} . It then cycles across the K bits in a bitwise fashion for N_{iter} iterations, using the probabilities calculated with (4) and (7) to determine state transitions. This process effectively guides the Gibbs sampler towards the more important regions. After a sufficient quantity of samples have been taken in this manner, the algorithm approximates the max-log MAP detector's extrinsic output LLR λ_k^e in (11) with (12). This procedure is outlined in Algorithm 1 for one Gibbs sampler, though N_{gibbs} can be used in parallel to increase sampling speed and increase sample diversity.

D. Stalling Problem

At high SNR, the MCMC detector derived here is known to stall. This means that the rate at which new samples are generated per iteration approaches zero [19]. This results in slow convergence and an enormous number of iterations needed to reach near MAP performance, potentially more than the calculations needed for MAP itself. To help understand stalling more thoroughly and the improvements proposed in

Algorithm 1 Basic MCMC Gibbs Sampler

```

1: Initialize  $\mathbf{x}$ .
2: for  $1$  to  $N_{\text{iter}}$  do
3:   for  $k = 1$  to  $N \log_2(N_{\text{qam}})$  do
4:     Update  $\min_{\mathbf{x} \in \mathbb{Z}^{k+}}(\cdot)$  and  $\min_{\mathbf{x} \in \mathbb{Z}^{k-}}(\cdot)$  for (12).
5:     Calculate  $P_{\text{gibbs}}$  with (4) and (7).
6:     Generate a uniform random variable  $0 \leq r \leq 1$ .
7:     if  $r < P_{\text{gibbs}}$  then  $x_k \leftarrow +1$  else  $x_k \leftarrow -1$ 
8: Compute output LLR with (12).

```

the subsequent sections, a detailed view of the underlying Gibbs samplers is presented that will lead to some insight.

In Fig. 2, we see the internal behavior of a pair of randomly initialized Gibbs samplers working on a 4×4 MIMO system with a 16 QAM symbol constellation, thus $4 \log_2(16) = 16$ bits per \mathbf{x} . The only difference between the subfigures is the magnitude of AWGN noise, 6 dB vs 12 dB E_b/N_0 , which is sufficient to strongly instigate the high SNR stalling problem. Each subplot has the bit index k over the horizontal axis from left to right, and the Gibbs iterations starting at the top and descending over time. Accordingly, the Gibbs sampler moves from left to right and top to bottom.

The left subplots show the probability of deterministic and non-deterministic flips where

$$\text{determinism} = |2P_{\text{gibbs}} - 1|. \quad (13)$$

With a gray-scale color mapping the determinism is shown as black = 1 (fully deterministic) and white = 0 (fully random).

The right subplots show the error in the Gibbs state \mathbf{x} when compared to the true transmitted bit sequence, where black indicates a bit error. State error is used instead of simply \mathbf{x} because it simultaneously shows if the state is changing and where the random walk is relative to the transmitted bit sequence.

In Fig. 2, we see the curious trend that at low SNR the Gibbs sampler probability is moderate and the sampler quickly converges to having only 1-bit error, whereas at high SNR the Gibbs probabilities are persistently at extreme values and no change is seen after the first iteration.

Stalling is not seen under all conditions. If a smaller QAM size of 2 to 16 is used at a low SNR, near the BER cliff, stalling will not be observed [20], [21]. Another condition where the full extent of stalling may not be obviously observed is when using well-conditioned, Rayleigh channels with uncorrelated i.i.d. Gaussian channel matrix elements [16]. Real-world channels, with correlation between antennas, display stalling much more strongly than the uncorrelated channels often used in the literature to simulate MCMC performance [4], [6], [7], [13]. There are many variations of the MCMC detector which attempt to mitigate the high SNR stalling problem. Most can be broadly classified as either hybrid or temperature scaling methods.

The idea behind the hybrid approaches is to combine the MCMC detector, which has decreasing performance as SNR increases, with a complementary detector, which has low performance at low SNR but high performance at high SNR.

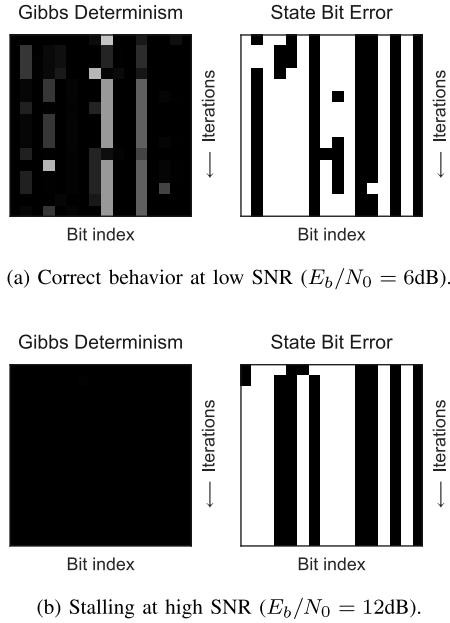


Fig. 2. Gibbs sampler detail with the original randomly initialized MCMC algorithm. Both subfigures have the exact same bit sequence, channel, seeded random number generator, and additive noise. The only difference is in noise magnitude which instigates the high SNR stalling problem. Parameters: 4 antennas, 16 QAM, $N_{\text{iter}} = 16$, WiFi TGn Model-D LOS channel.

There are many examples of this using the sphere-decoder such as [8] and [9] and with MMSE [7]. An interesting variation on this theme is to use successive-over-relaxation within each Gibbs sampler step calculation [22]. We note that although these methods can improve the performance compared to a standard MCMC sampler by itself, this class of detectors do not solve the high SNR stalling problem of the MCMC algorithm they use.

The temperature scaling class of MCMC detectors recognize that the superficial cause of stalling is that the P_{gibbs} values calculated are too extreme. They suggest that, as the SNR increases, σ_n^2 acts as a gain term in (7), therefore they heuristically add an additional scaling term α which either linearly scales or replaces σ_n^2 to counteract stalling [10], [13], [23]. A variation on this idea includes methods which use manually selected parameters to adjust the inclusion of additional randomization into the probability calculation or bit selection [20], [24]. Stalling will be further explained in the framework of the X-MCMC detector variation of Section IV-A.

E. MMSE Initialized MCMC

To better understand the improvements made by the X-MCMC detector that will be introduced in Section IV, an MCMC method is needed for comparison. We have selected the MMSE initialized MCMC detector described in [7]. This is a hybrid type of detector since it combines MMSE, which has good performance at high SNR, with MCMC to compensate for the high SNR stalling effect. The reason why we have selected this as apposed to some of the other excellent QRD-hybrid or temperature scaling varieties in the literature is

that those methods are generally loosely defined with parameters that need to be heuristically tuned to the application. The MMSE initialized variety is useful as a benchmark because it has an explicit implementation regardless of channel, SNR, number of antennas, or modulation order.

The only extension needed for the MMSE initialized MCMC hybrid algorithm is to initialize one of the parallel Gibbs samplers with the hard decision from an MMSE solution as in

$$\hat{\mathbf{s}}_{\text{MMSE}} = (\mathbf{H}^\dagger \mathbf{H} + 2\sigma_n^2 \mathbf{I})^{-1} \mathbf{H}^\dagger \mathbf{y} \quad (14)$$

where \dagger is the conjugate transpose. This method solves the issue where the MCMC detector does not converge for small QAM sizes at high SNR. It works because the MMSE solution will result in a correctly signed LLR output without any need of additional MCMC Gibbs iterations. For turbo iterations to work properly and to generate reasonable extrinsic information transfer (EXIT) charts, every iteration has one parallel Gibbs sampler initialized with the original MMSE solution.

As described in Section III-D, the MMSE initialized MCMC detector does not prevent stalling because it is a hybrid method, thus it is far from MAP performance when higher QAM sizes are used. Examples that show such behavior are presented in Section VI.

IV. X-MCMC DETECTOR

In this section, we present the two main components of the excited MCMC (X-MCMC) detector: the excited Gibbs sampler and LLR output conditioning. Both rely on the realization that there are error contributions from both noise and bits-in-error. Previous derivations have only included AWGN noise in their derivations. We found that the inclusion of error in the statistical model completely solves the high SNR stalling problem covered in Section III-D. It also provides a method to detect and mitigate poorly converged and overconfident output LLRs that otherwise confuse the decoder, potentially creating error floors in BER plots.

Previously in [16], we used a heuristic explanation to justify the dynamic scaling used in the X-MCMC detector. Simulations were matched to 8×8 MIMO testbed measurements with near max-log MAP performance. Those ideas will be expanded on here with a thorough theoretical understanding and some algorithmic improvements.

In the discussions and derivations that follow, the concept of distance is repeatedly used. It is the closeness of the current state \mathbf{x} to the transmitted sequence. To simplify its use, it will be defined as the square Euclidean distance

$$d = \|\mathbf{y} - \mathbf{Hs}\|^2 \quad (15)$$

where \mathbf{s} is the complex symbol mapped version of the bit state \mathbf{x} . If the k^{th} bit of \mathbf{x} is forced to a one or zero, then this can be indicated on all dependent variables and vectors with a superscript $k+$ or $k-$ as in

$$d^{k\pm} = \|\mathbf{y} - \mathbf{Hs}^{k\pm}\|^2. \quad (16)$$

A. Excited Gibbs Sampler

The cause of Gibbs sampler stalling in the MCMC detector is constant production of extreme P_{gibbs} values close to 0% or 100%. This creates a nearly deterministic walk which quickly stalls. Extreme probabilities are readily caused by (4) and (7). At just $|\gamma_k| = 3$ there is a mere 2% chance of a non-deterministic transition. In simulations of 4-antenna 64 QAM systems, γ_k values greater than 10 are commonly seen and values greater than 100 are not unusual at high SNRs.

In the MCMC derivation of Section III, an assumption is implicitly made that the only form of error is from the channel model's AWGN noise. It is a useful assumption because it allows the substitution of a Gaussian probability distribution into (5) that simplifies to the explicit γ_k of (7). However, this causes problems since the incorrect bits in \mathbf{x} also contribute to error and in general are much larger contributors, often by a factor of 1000 at high SNRs. This provides insight as to why the stalling problem only appears at high SNR, since at low SNRs the condition is less strongly violated. Therefore, the goal here is to find a way to calculate the needed probabilities in γ_k without this assumption.

We begin the derivation of the excited Gibbs sampler with (5) from the original MCMC derivation.

$$\gamma_k = \ln \frac{p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k})}{p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k})} + \lambda_k^a \quad (17)$$

The implicit statistical model used for the probability distributions in the original MCMC derivation is

$$\mathbf{n} = \mathbf{y} - \mathbf{Hs}. \quad (18)$$

But a more accurate model is needed when the Gibbs sampler is far from the true transmitted bit sequence. The error \mathbf{e}^{k*} must include both AWGN interference and the incorrect bits in the state. This is captured with the new model

$$\mathbf{e}^{k*} = \mathbf{y} - \mathbf{Hs}^{k*} = \mathbf{H}(\mathbf{s}_{\text{tx}} - \mathbf{s}^{k*}) + \mathbf{n} \quad (19)$$

where \mathbf{s}_{tx} is the transmitted symbol vector and \mathbf{s}^{k*} is the current Gibbs state with the k^{th} bit of \mathbf{x} correct, i.e. matches the respective transmitted bit. This model does not include error from the k^{th} bit in \mathbf{s} so that the distributions of the numerator and denominator of (17) will be the same.

It is not generally possible to separate the contributions of x_k and $\mathbf{x}^{\setminus k}$ to the error, so in the steps that follow we will initially assume that an oracle has access to such an error metric which excludes the k^{th} bit and finish with approximations sufficient for development of a detector that works well in practice.

To use the new error model, we first note that a Gaussian distribution is an appropriate approximation for the elements of \mathbf{e}^{k*} . Excluding the contribution of noise, each element is generated from the dot product of a row of the known \mathbf{H} and a column of unknown symbol errors. Therefore, each real and imaginary component of \mathbf{e}^{k*} is created from the sum of $4N$ uniform random variables, where the number of antennas has been multiplied by two for the complex nature of the symbols and two for the contribution of the independent s_{tx} and s^{k*} symbols. For the smallest MIMO size, this results in an $n = 8$ Irwin-Hall distribution which is already approximately

Gaussian and which becomes more Gaussian as the MIMO size increases. Though our QAM symbols take discrete values, noise and \mathbf{H} shift and scale the symbols sufficiently to retain the Gaussian approximation.

Here, we have assumed that \mathbf{s}_{tx} and \mathbf{s}^{k*} are uncorrelated which is valid when far from convergence. As the Gibbs sampler converges and the symbols become increasingly correlated, the true distribution becomes narrower than a Gaussian, thus, a Gaussian assumption results in underconfident probabilities near convergence. In Section V, we will find that pseudo-convergence is a stronger effect leading to overconfident probabilities during convergence. Therefore, the slight underconfidence here can be safely ignored.

Using a Gaussian distribution for the elements of \mathbf{e}^{k*} with variance σ_{k*}^2 per dimension results in

$$p(\mathbf{y}|x_k = \pm 1, \mathbf{x}^{\setminus k}) = (2\pi\sigma_{k*}^2)^{-N} \exp\left(-\frac{\|\mathbf{y} - \mathbf{Hs}^{k\pm}\|^2}{2\sigma_{k*}^2}\right). \quad (20)$$

Then, by substituting this Gaussian model into (17), we have

$$\gamma_k = \frac{d^{k-} - d^{k+}}{2\sigma_{k*}^2} + \lambda_k^a. \quad (21)$$

A single sample estimate of the error variance σ_{k*}^2 can be made with the current error vector \mathbf{e}^{k*} as in

$$\sigma_{k*}^2 \approx \frac{\|\mathbf{e}^{k*}\|^2}{2N} = \frac{d^{k*}}{2N} \quad (22)$$

where we have noted that the squared norm of the error vector is equivalent to having the distance d with the k^{th} bit known to be correct, $d^{k*} = \|\mathbf{y} - \mathbf{Hs}^{k*}\|^2$. We refer to this estimation of the error variance as the *oracle* method since it uses knowledge that is not generally available to the algorithm. Next, we will develop several potential approximations to the *oracle*.

We begin by using the confidence that either d^{k+} or d^{k-} has the k^{th} bit correct to generate a weighted average estimate \hat{d}^{k*} .

$$\begin{aligned} \hat{d}^{k*} &= d^{k+}P(x_k = +1|\mathbf{x}^{\setminus k}, \mathbf{y}) + d^{k-}P(x_k = -1|\mathbf{x}^{\setminus k}, \mathbf{y}) \\ &= \frac{d^{k+}p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k})P(x_k = +1|\mathbf{x}^{\setminus k})}{p(\mathbf{y}|\mathbf{x}^{\setminus k})} \\ &\quad + \frac{d^{k-}p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k})P(x_k = -1|\mathbf{x}^{\setminus k})}{p(\mathbf{y}|\mathbf{x}^{\setminus k})} \\ &= \frac{d^{k+}p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k}) + d^{k-}p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k})}{2p(\mathbf{y}|\mathbf{x}^{\setminus k})} \end{aligned} \quad (23)$$

where Bayes' rule is first applied and then $P(x_k = \pm 1|\mathbf{x}^{\setminus k}) = 1/2$ because of the independence between bits. This estimate can be bounded by the minimum and mean of $d^{\pm k}$ since $P(x_k = -1|\mathbf{x}^{\setminus k}, \mathbf{y}) = 1 - P(x_k = +1|\mathbf{x}^{\setminus k}, \mathbf{y})$ and the larger probability will correspond to the smaller distance.

$$\hat{d}_{\min}^{k*} = \min(d^{k+}, d^{k-}) \leq \hat{d}^{k*} \quad (24)$$

$$\hat{d}^{k*} \leq \frac{1}{2}(d^{k+} + d^{k-}) = \hat{d}_{\text{mean}}^{k*} \quad (25)$$

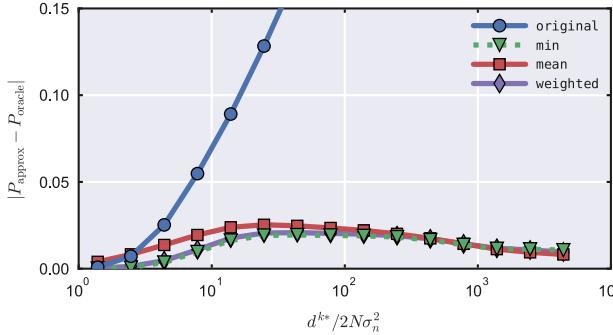


Fig. 3. Mean error from using various d^{k*} approximations to calculate P_{gibbs} with (4), (21), and (22). Parameters: 4 antennas, 64 QAM, $N_{\text{gibbs}} \times N_{\text{iter}} = 30 \times 30$, WiFi TGn Model-D LOS channel, $E_b/N_0 = 19\text{dB}$.

These bounds are useful rough approximations to \hat{d}^{k*} as they are readily available to the Gibbs sampler and are computationally efficient. We will refer to these as the *min* and *mean* approximations.

A more accurate approximation to the weighted average of (23) can be made by first replacing the denominator, which would require the unknown \mathbf{x}^{k*} , with the average of the probabilities with $x_k = \pm 1$.

$$p(\mathbf{y}|\mathbf{x}^{\setminus k}) = \frac{1}{2} \left(p(\mathbf{y}|x_k = +1, \mathbf{x}^{\setminus k}) + p(\mathbf{y}|x_k = -1, \mathbf{x}^{\setminus k}) \right) \quad (26)$$

Then, single sample estimates of the variances $\sigma_{k\pm}^2 \approx d^{k\pm}/2N$ are used to approximate the Gaussian distributions

$$\begin{aligned} p(\mathbf{y}|x_k = \pm 1, \mathbf{x}^{\setminus k}) &\approx \left(2\pi \sigma_{k\pm}^2 \right)^{-N} \exp \left(-\frac{d^{k\pm}}{2\sigma_{k\pm}^2} \right) \\ &\approx \left(\pi \frac{d^{k\pm}}{N} \right)^{-N} \exp \left(-\frac{d^{k\pm}}{d^{k\pm}/N} \right). \end{aligned} \quad (27)$$

This estimate is too crude to be used directly in (17), but by using (26) and (27), an approximation of (23) can be made

$$\begin{aligned} \hat{d}^{k*}_{\text{weighted}} &= \frac{d^{k+}(d^{k+})^{-N} + d^{k-}(d^{k-})^{-N}}{(d^{k+})^{-N} + (d^{k-})^{-N}} \\ &= \frac{d^{k+} + d^{k-}(d^{k+}/d^{k-})^N}{1 + (d^{k+}/d^{k-})^N}. \end{aligned} \quad (28)$$

For comparison, it is useful to include the original MCMC method in Section III in this framework. It implicitly uses the approximation

$$\hat{d}^{k*}_{\text{original}} = 2N\sigma_n^2. \quad (29)$$

The results of using the *min* (24), *mean* (25), *weighted* (28), and *original* (29) estimates of d^{k*} to generate P_{gibbs} are shown in Fig. 3, where the error is the mean of $|P_{\text{approx}} - P_{\text{oracle}}|$ over a small range of distances. Based on these plots, the new approximations are all a vast improvement over the *original* method. At large distances, all of the new approximations have similar performance since a single bit change has little impact on the total distance, and therefore, $d^{k+} \approx d^{k-}$ resulting in all of the approximations being roughly equivalent. At small distances, the *mean* method has the worst performance since it overweights the larger of $d^{k\pm}$, making the sampler too random.

Since the *min* and *weighted* approximations have similar performance, the less computationally complex *min* option is selected for σ_{k*}^2 in calculating

$$\gamma_k = \frac{d^{k-} - d^{k+}}{\hat{d}^{k*}_{\text{min}}/N} + \lambda_k^a. \quad (30)$$

The choice of $\hat{d}^{k*}_{\text{min}}$ will be verified in Section VI-A with BER plots in a complete system, but first, output LLR conditioning and pseudo-convergence enhancements will be introduced so that the combined effects can be accounted for.

B. Output LLR Overconfidence and Conditioning

As in the original MCMC algorithm in Section III, the output LLR can be calculated with (12). It uses the list $Z \subset \mathbb{X}$ of sampled bit permutations as an approximation to the full permutation list used by the MAP detector in (11). This is the method used by list based MIMO detectors such as MCMC [5], list sphere-decoding [3], and K-Best [4]. Notice that in order to use this method, one must have a list which accurately represents the statistics of the full permutation list. If the list is too short or poorly selected the statistics break down, the calculation becomes unreliable, and the output LLRs can be wildly inaccurate.

One can state that a requirement of the algorithm is for a sufficient number of samples to be taken, but there are two reasons why this is undesirable in practice. First, an implementable design has the goal of minimizing computational complexity and therefore real-world implementations need to use the minimum number of samples possible. As a result, a minority of realizations are likely to be poorly converged, leading to invalid output statistics. Second, many modern communication systems have multiple channel realizations per codeword, for example OFDM. When there are multiple channel realizations, some will be more ill-conditioned than others with a longer convergence time and therefore requiring more Gibbs iterations than the average. During our analysis we have observed situations where the worst channel, often during a deep fade, requires 10x more iterations than average while using a WiFi TGn Model-D LOS channel model, 4 antennas, 64 QAM. Using 10x more iterations consistently is undesirable, but if the slowly converging realizations are halted early while still statistically unstable, then their large incorrect soft-output values can easily corrupt the entire codeword despite comprising a small number of bit errors.

In the context of suboptimal iterative detection and decoding schemes, a constant positive scaling coefficient less than one is sometimes used to scale the extrinsic LLRs [25]–[27]. This has the effect of removing divergence and thus increase stability. Therefore, an alternative to using a large list to deal with outlier difficult channel realizations is to decrease the confidence of their respective LLRs so that they no longer have the strength to corrupt the entire codeword. Unlike other scaling methods which suggest a heuristic constant scaling coefficient to reduce LLR confidence, we propose using the statistics of the sample list to compute a dynamic coefficient that may be used with any list based detector.

If we allow that the sample list may be of poor quality, then there are two contributions to error: the AWGN noise in the system model and the quality of the list. The quality of the list can be computed using the probability that the list contains a sample that is the true transmitted signal without error. As in Section IV-A, if we assume that the combination of both forms of error are Gaussian distributions then the output LLR can be changed to

$$\begin{aligned} \lambda_{k, \text{X-MCMC}}^e &\approx \frac{1}{2} \min_{\mathbf{x} \in \mathbf{Z}^{k-}} \left(\frac{1}{\sigma_z^2} \left\| \mathbf{y} - \mathbf{Hs}^{k-} \right\|^2 - \mathbf{x}^{\setminus k} \cdot \boldsymbol{\lambda}^{a, \setminus k} \right) \\ &\quad - \frac{1}{2} \min_{\mathbf{x} \in \mathbf{Z}^{k+}} \left(\frac{1}{\sigma_z^2} \left\| \mathbf{y} - \mathbf{Hs}^{k+} \right\|^2 - \mathbf{x}^{\setminus k} \cdot \boldsymbol{\lambda}^{a, \setminus k} \right) \end{aligned} \quad (31)$$

where σ_z^2 represents the variance of the combined contributions of both AWGN noise and list error. This can be estimated with the minimum distance sample in the list similarly to (22) using \hat{d}_{\min}^{k*} . That is,

$$\sigma_z^2 \approx \frac{1}{2N} \min_{\mathbf{x} \in \mathbf{Z}} \left(\left\| \mathbf{y} - \mathbf{Hs} \right\|^2 \right). \quad (32)$$

Note that additionally we should limit $\sigma_z^2 \geq \sigma_n^2$. This protects against a sample over-fitting the noise and therefore having a distance less than the expectation of correct bit sequences, $2N\sigma_n^2$, which would incorrectly create overconfidence in the output LLR.

The corruption of codewords by rare poorly converged realizations can cause an error floor to appear on BER plots. Using output LLR conditioning reduces this effect and allows a short, fixed number of iterations to be used stably and reliably in real-world applications.

We also note that since the output conditioning is derived using the sample list, this method may also be applicable to other list based algorithms such as list sphere-decoding and K-Best. Likewise, extreme output LLR values in X-MCMC may be saturated to more moderate values similar to what is sometimes done in sphere-decoding [3], though we have found the dynamic output conditioning method suggested here to be preferable to the ad hoc saturation approach.

C. Summary

In practice, the X-MCMC algorithm is executed similarly to the original MCMC algorithm outlined in Algorithm 1. The changes include the use of the newly introduced dynamic scaling of $(d^{k-} - d^{k+})$ and output conditioning applied to $\lambda_{k, \text{X-MCMC}}^e$. The full X-MCMC algorithm is outlined in Algorithm 2 with the necessary equations specified. Note that many Gibbs samplers may be used in parallel to increase sampling speed and increase sample diversity. When using parallel Gibbs samplers each \hat{d}_{\min}^{k*} should be calculated independently for each Gibbs sampler whereas the output LLR should be calculated with a list \mathbf{Z} comprising a combination of all samplers. Moreover, σ_z^2 in (32) should be obtained by taking the minimum of \hat{d}_{\min}^{k*} of all the Gibbs samplers.

V. PSEUDO-CONVERGENCE

Now that the high SNR stalling problem is solved by the X-MCMC detector and the Gibbs sampler moves efficiently at

Algorithm 2 X-MCMC Gibbs Sampler

```

1: Initialize  $\mathbf{x}$ .
2: for  $1$  to  $N_{\text{iter}}$  do
3:   for  $k = 1$  to  $N \log_2(N_{\text{qam}})$  do
4:     Calculate  $d^{k+}$  and  $d^{k-}$  with (16).
5:     Update  $\min_{\mathbf{x} \in \mathbf{Z}}()$  for (32).
6:     Update  $\min_{\mathbf{x} \in \mathbf{Z}^{k+}}()$  and  $\min_{\mathbf{x} \in \mathbf{Z}^{k-}}()$  for (31).
7:     Calculate  $P_{\text{gibbs}}$  with (4) and (30).
8:     Generate a uniform random variable  $0 \leq r \leq 1$ .
9:     if  $r < P_{\text{gibbs}}$  then  $x_k \leftarrow +1$  else  $x_k \leftarrow -1$ 
10: Compute output LLR with (31), (32).

```

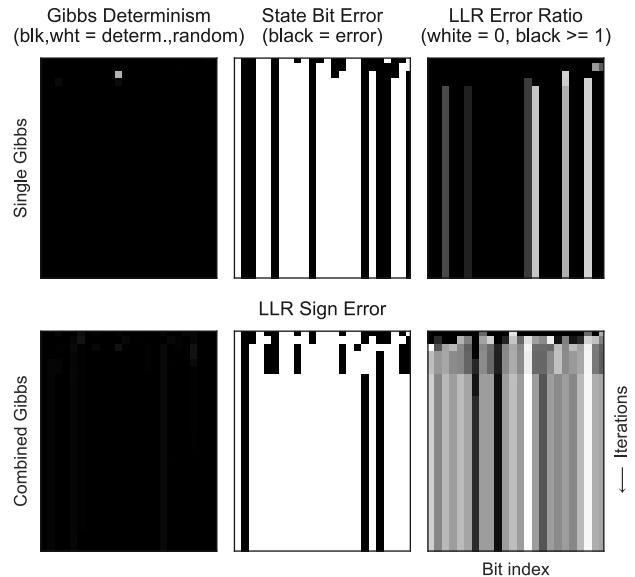


Fig. 4. Detail on Gibbs sampler behavior with the original random initialized MCMC. Notice the high SNR stalling problem resulting in no improvement after the first few iterations. Parameters: 4 antennas, 64 QAM, $N_{\text{gibbs}} \times N_{\text{iter}} = 30 \times 30$, WiFi TGn Model-D LOS channel, $E_b/N_0 = 19\text{dB}$.

all SNR levels, a new issue is encountered. The Gibbs sampler may stop moving due to an effect referred to as pseudo-convergence. This problem appears to be less understood in MIMO communications applications, but has been noted in the wider MCMC statistics literature [14]. Although symptomatically similar, it is different from the stalling problem discussed previously. It occurs when the posterior probability distribution is multi-modal with weak connections between modes, that is, important regions are weakly connected to other important regions through very low probabilities. Thus, the Gibbs sampler may stay in one mode (i.e. similar bit permutations connected with sufficient probability) and collect highly correlated samples. When encountered, pseudo-convergence decreases the sampling efficiency of MCMC, resulting in a need for a large number of iterations and/or parallel Gibbs samplers. Others appear to be encountering this effect as reports of performance improvements when adding random walk restarting have been observed [20], [23].

A. Gibbs Detail Plots

To understand the pseudo-convergence phenomena more thoroughly, we have developed the Gibbs detail plots of Fig. 4, 5, and 7, introduced partially in Section III-D.

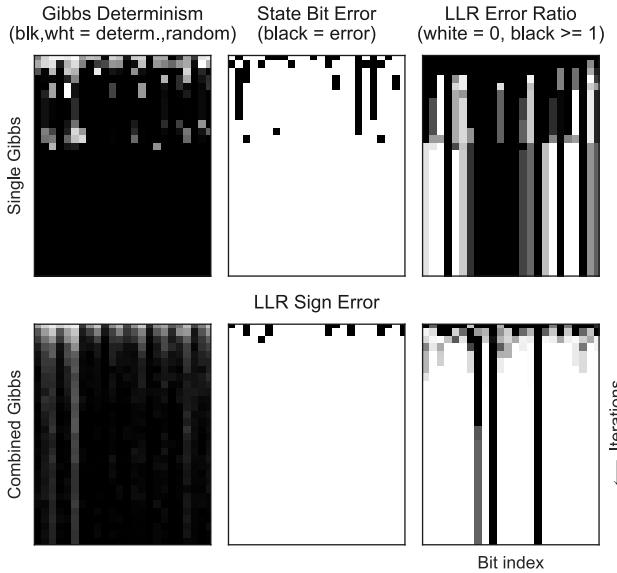


Fig. 5. Detail on Gibbs sampler behavior with partial X-MCMC ('x - -' = Gibbs excitement only). Stalling fixed but now there is a pseudo-convergence stopping issue. Parameters: 4 antennas, 64 QAM, $N_{\text{gibbs}} \times N_{\text{iter}} = 30 \times 30$, WiFi TGn Model-D LOS channel, $E_b/N_0 = 19\text{dB}$.

For consistency, these figures all use the same input data and initialized states. The first row of subplots show information on a single Gibbs sampler whereas the second row shows information on the combination of all parallel Gibbs samplers of the algorithm run on a single realization of transmitted data.

$$\text{LLR error ratio} = |\lambda_k^e - \lambda_{k,\text{MAP}}^e| / \text{mean}_k (|\lambda_{k,\text{MAP}}^e|) \quad (33)$$

is calculated against the optimal max-log MAP solution. The max-log MAP solution is what MCMC should converge to if run for an infinite number of iterations. In Fig. 4, 5, and 7, the LLR error ratio values are shown with a gray-scale color mapping, zero to one as white to black where values above one saturate to black.

The second row shows a combined view of all of the N_{gibbs} parallel Gibbs samplers. The first subplot is the average ‘‘determinism’’, whereas the second and third are the LLR sign error relative to Max-MAP and the LLR error ratio in (33).

The desired behavior of the first column is to show signs of the random walk being guided with a variation of determinism, not fully random or deterministic. In the second column, the single Gibbs sampler should not stay converged to any state and instead should continuously explore the state space, whereas the combined Gibbs samplers should converge to no LLR sign error relative to Max-MAP. Finally, the third column should continue to converge to the Max-MAP solution, displayed by white.

Fig. 4 shows that the original MCMC detector described in Section III is almost completely deterministic, is strongly stalled at this high SNR, and does not improve the output LLR after only a few iterations.

By using the excited Gibbs sampler described in Section IV-A, Fig. 5 shows a large improvement in behavior. It is no longer stalled and after a few Gibbs iterations the algorithm has mostly converged. In the second row of subplots,

after an initial very active period, the guided random walk slows as all parallel samplers become locked into an isolated posterior mode due to pseudo-convergence.

B. Detection and Escape

As pseudo-convergence is a byproduct of the structure of the posterior distribution, the algorithm is behaving correctly and as expected. Given enough iterations a single Gibbs sampler will eventually leave an isolated mode and sample others. Instead of waiting for the low probability of transition, we prefer a more efficient method of detecting when pseudo-convergence has occurred and then forcing state divergence. This allows the algorithm to collect more unique samples with fewer iterations, thus improving the sampling efficiency of the MCMC detector.

Two effective and computationally efficient methods to detect pseudo-convergence include what we refer to as the *distance* and *motion* methods. The *distance* method tracks the best (smallest) distance d sampled over time, including both d^{k+} and d^{k-} . If this distance does not improve in N_{motion} steps, then pseudo-convergence is detected. Alternatively, the *motion* method detects when no change has occurred in Gibbs state x for N_{motion} steps.

The choice of using the \hat{d}_{\min}^{k*} estimate in Section IV-A causes the Gibbs sampler to move slightly more slowly and deterministically, thus we have found that it tends to stop moving when in pseudo-convergence, therefore using the *motion* pseudo-convergence detection strategy works well with the choice of \hat{d}_{\min}^{k*} . For the detection threshold we use $N_{\text{motion}} = N \log_2(N_{\text{qam}})$ steps which is one full Gibbs iteration.

Once pseudo-convergence is detected, the most robust though not necessarily the best solution is to restart the Gibbs sampler with a new fully random state. This is a solution mentioned in the wider MCMC literature beyond MIMO communications applications [14]. A drawback of the full restart approach is that it requires re-initialization of the Gibbs sampler which may be an expensive and time consuming operation in VLSI implementations. We have found that a full restart is not necessary in a bitwise MCMC detector and under some circumstances is a bad choice. Instead, forcing a 1-bit state change in the next bit following pseudo-convergence detection can be a good solution. This can be thought of as adding an impulse of energy or excitation to the random walk which assists the Gibbs sampler in escaping the isolated posterior mode. Using this 1-bit strategy incurs no additional complexity from re-initializing the sampler and is trivial to implement in VLSI designs.

In Fig. 6 which uses *motion* based detection, we see that the full-restart escape method can actually degrade performance. This is due to the number of iterations needed to converge to a region of important samples being larger than the number needed to naturally leave a pseudo-converged region, thus, the full-restart is lowering overall sampling efficiency. For these test parameters, the full-restart creates too large of a divergence which wastes time in re-converging, whereas the next-bit forced change is sufficient to leave the pseudo-converged region without incurring a re-convergence penalty.

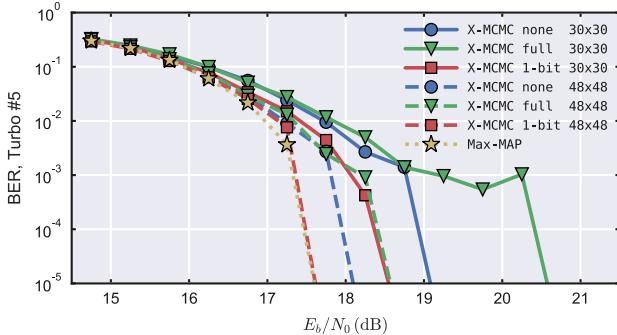


Fig. 6. Comparison of using no pseudo-convergence mitigation, full-restart escape, and next-bit forced change escape. Parameters: 4 antennas, 64 QAM.

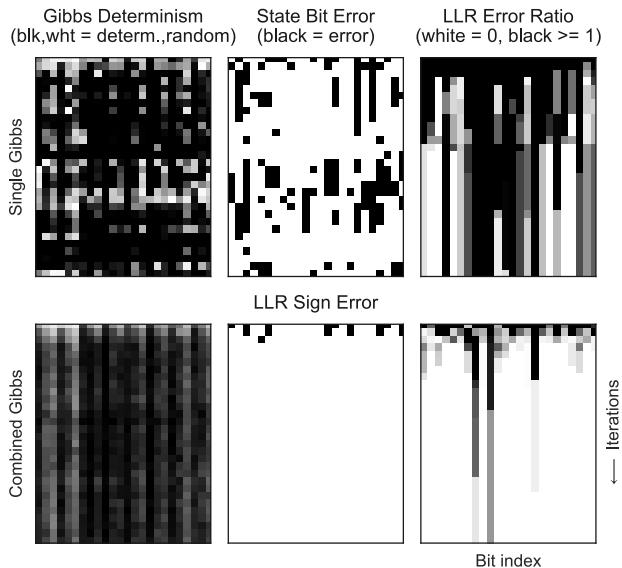


Fig. 7. Detail on Gibbs sampler behavior with partial X-MCMC ('x - p' = Gibbs excitement and pseudo-convergence stopping mitigation). Stalling and stopping are now resolved. Parameters: 4 antennas, 64 QAM, $N_{\text{gibbs}} \times N_{\text{iter}} = 30 \times 30$, WiFi TGn Model-D LOS channel, $E_b/N_0 = 19\text{dB}$.

By adding the *motion* based pseudo-convergence detection and next-bit forced change method to the excited Gibbs sampler we see the results presented in Fig. 7. Now the stalling problem seen in Fig. 4 and the stopping problem seen in Fig. 5 are resolved. Both the single and combined Gibbs determinism subplots show that the MCMC algorithm is consistently excited. The combined LLR quickly converges to a correct output bit sequence and then continues to improve the output LLR until near Max-MAP performance is achieved.

VI. RESULTS

One of the most important features of the results that follow is that the WiFi TGn Model-D LOS channel [15] has been used for \mathbf{H} . This realistic, correlated indoor model creates a much more challenging problem compared to the typical Gaussian i.i.d. channels commonly used in the literature [3]–[6], [13]. The main issue with using an unrealistically uncorrelated, and therefore unrealistically easy model, is that convergence is relatively easy to achieve, making stalling and stopping issues appear less severe. For more details on channel model selection, see the analysis and testbed results in [16].

We have found that producing BER curves using only the hard-decision of the MIMO detector by itself is insufficient to analyze MCMC performance. Since the MCMC detector can potentially have statistical stability issues, assessing the quality of the soft-output information is essential, therefore, we have used the LDPC 3/4 rate 1944 block length code from the 802.11n specification in generation of all of the following BER curves. It was selected as a moderate coding level among the 1/2, 2/3, 3/4, and 5/6 rates available in the WiFi specification.

For most of the BER results, five turbo loops are used, see Fig. 1, which enhances final BER performance by allowing the detector and decoder to iteratively exchange extrinsic soft-information [28]. Although not all applications have sufficient time in their latency budget to do turbo iterations, it is useful to do most of the analysis using these iterations because it allows for the testing of an algorithm's ability to use prior information correctly.

The figures in this section include an MMSE initialized MCMC detector for reference as described in Section III-D. The randomly initialized version is not generally included as it performs worse than the MMSE initialized version under all conditions. Max-MAP/Max-ML is shown as the optimal performance bound when possible since the channel capacity for non-Gaussian models is generally unknown. By using a highly optimized GPU implementation we are able to compute MAP/ML at up to 4 antennas with 64 QAM. For the 8 antenna with 256 QAM case we use a very large K-Best as an approximation of the Max-MAP limit since it is known to have near-MAP performance [4]. One moderate sized K-Best is generally included so that the reader may do some initial comparisons with the literature on K-Best.

A. d^{k*} Approximations

In the previous sections, the X-MCMC algorithm was developed with greatly enhanced performance. Now with all three components fully developed, we may revisit the selection of \hat{d}_{\min}^{k*} made in Section IV-A. The reason for doing this verification after the development of the other X-MCMC enhancements is that their combined interaction can impact the final choice. Therefore, in Fig. 8, we show BER plots comparing the suggested approximations to d^{k*} while also using the output LLR conditioning and pseudo-convergence enhancements. Both the first and final turbo iterations need to be shown to check for issues with the σ_{k*}^2 scaling relative to the prior λ^a in (30).

The BER curves show that the new approximations perform well and will approach near MAP performance given sufficient iterations, though not with equal speed or efficiency. As was also shown in the P_{gibbs} error analysis of Fig. 3, we find that the *min* approximation performs well. It remains the preferred approximation since it also has a much lower computational complexity compared to *weighted*.

B. EXIT Chart

To deepen the understanding of the MCMC and X-MCMC detectors, an extrinsic information transfer (EXIT) chart is presented in Fig. 9 with prior information randomly generated

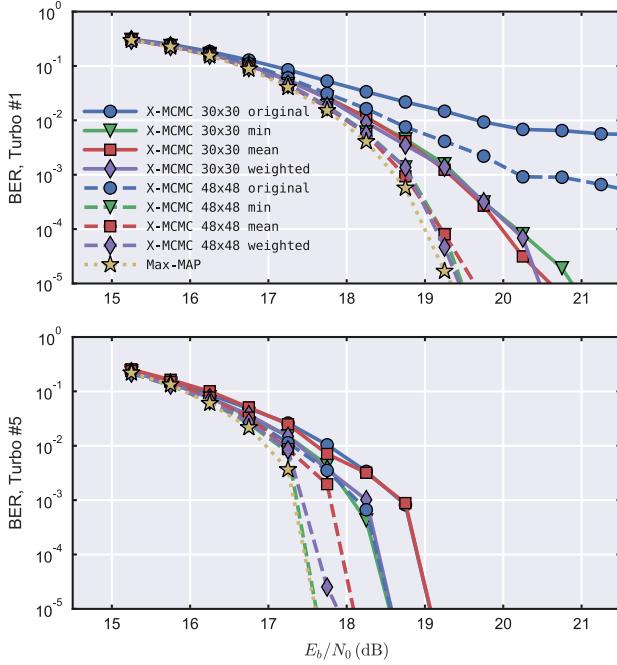


Fig. 8. Comparison of d^{k*} approximations (29), (24), (25), and (28) with all X-MCMC components enabled including Gibbs excitation, output LLR conditioning, and pseudo-convergence mitigation. Two plots are shown, the top without turbo iterations (Turbo#1) and with five turbo iterations (Turbo#5). Both are important as they show the effects without and with use of *a priori* information. Parameters: 4 antennas, 64 QAM.

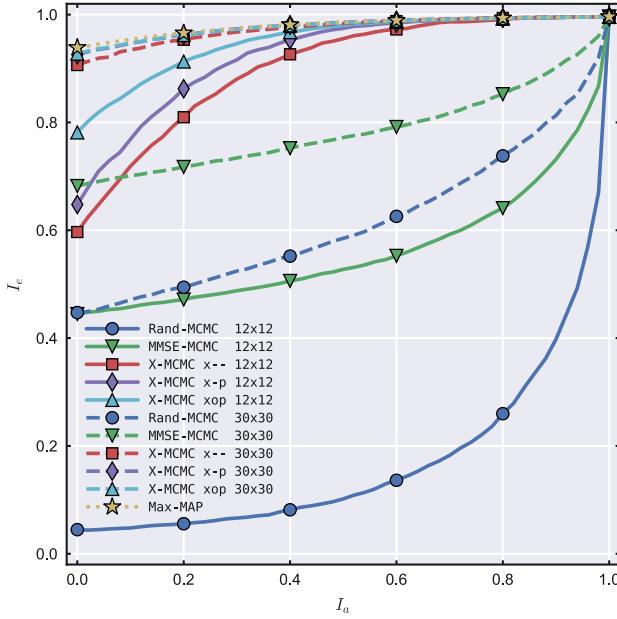


Fig. 9. EXIT chart showing random and MMSE initialized original MCMC methods versus X-MCMC components ('x o p' are flags representing inclusion of Gibbs excitation, output LLR conditioning, and pseudo-convergence mitigation). Parameters: 4 antennas, 64 QAM, $E_b/N_0 = 19$ dB, no coding.

as described in [29]. The EXIT chart is useful in evaluating MCMC performance since it is independent of code choice and shows both the ability of the detector to use input soft-information and generate output soft-information [30], [31]. Given an amount of input information I_a it shows how

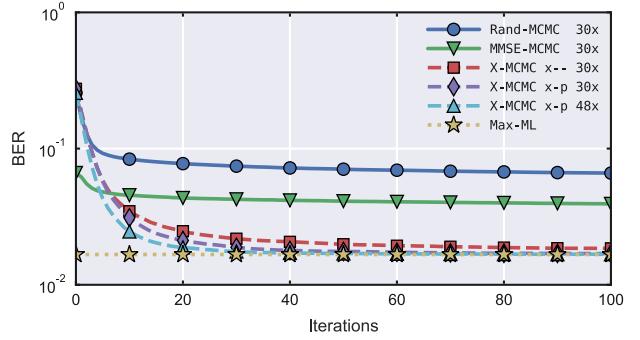


Fig. 10. Convergence analysis showing that X-MCMC converges quickly to the ML solution which is placed as a reference line, though it has no iterations. The number of parallel Gibbs samplers is specified in the legend as '#x'. For X-MCMC, 'x - p' are flags representing inclusion of Gibbs excitation and pseudo-convergence mitigation. Parameters: 4 antennas, 64 QAM, $E_b/N_0 = 19$ dB, no coding.

much extrinsic output information I_e a given method is able to produce, therefore, EXIT charts provide complementary analysis to detector-only BER curves which do not show the quality of LLR output.

As expected, the output extrinsic information I_e of the X-MCMC detector improves as the excited Gibbs sampler, output LLR conditioning, and pseudo-convergence enhancements are included in the algorithm. The introduction of the excited Gibbs sampler is the most important contribution as it fixes the unusual EXIT curve shapes presented by the random and MMSE initialized original MCMC methods, caused by high SNR stalling. This unusual shape has also been observed in [11]. The prolonged, flat shape with a sharp rise at the end of the curves is produced by the large d^{k*} underestimate from using (29). This means that the $(d^{k-} - d^{k+})/(\hat{d}_{\text{original}}^{k*}/2N)$ is far overweighted compared to the prior λ^a in (21). Once $I_a > 0.8$ the prior becomes strong enough to overcome the imbalance at this SNR. Because an MMSE initialization only provides the algorithm with a good starting point, it does not change the underlying problem.

C. Convergence

One of X-MCMC's primary advantages is that it has a fast convergence rate because it has no high SNR stalling problem and avoids pseudo-convergence. This can be seen in Fig. 10 where the hard-decision BER of several detectors are compared as the number of Gibbs iterations increases. The convergence time will generally become longer as the number of antennas increase, the constellation size becomes larger, and the channel becomes more correlated. Note that the output LLR conditioning of X-MCMC is not shown in the figure because scaling does not change the LLR sign and therefore has no effect on the detector BER.

D. BER Performance

The BER curves of Fig. 11 confirm the relationships shown in the EXIT chart of Fig. 9. There is an incremental improvement in performance as each of the Gibbs excitation, pseudo-convergence enhancement, and output LLR conditioning are

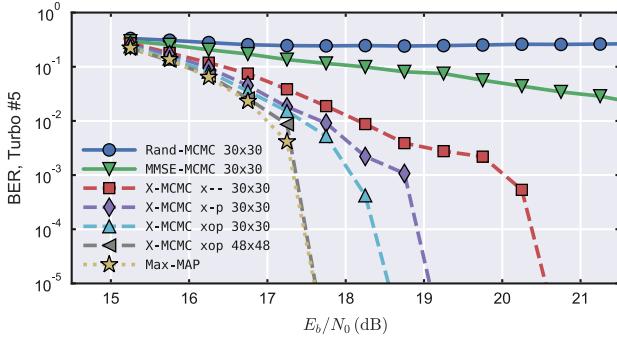


Fig. 11. BER curves comparing MCMC methods versus X-MCMC ('x o p' are flags representing inclusion of Gibbs excitation, output LLR conditioning, and pseudo-convergence mitigation). Parameters: 4 antennas, 64 QAM.

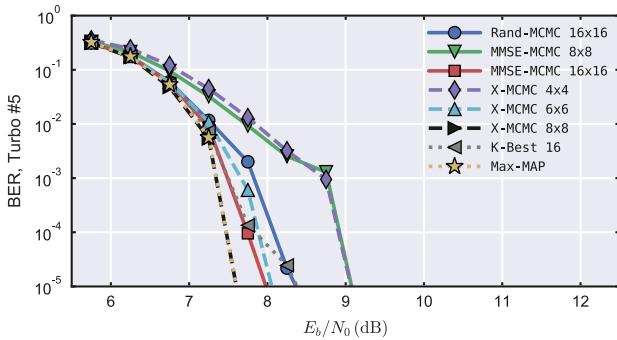


Fig. 12. BER curves showing that all methods function at low SNR with small constellation sizes, though X-MCMC is more efficient than previous MCMC methods. Parameters: 4 antennas, 4 QAM.

included. As predicted by the EXIT chart, a 48×48 X-MCMC detector achieves near Max-MAP performance, and a smaller 30×30 detector is within 1dB. The most interesting feature of these curves is the error floor seen in the X-MCMC curves without output LLR conditioning. This is caused by rare realizations with slow convergence that poorly converge with the fixed number of Gibbs iterations provided. The LLR overconfidence in the poorly converged cases are capable of corrupting entire codewords even when representing a small minority of realizations. For more details on how output LLR conditioning resolves this effect see Subsection IV-B.

The most important observation in the remaining BER figures is that the X-MCMC detector is capable of achieving near Max-MAP performance under all conditions tested. This is especially impressive at the maximum 802.11ac WiFi protocol size of 8 antenna MIMO with 256 QAM modulation shown in Fig. 14. Compared to MMSE-Initialized MCMC there is a massive >6 dB improvement. Similar results are seen in Fig. 13 with X-MCMC again achieving Max-MAP performance.

It is relatively easy to achieve near Max-MAP performance on low-order modulation with low-SNR, as seen in Fig. 12 and reported in [7] and [32]. Though MMSE-MCMC works under these conditions, it is at a lower efficiency than X-MCMC. This is predicted by our excited Gibbs derivation since the poor approximation $\hat{d}_{\text{original}}^{k*}$ in (29) becomes more accurate at lower SNRs and therefore the impact of stalling is limited.

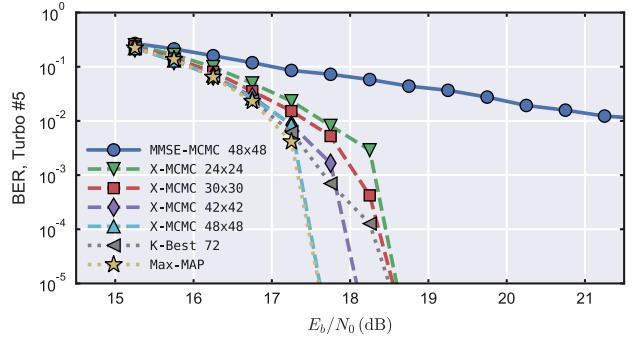


Fig. 13. BER curves showing near Max-MAP performance for X-MCMC. Parameters: 4 antennas, 64 QAM.

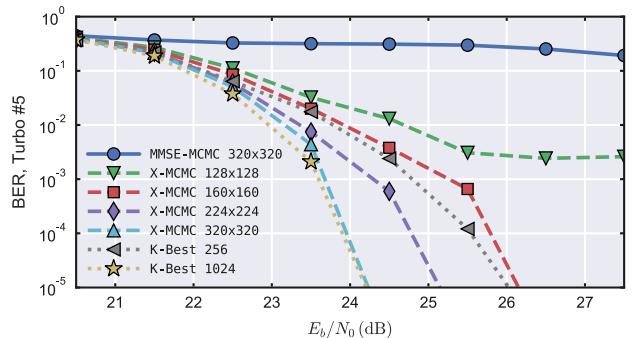


Fig. 14. BER curves showing X-MCMC achieves near Max-MAP performance (approximated with large K-Best) even at the maximum 802.11ac MIMO and QAM sizes. Parameters: 8 antennas, 256 QAM.

E. Complexity Analysis

To quantify complexity we will use the total number of multiplications as a metric. One can include more complex operations in the analysis, such as exponential functions and division by assuming a roughly equivalent number of multiplies [4], [33]. For the inverted exponential function used in the probability calculation of (4), in [6] it is shown that a lookup table (LUT) and adder can provide sufficient accuracy, therefore, its complexity is sufficiently less than a multiplier to be ignored. For the scaling operations of (30) and (31), we will temporarily set the division equal to C_{div} multiplies.

Note that using the number of multiplies as a complexity metric may be useful to roughly compare similar algorithms, but it is insufficient to compare X-MCMC and K-Best in particular. This is because a major contributor to K-Best complexity is list-sorting [4] which cannot be measured with a multiplication count. Therefore, a gate count from a VLSI design of X-MCMC will be needed to make a full complexity comparison with other MIMO detector techniques.

The proposed X-MCMC algorithm can be broken into four components for complexity analysis: frontend processing (FEP), backend processing (BEP), total number of bitwise steps including all parallel Gibbs samplers, and each bitwise step.

The FEP is used to calculate precomputed values, including $\mathbf{H}^\dagger \mathbf{H}$ and $\mathbf{H}^\dagger \mathbf{y}$, which simplifies each Gibbs sampler step [6]. Similarly, the BEP includes the calculation of the output

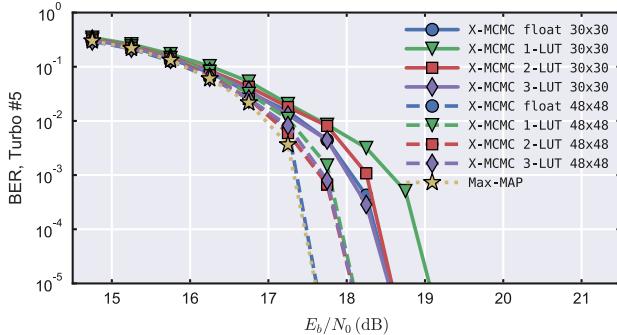


Fig. 15. Comparison of using full floating-point precision division vs a simple shift and LUT to replace the scaling operations in (30) and (31). Note that the 'X-MCMC float 30x30' curve is underneath the corresponding 3-LUT curve. Parameters: 4 antennas, 64 QAM.

LLR (31). Thus the FEP and BEP complexity can be represented as $C_{\text{FEP}} = (2N)^3 + (2N)^2$ and $C_{\text{BEP}} = C_{\text{div}} + NM$, where N is the number of antennas, and M is the number of bits in each complex valued matrix \mathbf{H} can be represented by a real version with dimension $2N \times 2N$ and likewise for \mathbf{y} with dimension $2N \times 1$. The BEP includes a single reciprocal operation to compute the scaling factor which is assumed to be equivalent to a division and then it is applied NM times.

The Gibbs sampler includes the total number of Gibbs sampler bitwise steps multiplied by the complexity of each step as in

$$C_{\text{gibbs}} = (NM(2NM)^2) \times (2N + 1 + C_{\text{div}}) \quad (34)$$

where we assume that there are NM bitwise steps per iteration, $2NM$ iterations per Gibbs sampler, and $2NM$ parallel Gibbs samplers are used. This number of iterations and parallel samplers is sufficient to achieve near-MAP performance, as shown in Fig. 13, depending on the channel and coding. Using the implementation of [6] as a foundation for the X-MCMC enhancements, the step calculation is dominated by a dot product of length $2N$ and the division introduced by our new excitation scaling factor σ_{k*}^2 in (30). It is assumed that the simple state machine necessary to implement the pseudo-convergence strategy of Section V is of negligible complexity as it requires no arithmetic operations.

It is clear that the division operation introduced to scale the calculation of γ in (30) can potentially have a large impact on the leading complexity terms $(8N^4M^3 + 4N^3M^3C_{\text{div}})$, therefore, we have explored a simple estimation of the division using a shift unit and n-bit LUT using the most significant bits of the scaling term. In Fig. 15, it is shown that a 3-bit LUT is sufficient to nearly match the floating-point division performance. This is not surprising as the scaling terms themselves are only rough approximations of the local and final error variances.

Since using a rough division approximation is acceptable, we conclude that a division complexity of $C_{\text{div}} \approx 1$ is appropriate. After this simplification, the X-MCMC algorithm adds little additional complexity to each Gibbs step compared to a standard MCMC detector while also significantly reducing

the number of needed steps. Thus, the algorithm complexity appears to be polynomial in N and scalable to large MIMO sizes.

VII. CONCLUSION

We have presented a new derivation of the MCMC detector which solves the high SNR stalling problem without use of hybridization or heuristic temperature scaling terms. Output LLR quality has been improved for poorly converged cases by conditioning output confidence on sample list statistics. Output LLR conditioning is shown to moderate soft-output overconfidence and allow a low complexity fixed length Gibbs sampler to be used in practice, eliminating error floors caused by rare slowly converging realizations. This conditioning may have application to other list based detectors such as list sphere-decoding and K-Best. Additionally, we have identified pseudo-convergence conditions which lower efficiency. The proposed 1-bit randomization procedure is shown as a low complexity alternative way to leave pseudo-convergence compared to using a full random-walk restart. Results show that the combined improvements allow near Max-MAP performance at all SNR regimes with large numbers of antennas and high-order modulation. This is true even with highly correlated, WiFi TGn Model-D LOS channels which are significantly more challenging than the Rayleigh channels with uncorrelated i.i.d. Gaussian elements commonly used in the literature. No heuristic optimizations are needed, making the proposed method straightforward to effectively implement in practice. A brief complexity analysis demonstrated that the X-MCMC enhancements requires little additional complexity compared to previous MCMC detectors while dramatically improving performance. A VLSI implementations should be possible with straightforward extensions of existing work.

MIMO sizes beyond our verified 8×8 sizes should be possible within the constraints of the suggested polynomial complexity growth of the X-MCMC algorithm. More correlated channels than the WiFi TGn Model-D LOS channels are possible [16], though with somewhat longer convergence times requiring more Gibbs iterations.

REFERENCES

- [1] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "802.11 with multiple antennas for dummies," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 19–25, Jan. 2010.
- [2] F. Boccardi *et al.*, "Multiple-antenna techniques in LTE-advanced," *IEEE Commun. Mag.*, vol. 50, no. 3, pp. 114–121, Mar. 2012.
- [3] B. M. Hochwald and S. ten Brink, "Achieving near-capacity on a multiple-antenna channel," *IEEE Trans. Commun.*, vol. 51, no. 3, pp. 389–399, Mar. 2003.
- [4] Z. Guo and P. Nilsson, "Algorithm and implementation of the K-Best sphere decoding for MIMO detection," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 3, pp. 491–503, Mar. 2006.
- [5] B. Farhang-Boroujeny, H. Zhu, and Z. Shi, "Markov chain Monte Carlo algorithms for CDMA and MIMO communication systems," *IEEE Trans. Signal Process.*, vol. 54, no. 5, pp. 1896–1909, May 2006.
- [6] S. A. Laraway and B. Farhang-Boroujeny, "Implementation of a Markov chain Monte Carlo based multiuser/MIMO detector," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 1, pp. 246–255, Jan. 2009.
- [7] X. Mao, P. Amini, and B. Farhang-Boroujeny, "Markov chain Monte Carlo MIMO detection methods for high signal-to-noise ratio regimes," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Nov. 2007, pp. 3979–3983.

[8] R. Peng, K. H. Teo, J. Zhang, and R.-R. Chen, "Low-complexity hybrid QRD-MCMC MIMO detection," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Nov. 2008, pp. 1–5.

[9] F.-L. Yuan, C.-H. Yang, and D. Marković, "A hardware-efficient VLSI architecture for hybrid sphere-MCMC detection," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Dec. 2011, pp. 1–6.

[10] M. Hansen, B. Hassibi, A. G. Dimakis, and W. Xu, "Near-optimal detection in MIMO systems using gibbs sampling," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Nov. 2009, pp. 1–6.

[11] M. Senst and G. Ascheid, "A Rao–Blackwellized Markov chain Monte Carlo algorithm for efficient MIMO detection," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2011, pp. 1–6.

[12] D. Auras, U. Deidersen, R. Leupers, and G. Ascheid, "VLSI design of a parallel MCMC-based MIMO detector with multiplier-free gibbs samplers," in *Proc. IEEE Int. Conf. Very Large Scale Integr. (VLSI-SoC)*, Oct. 2014, pp. 1–6.

[13] B. Hassibi, M. Hansen, A. G. Dimakis, H. A. J. Alshamary, and W. Xu, "Optimized Markov chain Monte Carlo for signal detection in MIMO systems: An analysis of the stationary distribution and mixing time," *IEEE Trans. Signal Process.*, vol. 62, no. 17, pp. 4436–4450, Sep. 2014.

[14] S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, *Handbook of Markov Chain Monte Carlo*. Boca Raton, FL, USA.: Chapman & Hall, 2011.

[15] V. Erceg *et al.*, *IEEE P802.11 Wireless LANs—Indoor MIMO WLAN TGn Channel Models*, document IEEE 802.11-03/940r4, 2004.

[16] J. C. Hedstrom, C. H. G. Yuen, R.-R. Chen, and B. Farhang-Boroujeny, "Excited Markov chain Monte Carlo MIMO detector with 8-antenna 802.11ac testbed demonstration," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2017, pp. 1–7.

[17] J. Hagenauer, E. Offer, and L. Papke, "Iterative decoding of binary block and convolutional codes," *IEEE Trans. Inf. Theory*, vol. 42, no. 2, pp. 429–445, Mar. 1996.

[18] P. Robertson, E. Villebrun, and P. Hoeher, "A comparison of optimal and sub-optimal MAP decoding algorithms operating in the log domain," in *Proc. IEEE Int. Conf. Commun. (ICC)*, vol. 2, Jun. 1995, pp. 1009–1013.

[19] H. Zhu, Z. Shi, and B. Farhang-Boroujeny, "MIMO detection using Markov chain Monte Carlo techniques for near-capacity performance," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, vol. 3, Mar. 2005, pp. III-1017–III-1020.

[20] A. Kumar, S. Chandrasekaran, A. Chockalingam, and B. S. Rajan, "Near-optimal large-MIMO detection using randomized MCMC and randomized search algorithms," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2011, pp. 1–5.

[21] Z. Shi, H. Zhu, and B. Farhang-Boroujeny, "Markov chain Monte Carlo techniques in iterative detectors: A novel approach based on Monte Carlo integration," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, vol. 1, Nov. 2004, pp. 325–329.

[22] J. Choi, "An MCMC–MIMO detector as a stochastic linear system solver using successive overrelaxation," *IEEE Trans. Wireless Commun.*, vol. 15, no. 2, pp. 1445–1455, Feb. 2016.

[23] T. Datta, N. A. Kumar, A. Chockalingam, and B. S. Rajan, "A novel MCMC algorithm for near-optimal detection in large-scale uplink multiuser MIMO systems," in *Proc. IEEE Inf. Theory Appl. Workshop (ITA)*, Feb. 2012, pp. 69–77.

[24] S. He, J. Zhou, J. Hu, and J. Chen, "A low complexity MCMC algorithm for MIMO system with bias technique," in *Proc. IEEE Int. Conf. ASIC (ASICON)*, Nov. 2015, pp. 1–4.

[25] C. Douillard *et al.*, "Iterative correction of intersymbol interference: Turbo-equalization," *Eur. Trans. Telecommun.*, vol. 6, no. 5, pp. 507–511, 1995.

[26] A. Elkhazin, K. N. Plataniotis, and S. Pasupathy, "Reduced-dimension MAP turbo-BLAST detection," *IEEE Trans. Commun.*, vol. 54, no. 1, pp. 108–118, Jan. 2006.

[27] G. Colavolpe, D. Fertonani, and A. Piemontese, "SISO detection over linear channels with linear complexity in the number of interferers," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 8, pp. 1475–1485, Dec. 2011.

[28] P. Robertson, P. Hoeher, and E. Villebrun, "Optimal and sub-optimal maximum *a posteriori* algorithms suitable for turbo decoding," *Eur. Trans. Telecommun.*, vol. 8, no. 2, pp. 119–125, 1997.

[29] S. ten Brink, "Convergence behavior of iteratively decoded parallel concatenated codes," *IEEE Trans. Commun.*, vol. 49, no. 10, pp. 1727–1737, Oct. 2001.

[30] R.-R. Chen, R. Peng, A. Ashikhmin, and B. Farhang-Boroujeny, "Approaching MIMO capacity using bitwise Markov chain Monte Carlo detection," *IEEE Trans. Commun.*, vol. 58, no. 2, pp. 423–428, Feb. 2010.

[31] M. Senst, G. Ascheid, and H. Lüders, "Performance evaluation of the Markov chain Monte Carlo MIMO detector based on mutual information," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2010, pp. 1–6.

[32] J. C. Hedstrom, C. H. Yuen, and B. Farhang-Boroujeny, "Markov chain Monte Carlo based multiuser/MIMO detector: 802.11ac implementation and measurement," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 4846–4852.

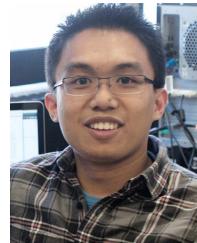
[33] K. Amiri, P. Radosavljevic, and J. R. Cavallaro, "Architecture and algorithm for a stochastic soft-output MIMO detector," in *Proc. IEEE Asilomar Conf. Signals, Syst. Comput. (ACSSC)*, Nov. 2007, pp. 1034–1038.



Jonathan C. Hedstrom (M'17) received the B.S. degree in electrical engineering, physics, and mathematics from the University of Wisconsin–Madison, Madison, WI, USA, in 2008, the M.S. degree in electrical engineering from Boston University, Boston, MA, USA, in 2011, and the Ph.D. degree in electrical engineering from the University of Utah, Salt Lake City, UT, USA, in 2017.

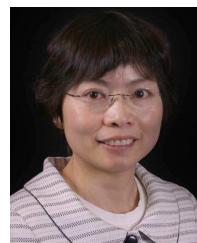
From 2008 to 2013, he was a Technical Lead and Project Manager with the MIT Lincoln Laboratory, with a focus on rapid prototyping of RF systems.

Since 2015, he has been a Technical Lead and the CEO with Farhang Wireless Inc. His current research interests include real-world signal processing applications and using statistical techniques to solve complex problems, with an emphasis on wireless MIMO detectors.



Chung Him (George) Yuen (M'12) was born in Hong Kong in 1984. He received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Utah, Salt Lake City, UT, USA, in 2007, 2011, and 2015, respectively. During his graduate studies, he was a Graduate Research Assistant with the Wireless Communications Laboratory, under the supervision of Ph.D. Advisor Prof. B. Farhang-Boroujeny. His research focused on wireless communications signal processing with emphasis on OFDM, filter bank multicarrier communication techniques, power reduction methods in SC-FDMA uplink of LTE, and time scaling compensation in Doppler scaling channels.

Since 2015, he has been a Senior Research Engineer with Farhang Wireless Inc., where he has been involved in next generation wireless communication technologies. His research interests include signal processing, multiple input multiple output detection, and general areas of wireless communication techniques.



Rong-Rong Chen is currently an Associate Professor with the Department of Electrical and Computer Engineering, University of Utah. Her main research interest include communication systems and networks, with current emphasis on statistical signal processing, caching networks, channel coding, multiple antenna communications, and underwater acoustic communications. She has also been involved in signal processing for medical imaging. She was a recipient of the M. E. Van Valkenburg Graduate Research Award for excellence in doctoral research in the ECE Department, University of Illinois at Urbana–Champaign, and the prestigious National Science Foundation Faculty Early Career Development Award in 2006. She was rated among the top 15% instructors of the College of Engineering, University of Utah, in 2017, 2011, 2010, 2006, and 2005. She has served as an Associate Editor of the *IEEE TRANSACTIONS ON SIGNAL PROCESSING* and as a Guest Editor of *IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING*. She has served on the technical program committee of leading international conferences in wireless communication and networks.



Behrouz Farhang-Boroujeny (M'84–SM'00) received the B.Sc. degree in electrical engineering from the University of Tehran, Tehran, Iran, in 1976, the M.Eng. degree from the Institute of Science and Technology, University of Wales, Cardiff, U.K., in 1977, and the Ph.D. degree from Imperial College London, University of London, London, U.K., in 1981. From 1981 to 1989, he was with the Isfahan University of Technology, Isfahan, Iran.

From 1989 to 2000, he was with the National University of Singapore, Singapore. Since 2000, he has been with the University of Utah, Salt Lake City, UT, USA. He is an expert in signal processing. He has authored the books *Adaptive Filters: Theory and*

Applications (Wiley, 1998) and *Signal Processing Techniques for Software Radios* (Lulu, 2009 and 2010, second edition). In the past, he was involved with and made significant contributions to the areas of adaptive filters theory, acoustic echo cancellation, magnetic/optical recording, and digital subscriber line technologies. His current research interests include adaptive filters, multicarrier communications, detection techniques for space-time-coded systems, and cognitive radio.

Dr. Farhang-Boroujeny received the UNESCO Regional Office of Science and Technology for South and Central Asia Young Scientists Award in 1987. He has been involved in various IEEE activities, including the chairmanship of the Signal Processing/Communications Chapter of the IEEE of Utah in 2004 and 2005. He served as an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING from 2002 to 2005 and the IEEE SIGNAL PROCESSING LETTERS from 2008 to 2010.