Phase noise analysis for mmwave massive MIMO: a design framework for scaling via tiled architectures

Maryam Eslami Rasekh*, Mohammed Abdelghany†, Upamanyu Madhow‡, Mark Rodwell §
Department of Electrical and Computer Engineering
University of California Santa Barbara
Email: {*rasekh, †mabdelghany, ‡madhow, §rodwell}@ucsb.edu

***Abstract*—We consider a tiled architecture for scaling a millimeter wave (mmWave) massive MIMO uplink to support a large number of simultaneous users, targeting per-user data rates of multi-Gbps. Our goal in this paper is to evaluate the impact of phase noise, which is widely considered to be a significant bottleneck at high carrier frequencies and large bandwidths, on MIMO performance. Our analysis provides a cross-layer design framework which can be employed by hardware designers to determine allowable masks for the phase noise power spectral density for different circuit components. For typical settings, we conclude that scaling up to a large number of antenna elements by increasing the number of tiles (each tile supporting a fixed number of antenna elements) is attractive in two ways: (1) for a fixed number of users, reducing the load factor (ratio of number of users to number of antennas) improves performance with phase noise, and (2) the independent phase noise generated at each tile gets averaged across tiles. We illustrate our numerical results for a nominal 140 GHz system with per-user data rates of 10 Gbps.**

***Index Terms*—Millimeter Wave, THz, Multi User, Phase Noise**

I. Introduction

Millimeter wave (mmWave) bands offer potential channel bandwidths of several GHz. Furthermore, the small carrier wavelength (e.g., 5 mm at 60 GHz, and $\approx$ 2 mm at 140 GHz) enables us to fit antenna arrays with a very large number of elements in compact form factors, enabling the synthesis of electronically steerable pencil beams. While there have been significant advances in low-cost mmWave hardware over the past decade, existing mmWave hardware is typically based on RF beamforming, which enables synthesis of only one beam at a time. However, we expect fully digital beamforning to become feasible in the relatively near future, with one RF chain per antenna element: this will enable the synthesis of multiple simultaneous beams, and lead to truly massive MIMO with 100s or 1000s of antenna elements, supporting 100s of simultaneous users, each at rates of Gbps. In this paper, we analyze the impact of one potential bottleneck in such a system: phase noise. A systematic framework for determining the impact of phase noise on MIMO performance is important because oscillators at higher carrier frequencies typically exhibit higher noise power spectral density, and the increased system bandwidth that we must integrate over contributes further to increased distortion in the phase.

In order to scale to a large number of elements, we consider a modular approach in which a large array is built from subarrays, or tiles, with separate RF processing on each tile. We focus on uplink massive MIMO, hence the RF processing in a tile consists of downconversion to baseband, followed by quantization. Thus, only digital baseband signals are communicated from each tile to a central processor. While this approach greatly simplifies RF hardware design (each tile can be controlled by a separate RF chip) and produces an easily expandable system, synchronizing the subarrays with each other requires distribution of a common clock. Distributing a high-frequency reference over a large area incurs high propagation loss and is therefore inefficient. We consider here an alternative approach, in which a lower frequency clock is distributed to the tiles, with each tile synthesizing its local oscillator using, for example, a harmonic multiplier or a phase locked loop, as shown in Fig. 2.

Contributions: We analyze the effect of phase noise on a tiled architecture, where phase noise is generated both in the reference oscillator and in independent oscillators that are locked to the reference on each tile. We assess system-level impact by considering a MIMO uplink with K simultaneous users communicating with a base station using a tiled phased array, with singlecarrier modulation and linear MMSE (LMMSE) reception. While our approach applies to a variety of frequency multiplication techniques, for concreteness, we focus on PLL-based generation of the local oscillator in each tile. We model propagation of phase noise through the two-step clock generation using a linearized PLL model. Our conclusions are summarized as follows:
(a) The **common phase noise** arising from the low-frequency clock "passes through" the LMMSE receiver, and contributes to symbol rotation after interference suppression. Since the PLL in each tile acts as a lowpass filter, the phase rotations across symbols vary slowly and can be tracked effectively. Thus, for typical settings, the impact of common phase noise is small.
(b) The **independent phase noise** from the voltage controlled oscillator (VCO) in each tile sees a highpass response from

978-1-7281-1151-3/19/$31.00

the PLL, and cannot be tracked at the symbol time scale. The effect of independent VCO phase noise can be summarized as follows:

- Self-interference primarily in the form of a phase rotation in the received signal; this effect is averaged over tiles and therefore diminishes as number of tiles is increased.
- Multiuser interference or "crosstalk"; this phenomenon is insensitive to the number of tiles but proportional to load factor (ratio of number of users to array size).

We provide an approximate analysis which shows that, for typical settings, we get substantial benefit by averaging across tiles. Simulation results are provided to validate these predictions.

(c) The dependence of the phase noise PSD on frequency is often modeled as $L(f) = a_0 + \frac{a_1}{f} + \frac{a_2}{f^2} + \frac{a_3}{f^3}$, with parameters $\{a_i\}$ that depend on the oscillator structure and design. We note that, under our linearlized PLL model, these parameters can be directly related to the output phase noise variances. We use this observation to translate output performance criteria such as bit error rate (BER) back to "masks" that specify the maximum allowable envelope of each term of the form $\frac{a_i}{f^i}$, thus providing a framework for circuit design aimed at controlling the phase noise PSD.

Related Work: In [1], the authors consider a MIMO channel in the presence of phase noise in both transmitter and receiver. They derive and heuristically solve an approximation of the maximum likelihood function of the transmitted signal, but temporal correlations of phase noise are ignored. The authors of [2] and [3] consider a massive MIMO uplink/downlink system with a common oscillator at the base station and hence common phase noise at all the antenna elements. They proposed a zero-forcing based low-complexity phase noise suppression detection/precoding algorithm. In [4], the authors studied the effect of the phase noise on uplink with imperfect channel state information. They showed that the system performance degraded as the number of oscillators increases at the base station. In [5], PLL phase noise in the uplink of a MIMO OFDM system was analyzed, showing that the array averages out the uncorrelated part of the phase noise. Both cases of synchronous and asynchronous operation were considered. However, the tiled architecture was not examined.

II. System model

The mmWave uplink is depicted in Fig. 1. As a running example, we consider a carrier frequency of 140 GHz, a target bit rate (per user) of 10 Gbps, and either QPSK or 16QAM modulation, corresponding to 5 GHz and 2.5 GHz of bandwidth, respectively, ignoring excess bandwidth. In our simulations, we assume coarse power control in which the received power from different mobiles are within 5 dB of each other.

Since we are interested mainly in horizontal scanning, we consider a linear array with $N = 256$ elements. The BS antenna array is arranged in N_t tiles, each with N_0 elements, i.e., $N = N_t \times N_0$ elements in total. Our nominal system employs $N_t = 16$ tiles, each with $N_0 = 16$ elements. We define the load factor as $\beta = \frac{K}{N}$, where K denotes the number of users. We vary β between $1/256$ (i.e., single user) and $1/2$ to obtain insight into the impact of phase noise on spatial multiplexing.

Phase Noise Model

Phase noise is the random fluctuation in the phase of a signal. For a noisy output sinusoid,

$$c(t) = e^{j2\pi f_c t + j\phi(t)},$$

$\phi(t)$ denotes the phase noise random process with power spectral density $L(f)$. The phase noise PSD contains several components; a constant corresponding to additive white thermal noise, and $1/f$, $1/f^2$, and $1/f^3$ components that are the result of different effects in active elements of the oscillator as determined by Leeson's formula [6]. The phase noise spectral density is modeled as

$$L(f) = a_0 + \frac{a_1}{f} + \frac{a_2}{f^2} + \frac{a_3}{f^3}. \quad (1)$$

The PSD function, $L(f)$, is expressed in units of W/Hz or J, meaning a_0, a_1, a_2 and a_3 have units J, J/s, J/s^2, and J/s^3, respectively. These parameters depend on the oscillator structure and design, and we wish to provide guidance on the maximum allowable values of these parameters for a desired system performance.

The reference XO and tile VCO's can have different noise statistics. We denote by $L_X(f)$ and $L_V(f)$ the PSD of reference oscillator and tile VCO's respectively. The phase noise at each element is a linear combination of the two after being passed through the PLL. The former is common over the entire array while the latter is identical for elements on the same tile but independent for different tiles.

Downconversion and Interference Suppression

The received signal at each element is down-converted using the frequency multiplied clock at the respective tile, and the phase noise present at each element translates directly to phase noise in the demodulated complex baseband signal. Assuming that the system operates in a high SNR regime, the LMMSE estimate of the K channels (assuming perfect

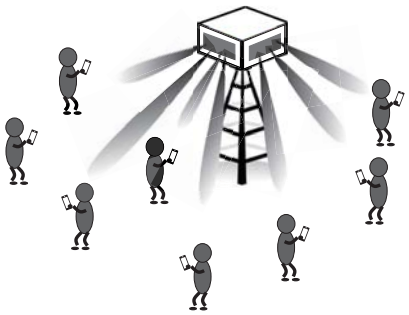

Fig. 1: Mmwave base station supporting many users with large array forming narrow pencil beams.

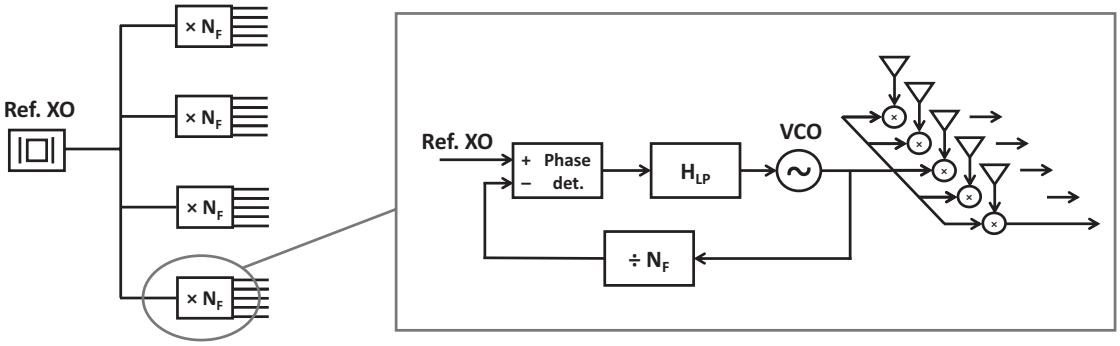

Fig. 2: Clock distribution in the tiled structure.

knowledge of CSI) is close to the zero-forcing receiver response. The high-SNR assumption simplifies the analysis and allows us to isolate the effect of phase noise from other sources of distortion such as additive receiver noise.

The channel from each user is assumed to be a single line-of-sight path identified by its angle of arrival. Assuming users are located at angles $\{\theta_k, k = 1, ..., K\}$ corresponding to spatial frequencies $\{\omega_k = \frac{2\pi d}{\lambda} \sin\theta_k\}$, where d is the element spacing and λ is the carrier wavelength, the n'th entry of the N-dimensional signal on the array is given by

$$x_n = \sum_{k=1}^{K} A_k e^{j\omega_k n} s_k,$$

where s_k is the complex symbol sent by the k'th user. Here, A_k is the complex amplitude of the k'th user's channel. Since power leveling is assumed, this value is chosen in simulations such that its magnitude is uniformly distributed over the range of $[-5, 0]$ dB.

Constellation Derotation

Lowpass components of PSD cause slow-changing phase drift that blows up over time, but can be controlled by tracking phase noise across symbols. Assuming correct symbol decoding, we can estimate the phase noise seen by that symbol. We average such estimates over a window W of past symbols to predict the phase noise in the current symbol, and derotate the received sample by the estimated amount. The residual phase noise can be modeled as a filtered version of the input phase noise, with filter impulse response,

$$h_W[t/T_s] = \delta[t/T_s] - \frac{1}{W}\sum_{d=1}^{W} \delta[t/T_s - d].$$

The frequency response of this filter is plotted in Fig. 3 for different values of W. The derotation filter enhances high pass noise components and suppresses the low pass portion of noise power. Larger windows result in less high-pass noise enhancement, but have a smaller suppression bandwidth.

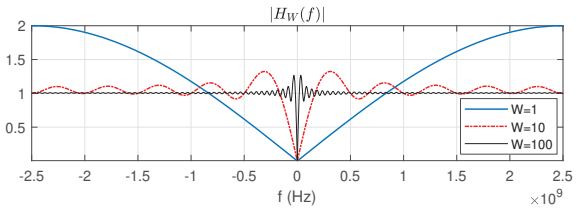

Fig. 3: Frequency response of derotation filter for different window sizes. $W = 1$ is equivalent to differential modulation.

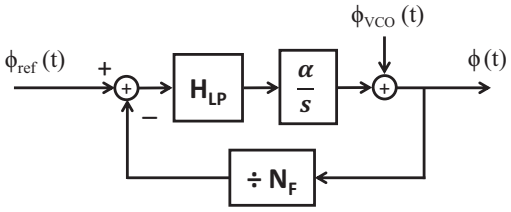

Fig. 4: LTI system model of PLL for phase noise.

III. Phase Noise in a SISO System

As a first step, we consider a single user received by a single antenna. Using the linearized PLL model of Fig. 4, the PLL transfer function can be obtained as,

$$\phi(s) = H_{VCO}^{PLL}(s)\phi_{VCO}(s) + H_{ref}^{PLL}(s)\phi_{ref}(s) \qquad (2)$$

where H_{ref}^{PLL} and H_{VCO}^{PLL} identify the output phase response to signals ϕ_{ref} and ϕ_{VCO} respectively. The specific PLL parameters used in simulations are described in Section VI.

The PLL applies a lowpass filter to the reference phase noise and a highpass filter to the VCO phase noise. This is followed by the derotation filter $H_W(f)$ discussed earlier. The total phase noise at the receiver output, that rotates the measured symbol, has PSD $S_\phi(f)$ that is the sum of reference and VCO contributions,

$$S_\phi(f) = S_X(f) + S_V(f),$$

where

$$S_X(f) = L_X(f)\,|H_{ref}^{PLL}(f)H_W(f)|^2,$$

$$S_V(f) = \; L_V(f)\,|H_{VCO}^{PLL}(f)H_W(f)|^2.$$

The output phase noise is modeled as a Gaussian random variable $\phi = \phi^X + \phi^V$, with variance $\sigma_\phi^2 = \sigma_X^2 + \sigma_V^2$ derived as,

$$\sigma_\phi^2 = R_\phi(0) = \int_{-B/2}^{B/2} S_\phi(f)df \tag{3}$$

$$\sigma_X^2 = \int_{-B/2}^{B/2} S_X(f)df, \qquad \sigma_V^2 = \int_{-B/2}^{B/2} S_V(f)df.$$

We observe that the value of σ_ϕ^2, which determines system performance, is a linear function of the a_i parameters that describe the phase noise PSDs for the reference clock and the tile VCOs. Thus, constraints on σ_ϕ^2 imposed by system performance requirements can be translated to a phase noise PSD mask based on maximum allowable values of $\{a_i\}$. An interesting observation is that the contribution of common (reference) phase noise to the overall σ_ϕ^2, that is, the phase noise variance seen at the receiver output, turns out to be negligible. Between lowpass filtering of the phase-locked loop and DC rejection due to the derotation filter, almost all of the reference phase noise is suppressed. Thus, VCO phase noise is greatly dominant in the final measured symbol.

IV. Phase Noise in Single User Array

When receiving the signal of a single user on the tiled array, LMMSE or zero-forcing reception reduces to receive beamforming (i.e., matched filtering). Denoting the transmitted symbol by s and the LMMSE output as y, the perfectly noiseless receiver array would achieve $y = Ns$ as the signal of all elements are combined with perfect phase alignment. Phase noise causes phase rotation of the complex signal at each element; as long as these rotations are identical on all elements the overall response is simply a phase-rotated version of the input signal. Variation of phase noise terms over array elements rotates element signals in different directions and disturbs the phase alignment of the combining vectors. This results in fading in the amplitude of the output. Thus the amplitude and phase distortion caused by phase noise can be modeled by a random complex distortion factor $\alpha = A_{SU}e^{\Phi_{SU}}$ that describes the output as $y = \alpha s$. The distributions of A_{SU} and Φ_{SU} can be found by applying the small angle approximation, $e^{j\phi_n} = \cos\phi_n + j\sin\phi_n \approx 1 + j\phi_n$, which allows us to express α as,

$$\alpha = \sum_{n=1}^{N} e^{j\phi_n} \approx N + j\sum_{n=1}^{N}\phi_n = N\left(1 + j\frac{1}{N}\sum_{n=1}^{N}\phi_n\right)$$

Denoting the average phase noise across antennas as $\bar{\phi} = \frac{1}{N}\sum_{n=1}^{N}\phi_n$, we therefore obtain that $\alpha \approx N(1 + j\bar{\phi}) \approx Ne^{j\bar{\phi}}$, so that $\Phi_{SU} = \angle\alpha \approx \bar{\phi}$. Here the independent VCO phase noise terms of different tiles are averaged at the output resulting in lower phase noise variance than the single chain structure.

$$\Phi_{SU} \approx \phi^X + \frac{1}{N_t}\sum_{i=1}^{N_t}\phi_i^V \sim \mathcal{N}(0, \sigma_X^2 + \frac{\sigma_V^2}{N_t}) \tag{4}$$

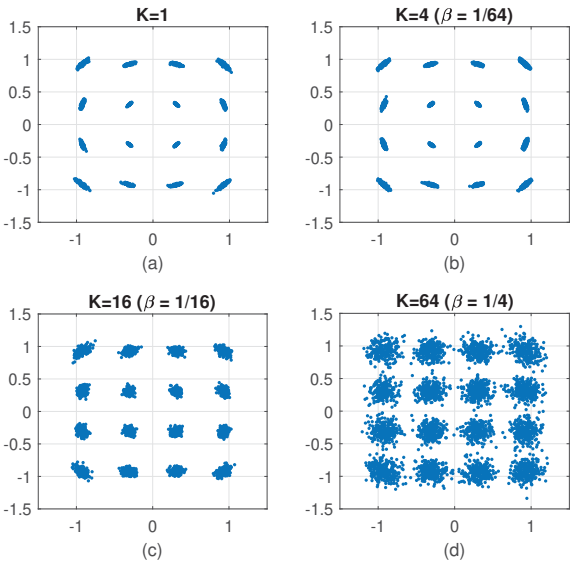

Fig. 5: Scatter plot of received 16-QAM symbols on a 256-element array for different loading factors. Phase noise in the tiled structure results in angular rotation of symbols as well as crosstalk between users in the form of additive complex Gaussian interference. The latter becomes dominant as the number of users, or loading factor, increases.

Despite averaging of VCO noise terms, the reference phase noise, ϕ_X, is still negligible for typical parameters, making up less than 1% of total noise power after derotation, even with maximum averaging, i.e., $N_t = N$.

While we have used $A_{SU} = |\alpha| \approx N$ when estimating the phase, we can obtain a more refined estimate under our small angle approximation as follows:

$$A_{SU} \approx \sum_{n=1}^{N}\cos\phi_n \approx \sum_{n=1}^{N}\left(1 - \phi_n^2/2\right)$$

Since ϕ_n^2 has chi-squared distribution with mean σ_ϕ^2 and variance $2\sigma_\phi^4$, using the central limit theorem, A_{SU} can be approximated as Gaussian:

$$A_{SU} \sim \mathcal{N}(N(1 - \frac{\sigma_\phi^2}{2}), \frac{N_t N_0^2}{2}\sigma_\phi^4). \tag{5}$$

For small σ_ϕ^2, the distortion in amplitude is negligible and the much larger phase noise is dominant in determining BER. Overall, the single-user distortion from phase noise leads primarily to spreading of constellation symbol phases as shown in Fig. 5a.

V. Phase Noise in Multiuser MIMO

When receiving the signal of K different users, the receiver uses knowledge of the channel to each user to provide LMMSE estimates of the K symbols. Since we assume high-SNR operation, the LMMSE receiver is close to the zero-forcing (ZF) solution, which we consider henceforth. The ZF solution consists of applying the $K \times N$ matrix H_{ZF} to the $N \times 1$ received signal, where $H_{ZF} = H_C^+$, the pseudoinverse of the channel matrix.

The output symbol y_p for user p is the inner product of the p'th row of H_{ZF} and the phase distorted array response. Therefore y_p will be a linear combination of input signals, $\{s_q\}$. The response of y_p to s_p is predominantly the phase rotation defined as Φ_{SU} in (4). The response to all other symbols acts as *interference* that is added to the output symbol.

Take the user pair p and q at angles θ_p and θ_q. The response of each on the array is $\mathbf{x}(\omega_p) = s_p[e^{j\omega_p 0}, ..., e^{j\omega_p(N-1)}]^T$ and $\mathbf{x}(\omega_q) = s_q[e^{j\omega_q 0}, ..., e^{j\omega_q(N-1)}]^T$ respectively. Assuming they are sufficiently separated in angular domain, we have the following:

$$\langle H_{ZF}(p,.), \mathbf{x}(\omega_p)\rangle \approx N s_p$$
$$\langle H_{ZF}(p,.), \mathbf{x}(\omega_q)\rangle = 0$$

The second equation can be written as a sum of complex vectors that sum to 0, as per (6).

$$\langle H_{ZF}(p,.),\mathbf{x}(\omega_q)\rangle = \sum_{n=1}^{N} H_{ZF}(p,n)e^{j\omega_q n} = \sum_{n=1}^{N} Z_n$$
$$= \sum_{t=1}^{N_t}\sum_{m=1}^{N_0} Z_n = \sum_{t=1}^{N_t} C_t = 0. \quad (6)$$

wherein C_t is the *subarray* response of row p of the zero-forcing matrix. In the present structure, tile t is distorted by phase noise ϕ_t, so that the cross interference, I_{pq} that satisfies $y_p = I_{pq}s_q$ is no longer zero, but rather can be found by

$$\begin{aligned} I_{pq} &= \sum_{t=1}^{N_t} C_t e^{j\phi_t} \\ &\approx e^{j\phi^X}\sum_{t=1}^{N_t} C_t(1+j\phi_t^V) \\ &= e^{j\phi^X}\sum_{t=1}^{N_t} C_t + e^{j\phi^X}\sum_{t=1}^{N_t} C_t j\phi_t^V \\ &= e^{j\phi^X}\sum_{t=1}^{N_t} C_t j\phi_t^V \sim \mathcal{CN}(0, N_t\mathbb{E}[|C_t|^2]\sigma_V^2). \end{aligned}$$

To determine the interference power, an estimate for the value of C_t^2 is required. It can be shown that the *subarray* response of zero-forcing is very close to that of matched filtering. In fact, assuming angular separation between the desired transmitter and interferer is uniformly distributed, the expected value of subbaray response magnitude varies by less than 5% between the zero-forcing and matched filtering response in the worst case of $N_t = 2$. Thus the magnitude of C_t is approximately equal to N_0 times the normalized subarray pattern gain,

$$G_0(\omega_p - \omega_q) := \left|\frac{1}{N_0}\sum_{n=1}^{N_0} e^{j(\omega_p-\omega_q)n}\right| = \left|\frac{\sin(N_0(\omega_p-\omega_q))}{N_0\sin(\omega_p-\omega_q)}\right|.$$

The variance of the interference term I_{pq} is approximately equal to

$$\text{var}(I_{pq}) \approx N_t N_0^2 G_0^2(\omega_p-\omega_q)\sigma_V^2$$

and the ratio of interference power from user q to signal power of user p is thus equal to

$$\frac{\text{var}(I_{pq})}{S_p} \approx \frac{N_0 G_0^2(\omega_p-\omega_q)\sigma_V^2}{N}.$$

The total interference power is the sum of interference from all other users,

$$\frac{\sigma_I^2}{S_p} \approx \frac{(K-1)N_0\overline{G_0^2}\sigma_V^2}{N} \approx \beta\sigma_V^2. \quad (7)$$

Here we have approximated the average subarray gain $\overline{G_0^2}$ by its statistical mean, $1/N_0$. A key observation is that multiuser interference is a direct result of tiling, and would be zero if all tiles had identical phase noise. Thus while tiling provides an averaging gain in phase noise as described by (4), it adds crosstalk between channels.

At the receiver output for each channel, the distortion caused by multiuser interference is present alongside the angular rotation of the desired symbol as described in the previous section. Fig. 5 shows scatter plots of received samples for a 16-QAM signal under low, medium and high loading. For $\beta = 1/64$, phase noise is dominant over interference, resulting in near SIMO performance. For larger loading factors interference becomes dominant and the scatter cloud looks more circular.

VI. Results and Discussion

In this section we provide simulation results to gain insight into the requirements for oscillator quality and phase noise mask given certain system design choices. Simulation parameters are as follows. Frequency multiplication is by a factor of $N_f = 14$, and the multiplier is a type-2 PLL with loop resonance frequency of 1 MHz and damping factor of 0.707. While the VCO and reference oscillators can have different characteristics, we assume they have identical phase noise PSD. In each case, the $\{a_i\}$ parameters are chosen such that all terms in (1) have equal contribution on output SINR under nominal array size (16×16) and loading factor ($\beta = 1/4$). These values are reported in Fig. 6. The derotation filter has window size $W = 10$.

One interesting comparison is between constellation size and bandwidth; higher order constellations require lower bandwidth to obtain the same bit rate, and while increasing bandwidth intensifies phase noise, a larger constellation is more dense and provides smaller margin of error. To support a bit rate of 10 Gbps, we consider two scenarios: Case I, in which QPSK modulation is used over a bandwidth of 5 GHz, and Case II, where 16-QAM signaling is deployed with a smaller bandwidth of 2.5 GHz. In both cases a moderate to high loading factor of $1/4$ was chosen. Fixing the shape of the PSD curve, the required phase noise performance in each case is depicted in Fig. 6, which shows a less stringent phase noise mask for QPSK. Thus in systems where phase noise is a bottleneck, choosing a smaller constellation size at the cost of larger bandwidth is beneficial. Note that the specific shape of the PSD curve can be tuned to provide the desired trade-off between low and high frequency mask;

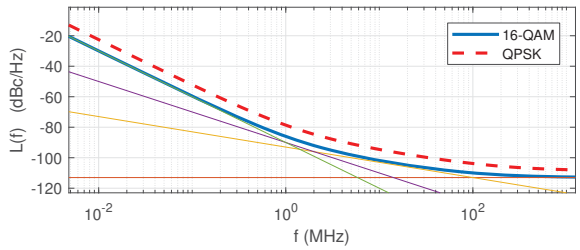

Fig. 6: $L(f)$ curves providing $P_E \leq 10^{-3}$ in two cases with $\beta = 1/4$. Narrow lines show the a_i/f^i components for case II. Case I: $a_0 = 1.4 \times 10^{-11}$ J, $a_1 = 2.75 \times 10^{-3}$ J/s, $a_2 = 5.5 \times 10^3$ J/s^2, $a_3 = 5.5 \times 10^9$ J/s^3. Case II: $a_0 = 5 \times 10^{-12}$ J, $a_1 = 5 \times 10^{-4}$ J/s, $a_2 = 10^3$ J/s^2, $a_3 = 10^9$ J/s^3.

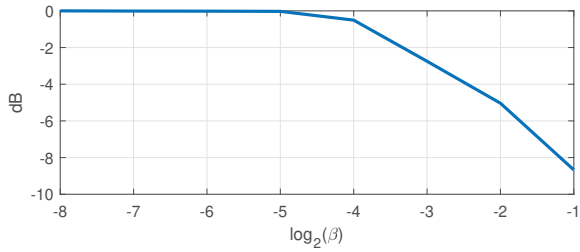

Fig. 7: Tolerable phase noise for $P_E < 10^{-3}$ with QPSK modulation, as a function of loading factor. Values are normalized to single user tolerance, and array size is 256.

e.g., by decreasing the strength of the $1/f$, $1/f^2$, and $1/f^3$ components, a higher constant (noise floor) can be permitted.

Another important design choice is loading factor. One significant benefit of modular designs is that they reduce the cost and complexity of increasing array size. Thus in many applications, in order to serve a fixed number of users it may be preferable to deploy larger arrays with less stringent per-element specifications. To see the effect of loading factor on oscillator phase noise requirements, the maximum allowable PSD mask has been depicted in Fig. 7 as a function of loading factor, normalized to the single user mask. The mask shape is fixed and its peak is reduced until BER is below 10^{-3}. We see here that by reducing loading from $1/2$ to $1/8$, the allowable phase noise mask is lifted by about 8 dB. We also observe that if the number of tiles, N_t, is fixed to 16 and array size is scaled to $N = 512$ or 1024, then as long as the loading factor $\beta = K/N$ is maintained at $1/4$, the exact same performance is achieved when using the spectral mask reported in Fig. 6. This is in agreement with the result in (7).

Fig. 8 quantifies the effect of tiling on system performance. For high loading factor, crosstalk interference is more pronounced and offsets the gain of phase noise averaging over tiles. At low loading, however, the tiled structure is beneficial as self-interference is dominant over crosstalk. We may conclude that using a modular structure not only allows for greater scaling but improves performance and loosens oscillator quality requirements. This conclusion must be taken with a grain of salt, however, keeping in mind that the current analysis is limited to a relatively low phase noise linearizable regime.

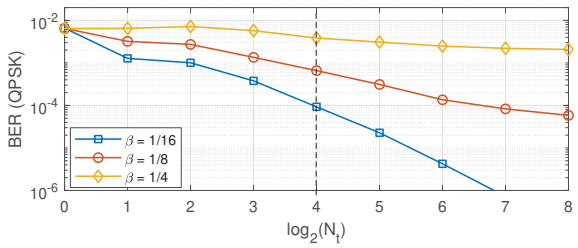

Fig. 8: Effect of tiling on performance of 256 element array (200 realizations, minimum user separation $\Delta\omega_{\min} = \pi/N$).

VII. Conclusions

Our results indicate that phase noise is not a bottleneck in scaling up the massive MIMO system using a tiled architecture. For the two-step clock generation scheme considered here, the impact of the reference oscillator phase noise is negligible for typical settings, and is in any case independent of scaling up the number of users: it simply appears as a phase shift at the output of MIMO processing, and can be tracked effectively due to its lowpass nature. The self-interference due to phase noise generated at the tile VCOs gets averaged across tiles, and hence is reduced as the number of tiles increases, whereas the multiuser interference is insensitive to the number of tiles, and scales with the load factor. Thus, starting with tiles with a fixed number of antennas, we can scale up the number of supported users, while maintaining a desired load factor, by scaling up the number of tiles.

Acknowledgments

This work was supported in part by ComSenTer, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

The authors wish to thank Ali Farid, Ahmed Ahmed, and Arda Simsek for their valuable insight and helpful discussions.

References

[1] R. Combes and S. Yang, "An approximate ML detector for MIMO channels corrupted by phase noise," *IEEE Transactions on Communications*, vol. 66, no. 3, pp. 1176–1189, 2018.

[2] Y.-F. Wang and J.-H. Lee, "A simple phase noise suppression scheme for massive MIMO uplink systems," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 6, pp. 4769–4780, 2017.

[3] ——, "A ZF-based precoding scheme with phase noise suppression for massive MIMO downlink systems," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 2, pp. 1158–1173, 2018.

[4] Y. Fang, X. Li, and L. Qiu, "Asymptotic equivalent performance of uplink massive MIMO systems with phase noise," in *2018 IEEE International Conference on Communications (ICC)*.

[5] A. Puglielli, G. LaCaille, A. M. Niknejad, G. Wright, B. Nikolić, and E. Alon, "Phase noise scaling and tracking in OFDM multi-user beamforming arrays," in *Communications (ICC), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1–6.

[6] D. Lesson, "A simple model of feedback oscillator noise spectrum," *proc. IEEE*, vol. 54, no. 2, pp. 329–330, 1966.