Phase Noise Scaling and Tracking in OFDM Multi-user Beamforming Arrays

Antonio Puglielli*, Greg LaCaille*, Ali M. Niknejad*, Gregory Wright†, Borivoje Nikolić*, Elad Alon*
*University of California, Berkeley, CA 94704 USA
†Alcatel-Lucent Bell Labs, Holmdel, NJ 07733 USA

***Abstract*—Many-element antenna arrays, used for multi-user MIMO, are expected to be one of the cornerstone technologies for 5G wireless systems. Large arrays also offer the opportunity to average out some of the transceivers' analog imperfections, potentially enabling a lower-power implementation. In this paper we study the effect of local oscillator phase noise on beamforming MU-MIMO-OFDM systems. We show that the array does average out uncorrelated phase noise at each element. Exploiting this, we propose scaling the per-element phase noise specification proportionally to the array size, thereby maintaining constant array-level performance with lower power consumption. However, if the phase noise is entirely uncorrelated, this scaling causes a substantial degradation in the recovered signal energy. If, instead, some correlated low-frequency phase noise is introduced at each element, we show that phase noise scaling incurs no performance loss. In fact, under these conditions, a single, global pilot tracking loop can replace carrier recovery at each element. Additionally, this level of phase noise correlation eliminates the phase noise-induced channel aging effect. This type of correlation can be achieved by distributing a common reference and optimizing the bandwidth of the PLL.**

I. Introduction

Many-antenna arrays using multi-user (MU) beamforming techniques are expected to play an important role in 5G systems [1]. One technology which has received much recent interest is massive MIMO, where the base-station uses an array with a large number of elements to communicate with a smaller number of user terminals. In this regime, linear beamforming asymptotically achieves the capacity of the multi-user channel [2]–[4]. Another promising technology is the use of mm-wave bands [1]. Mm-wave systems may use large arrays to achieve high beamforming gain, compensating for link budget challenges at high frequency [5].

From an engineering perspective, large arrays can be used to relax hardware design requirements by exploiting averaging of uncorrelated analog imperfections. This has been well-studied in the context of massive MIMO [6]–[8]. By exploiting this averaging, one can envision constructing a massive MIMO array out of low-performance, energy-efficient transceivers while recovering the desired performance at the array level, thus lowering total power consumption.

The impact of phase noise on orthogonal frequency-division multiplexing (OFDM) receivers has been thoroughly studied in the literature [9]–[16]. Early studies focused on characterizing the signal to interference and noise ratio (SINR) [9], [11], [15] and the bit error rate [9], [12]. Later works investigated the loss of orthogonality between subcarriers [13]–[16], and proposed techniques to correct for these effects [10], [14]–[16]. Though phase noise is commonly modeled in the literature as a Wiener process arising from a free-running oscillator, Petrovic *et al.* [16] extended the analysis to the case of a phase-locked-loop (PLL), which reflects a more realistic implementation scenario.

Recently, there has been interest in extending these results to arrays. In [17], the authors demonstrated that phase noise which is uncorrelated between transmitters is averaged by the spatial combining. In [18]–[20], these results are extended to a multi-user massive MIMO scenario with single-carrier (SC) modulation and the authors develop expressions for the SINR. [21] studies the capacity of MISO and SIMO channels. All of these studies used a free-running oscillator model instead of a PLL and only considered either fully correlated or fully uncorrelated phase noise. They also did not consider the use of phase noise tracking or carrier recovery techniques.

In this paper, we analyze the effect of phase noise on beamforming arrays with realistic frequency generation schemes and phase noise tracking. Our main contributions are: (1) analyzing phase noise in a massive MU-MIMO-OFDM system with PLL-based frequency generation; (2) exploiting array gains to relax the per-transceiver phase noise specifications and reduce implementation power consumption; and (3) studying carrier recovery in massive arrays and deriving PLL design guidelines. Section II reviews PLL operation and introduces the system model. In Section III we study the effect of phase noise on MU beamforming and propose phase noise scaling to relax frequency generation specifications. Section IV studies the use of a global pilot tracking loop with different frequency generation schemes and derives key PLL design guidelines. Finally, in Section V we validate our findings with simulation results and offer concluding thoughts in Section VI.

II. System Model

In this section, we survey important characteristics of the phase noise of an oscillator within a PLL and introduce two frequency generation models for large arrays.

A. Phase Noise in PLLs

A PLL is used to generate a high frequency oscillation locked to a stable low-frequency reference. Due to their tunability, voltage-controlled oscillators (VCOs) have poor frequency and phase stability. On the other hand, crystal oscillators (XOs) generate very precise and stable oscillations

978-1-4799-6664-6/16/$31.00 ©2016 IEEE

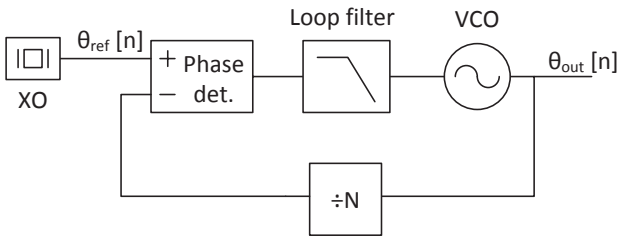

(a)

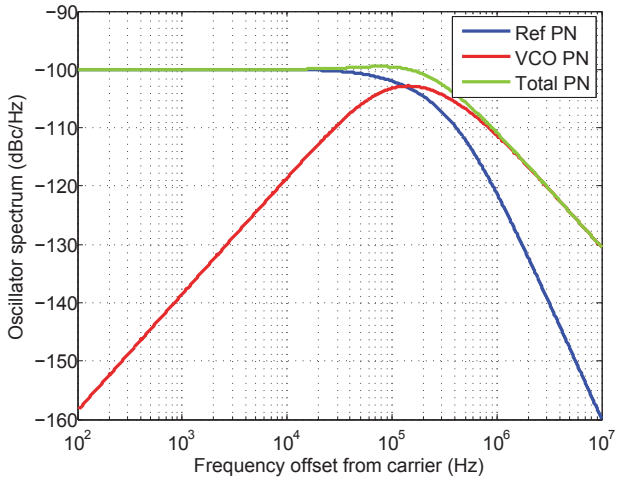

(b)

Fig. 1. (a) Block diagram of a PLL. The output phase of a VCO is divided, compared with an accurate reference crystal, filtered, and fed back to the VCO's control input. (b) Representative PSD of the PLL's phase noise, indicating contributions from both the reference and VCO.

but at a fixed, and typically low, frequency. By placing a high-frequency VCO in a feedback loop referenced to an XO (Fig. 1(a)), a PLL can transfer the phase and frequency stability of the crystal to the VCO. Within the PLL bandwidth, B_{PLL}, the output phase tracks the reference while any VCO noise is suppressed. Outside of the PLL bandwidth, the VCO operates in open loop and all of its noise therefore appears at the output. As an example, the phase noise contributions of the reference and the VCO are shown in Fig. 1(b) for a type-II, third-order PLL with 150kHz bandwidth. A detailed discussion of the common charge-pump PLL can be found in [22].

The phase variance can be written in terms of the reference and VCO phase variances. Since the reference is an open-loop oscillator, its phase noise follows a Wiener process and the phase variance is well-known: $\sigma^2_{ref}[n] = n\sigma^2_{ref}$ [14]–[16], [18]–[20]. In contrast, the VCO phase noise is filtered by the PLL transfer function and is therefore wide-sense stationary [16], [23], [24]. The total phase variance is then:

$$\sigma^2_\theta = n\sigma^2_{ref} + \sigma^2_{VCO} \tag{1}$$

The phase noise $\theta[n]$ impacts the transceiver via the oscillator voltage $v[n] = e^{j\theta[n]}$. The statistics of this process have been studied in [16], [23], [24], where the power spectral density (PSD), $S_v(\omega)$, is derived. In this paper, we use the prototype PSD in Fig. 1(b), which is representative of a wide range of PLL architectures, to analyze the effect of phase noise on communication systems that use PLL-based frequency generation. This allows us to make design decisions in terms of spectra which are well-characterized and understood.

B. Frequency Generation and Distribution in Large Arrays

There are three main options for frequency generation and distribution in a large array:

1) Generate a single high-frequency local oscillator (LO) signal centrally and distribute it throughout the array.
2) **Synchronous distribution**: Distribute a single XO reference to all transceivers which independently generate LOs using separate PLLs.
3) **Asynchronous distribution**: Use independent references at each element, generating fully independent clocks from a PLL at each transceiver.

The first option is undesirable since global distribution of a high-frequency clock consumes excessive power and can lead to unacceptable waveform degradation. Instead, to avoid these problems, either the second or third option should be used. If a single reference is shared across the array, its low-frequency phase noise components will appear at each element. Phase noise at frequencies above the PLL bandwidth arises from the free-running evolution of each individual oscillator and will be independent across the array. On the other hand, if separate references are used at each element, then the phase noise processes of each LO will be entirely uncorrelated.

C. Single-Element Receiver Model

Consider an OFDM communication system using N_{sc} subcarriers, each with bandwidth B_{sc}, for a total channel bandwidth of $B = N_{sc}B_{sc}$. All symbols have energy E_s and are statistically independent. The transmitter modulates a data and pilot sequence $\{d_k\}$ onto the N_{sc} subcarriers. The resulting time-domain sequence $s[n]$ is transmitted through a channel with impulse response $h[n]$ and additive white Gaussian noise $w[n]$ with variance σ^2. The receive signal is corrupted by a unit energy oscillator sequence drawn from the process $v[n]$. Since the phase noise does not alter the statistical properties of the additive white noise, we can write

$$y[n] = (s[n] \odot h[n])v[n] + w[n] \tag{2}$$

The multiplication by $v[n]$ creates a circular convolution in the frequency domain between the data and oscillator spectra. Denoting the DFT of the phase noise sequence, channel, and white noise as $\{Q_k\}, \{H_k\}$, and $\{W_k\}$, respectively, the received frequency domain samples $\{Y_k\}$ are

$$Y_k = d_k H_k Q_0 + \sum_{j \neq k} d_j H_j Q_{i-j} + W_k \tag{3}$$

Phase noise has two main effects. First, each symbol is affected by a common phase error (CPE) Q_0, which appears as the multiplication of the complex channel gain equally across all subcarriers. Second, the symbols are corrupted by inter-carrier interference (ICI): mixing of data streams caused by loss of orthogonality between subcarriers.

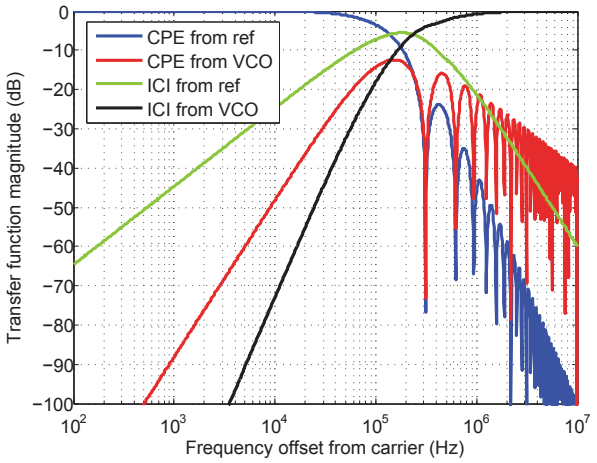

Fig. 2. Cascaded transfer function of reference and VCO phase noise to the CPE and ICI, with $B_{PLL} = 200kHz$ and $B_{sc} = 312kHz$. These results validate the conclusion that reference noise dominates the CPE while VCO noise dominates the ICI generation.

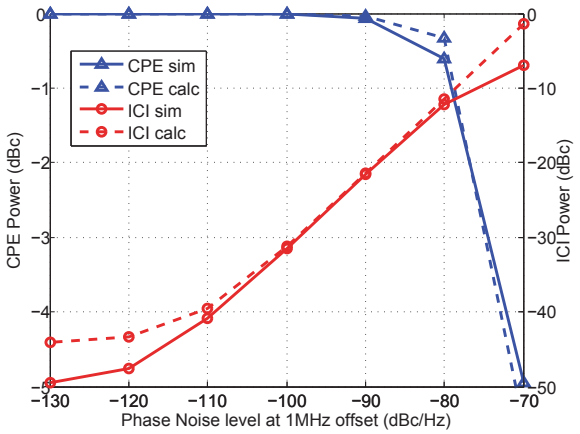

Fig. 3. Simulated and predicted results of ICI and CPE power based on the model in (7) and (8), for a PLL with 200kHz bandwidth and B_{sc} of 624kHz.

It is common in OFDM systems to reserve a small number of subcarriers for pilots [10], [14] which are used at the receiver to estimate various imperfections. Using the N_p pilots, found at the set of indices $\mathcal{P}$, the CPE can be estimated as:

$$\hat{Q}_0 = \frac{1}{N_p} \sum_{k \in \mathcal{P}} \frac{Y_k}{d_k} \tag{4}$$

The CPE can then be compensated by dividing all received symbols by this estimate. Ideally, this cancels the sample mean of the phase noise process on a symbol-by-symbol basis, giving an average SINR of

$$SINR = \frac{\mathbb{E}[|Q_0|^2]}{\mathbb{E}[|ICI|^2] + \frac{\sigma^2}{E_s}}. \tag{5}$$

D. Impact of Subcarrier and PLL Bandwidths

To gain a deeper understanding of the SINR, the CPE and ICI can be characterized in terms of the phase noise PSD. Define $\mathbf{1}_{N_{sc}}$ as a rectangular window of N_{sc} ones. Then the moving average filter of length N_{sc} has impulse response $h_{N_{sc}} = \frac{1}{N_{sc}} \mathbf{1}_{N_{sc}}$ and continuous-time frequency response

$$H_{N_{sc}}(\omega) = \frac{sin(N_{sc} \frac{\omega}{B})}{N_{sc}^2 \frac{\omega}{B}} \tag{6}$$

The width of this filter's main lobe is B_{sc}. Since the CPE is the DC component of the windowed phase noise, its energy is simply the energy of the random process generated by filtering $v[n]$ with $h_{N_{sc}}$. This can be expressed as:

$$\mathbb{E}[|Q_0|^2] = \int_{-\infty}^{\infty} S_v(\omega) |H_{N_{sc}}(\omega)|^2 d\omega \tag{7}$$

Since $v[n]$ has unit energy, the ICI energy is:

$$\mathbb{E}[|ICI|^2] = 1 - \mathbb{E}[|Q_0|^2] = \int_{-\infty}^{\infty} S_v(\omega)(1 - |H_{N_{sc}}(\omega)|^2) d\omega \tag{8}$$

This reveals a key insight: the CPE is contributed primarily by oscillator energy below the subcarrier bandwidth while ICI comes mainly from energy above the subcarrier bandwidth. This provides an intuitive explanation for how the subcarrier bandwidth affects the SINR. If B is kept fixed as N_{sc} is increased, the bandwidth of $h_{N_{sc}}$ is reduced. This will reduce the CPE energy and increase ICI, degrading the SINR.

Fig. 3 shows the simulated CPE and ICI power as a function of the VCO phase noise compared with the predictions in (7) and (8). These match well, including the signal energy degradation at very high phase noise levels. Also, over a wide range of phase noise levels, the ICI power increases linearly with the VCO's phase noise (1 dB/dB).

The CPE and ICI transfer functions can be cascaded with the PLL transfer functions to isolate the contributions of both reference and VCO to the CPE and ICI. Since the reference transfer function is lowpass, its phase noise will mostly contribute CPE and little ICI. Similarly, the VCO's phase noise will primarily add ICI rather than CPE. These conclusions are illustrated in Fig. 2 which captures the cascade of the PLL transfer functions with $H_{N_{sc}}(\omega)$. In fact, the PLL bandwidth can be optimized relative to the subcarrier bandwidth to tune these transfer functions (Section IV-B).

III. Phase Noise Scaling in Beamforming Arrays

In this section we analyze the impact of phase noise on beamforming OFDM arrays. The array considered here has M antennas and RF transceivers and performs beamforming independently on each subcarrier. For simplicity we assume perfect channel estimation and synchronization, but note that the impact of channel estimation errors in the presence of phase noise has been considered elsewhere [18]. The analysis is restricted to the uplink, though the downlink is very similar. We also do not include the phase noise of the user terminal since its effect is no different than in a SISO system. For notational simplicity, we drop subcarrier indices.

The K users transmit frequency-domain symbols $\mathbf{y}$ through $M \times K$ channel $\mathbf{H}$. The array uses zero-forcing beamforming matrix $\mathbf{G_{zf}}$ to estimate the user data streams. Splitting the CPE

at each element into correlated and uncorrelated components, $\mathbf{Q_{0,i}} = Q_0^{(c)}\mathbf{I_M} + diag\{Q_{0,i}^{(u)}\}$, the received signal is:

$$\mathbf{r} = \mathbf{G_{zf}}\mathbf{Q_{0,i}}\mathbf{Hy} + \mathbf{G_{zf}}(\mathbf{ICI} + \mathbf{w}) \tag{9}$$

It is important to study how the ICI behaves under beamforming. Consider just the ICI inflicted by subcarrier $k+1$ onto subcarrier k:

$$ICI_{k+1\to k} = \mathbf{G_{zf}^{(k)}} diag\{Q_{1,i}\}\mathbf{H^{(k+1)}}\mathbf{y_{k+1}} \tag{10}$$

The ICI experiences an effective phase noise channel composed of the true channel $\mathbf{H^{(k+1)}}$ scrambled by the uncorrelated ICI coefficients $\{Q_{1,i}\}$. This has two important effects. First, it decorrelates the ICI across the array. Second, it causes the ICI channels to lose orthogonality, creating *inter-user ICI* in addition to the self-ICI. Therefore, the per-user SINR with genie-aided CPE estimation is

$$SINR^{(g)} = \frac{M\mathbb{E}[|Q_0|^2]}{K\mathbb{E}[|ICI|^2] + \frac{\sigma^2}{E_s}} \tag{11}$$

This result is identical to (5) apart from an array gain due to averaging of thermal noise and ICI. Equivalently, the effective array-level LO experiences an M/K-fold reduction in its far-out phase noise compared to the LO at any element.

The gain in the signal to ICI ratio comes at the cost of an M-fold increase in frequency generation power, since the array possesses M high-performance VCOs. However, it is possible to save power by relaxing the phase noise specification of each element while exploiting the array gain to achieve high array-level performance. Each element's ICI can be increased proportionally with M/K to maintain constant total ICI energy $E_{ICI,0}$:

$$\mathbb{E}[|ICI|^2] = \tfrac{M}{K} \times E_{ICI,0} \tag{12}$$

$$\mathbb{E}[|Q_0|^2] = 1 - \tfrac{M}{K} \times E_{ICI,0} \tag{13}$$

This can be achieved by increasing the VCO's phase noise by 3dB for every factor of 2 in M/K, reducing the VCO power consumption proportionally.

Phase noise scaling maintains constant signal to ICI ratio over a wide range of array sizes, up to a limit. Since the oscillator has unit energy, when the ICI is a sizeable fraction of the total energy, $\mathbb{E}[|Q_0|^2]$ will begin to deviate significantly from unity. Nevertheless, the scaling range is fairly wide in practice (Section V). This idea is confirmed by Fig. 3, which shows that ICI power scales linearly with the VCO phase noise over a large range of phase noises. Only at very high phase noise levels does the CPE energy begin to drop.

IV. CPE Estimation in MU-MIMO Arrays

In practice, beamforming must precede CPE estimation in order to suppress MU interference. Moreover, there are several reasons to prefer using only a global CPE estimator. First, this exploits the full array gain to maximize the estimator's SINR. Second, it eliminates the complexity and overhead of multiple pilot tracking loops. In this section we study the use of global pilot tracking with both synchronous (where all transceivers share the same XO reference) and asynchronous (where each element has its own XO) modes of operation.

A. Asynchronous Array

In the asynchronous mode, the phase noise is fully uncorrelated between elements. The CPE at each element can be expressed in terms of its mean and deviation: $Q_{0,i}^{(u)} = \mathbb{E}[Q_0^{(u)}] + \delta_i$. It is reasonable to assume that the array possesses statistical phase noise information but cannot track the instantaneous deviations. After beamforming, the received signal is:

$$\mathbf{r} = M\mathbb{E}[Q_0^{(u)}]\mathbf{y} + \mathbf{G_{zf}}(diag\{\delta_i\}\mathbf{Hy} + \mathbf{ICI} + \mathbf{w}) \tag{14}$$

The mean CPE can be found as follows:

$$\mathbb{E}[Q_0^{(u)}] = \mathbb{E}[\sum_{n=1}^{N_{sc}} \frac{1}{N_{sc}} e^{j\theta[n]}] = \mathbb{E}[e^{j\theta[n]}] = e^{-\frac{\sigma_\theta^2}{2}} \tag{15}$$

where σ_θ^2 is the variance of the phase. The last step uses the characteristic function of the phase noise, derived in [16]. When exploiting VCO phase noise scaling, $\sigma_{VCO}^2 = (M/K)\sigma_{VCO,0}^2$.

Using (1), the per-user SINR *as a function of time index* n is

$$SINR^{(a)} = \frac{Me^{-n\sigma_{ref}^2} e^{-\frac{M}{K}\sigma_{VCO,0}^2}}{\mathbb{VAR}[\mathbf{I_{eff}}] + K\mathbb{E}[|ICI|^2] + \frac{\sigma^2}{E_s}} \tag{16}$$

where $\mathbf{I_{eff}} = \mathbf{G_{zf}} diag\{\delta_i\}\mathbf{Hy}$.

There are two key effects are in the asynchronous array. First, we observe the well-known channel aging effect whereby the SINR degrades over time [18]–[21]. This occurs because if the phase drifts independently at each element, it causes the channel to change with time and imposes a limit on the coherence time. Second, when phase noise scaling is used the signal energy loss is exacerbated. By deliberately increasing the phase variance, the received signals are progressively decorrelated, causing a loss in signal energy.

B. Synchronous Array

In the ideal synchronous array, the CPE is fully correlated and ICI fully uncorrelated. The beamformed receive signal is

$$\mathbf{r} = MQ_0^{(c)}\mathbf{y} + \mathbf{G_{zf}}(\mathbf{ICI} + \mathbf{w}) \tag{17}$$

The global effective CPE is simply $Q_0^{(c)}$, identical to the per-element CPE, so the synchronous per-user SINR is

$$SINR^{(s)} = \frac{M\mathbb{E}[|Q_0^{(c)}|^2]}{K\mathbb{E}[|ICI|^2] + \frac{\sigma^2}{E_s}} \tag{18}$$

There are three clear results. First, the synchronous SINR is identical to the genie-aided case (11) despite the use of only a *single, global pilot tracking loop*. In other words, by ensuring that the CPE is the same at every element, there is no penalty in estimating and tracking it after beamforming[1]. Second, the performance depends only on the *array-level* SINR and is therefore robust to phase noise scaling. Contrasting with (16), it is clear that synchronous operation is a key enabler of phase noise scaling. Third, the channel aging effect is completely

[1]In fact, the practical performance of global pilot tracking is superior since the estimator SINR is enhanced by the array gain.

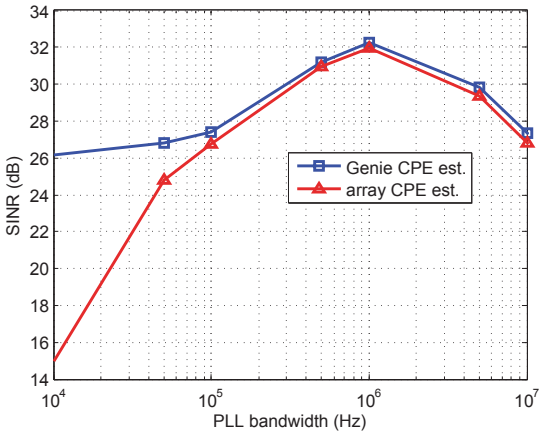

Fig. 4. SINR as a function of PLL bandwidth for a 64-element array with B_{sc} of 624kHz and effective phase noise of -100dBc/Hz at 1MHz offset.

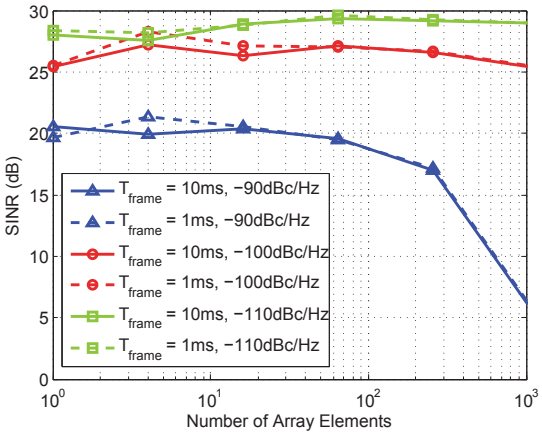

Fig. 6. Simulated SINR with global CPE estimation schemes in the synchronous array, for different frame lengths. $B_{PLL} = 500kHz$, $B_{sc} = 624kHz$, and $N_{sc} = 64$

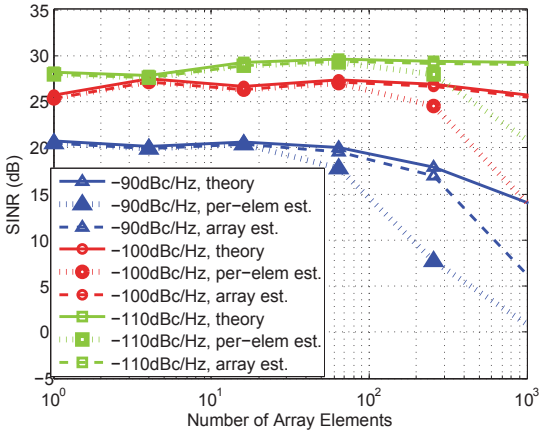

Fig. 5. Simulated and predicted SINR for synchronous array with various levels of phase noise. $B_{PLL} = 500kHz$, $B_{sc} = 624kHz$, and $N_{sc} = 64$.

eliminated since the LO phase drift is the same at every element and is tracked at the array level.

How can this level of CPE correlation be achieved? As discussed in Section II-D and Fig. 2, the PLL's action ensures that CPE is dominated by the reference's phase noise. Consequently, the mere fact of distributing a common reference to every transceiver ensures a relatively high level of CPE correlation while keeping ICI correlation low. Additionally, the PLL bandwidth can be optimized to maximize CPE correlation and minimize ICI correlation:

$$CPE_{uncorr} = \int_{-\infty}^{\infty} S_v(\omega)|H_{N_{sc}}(\omega)|^2|H_{VCO}(\omega)|^2 d\omega$$
$$ICI_{corr} = \int_{-\infty}^{\infty} S_v(\omega)(1-|H_{N_{sc}}(\omega)|^2)|H_{ref}(\omega)|^2 d\omega \quad (19)$$

where H_{VCO} is the VCO transfer function and H_{ref} is the reference transfer function to the PLL output. The PLL bandwidth can be found by minimizing the appropriate cost function $C(CPE_{uncorr}, ICI_{corr})$. In practice, this is difficult to solve analytically; however, numerical analysis reveals that the optimal B_{PLL} is close to B_{sc}.

V. Simulation Results

To validate these reults, numerical simulations were conducted in both synchronous and asynchronous modes. Random phase noise sequences were generated for the reference and VCO and filtered through an s-domain model of the PLL. The data sequence consisted of 1000 802.11n-like symbols with 20MHz bandwidth, 64 subcarriers, and 16-QAM data symbols and BPSK pilots, at pilot density of $1/8$. All phase noise specifications are quoted at 1MHz offset from the carrier.

Fig. 4 shows how the SINR varies with PLL bandwidth for the single-user synchronous array with $B_{sc} = 624kHz$. There is a clear optimum close to the subcarrier bandwidth; furthermore, this optimum is quite shallow, allowing flexibility in choosing B_{PLL} to meet other design considerations. For low PLL bandwidth, the majority of the CPE is uncorrelated so the SINR approaches that of an asynchronous array (16) and post-beamforming CPE estimation performs very poorly. On the other hand, when the PLL bandwidth is large, the reference phase noise adds correlated ICI which cannot be averaged.

Fig. 5 shows the SINR for a single-user, synchronous array with $B_{PLL} = 500kHz$, $B_{sc} = 624kHz$ as the per-element phase noise is scaled with M. The beamformed SNR is constant at 30dB. The simulated results are compared with the predictions of (18). This figure reveals two key results. First, phase noise scaling succeeds in maintaining constant array-level performance over a wide range of array sizes. Second, post-beamforming CPE estimation is necessary to enable phase noise scaling, tracking the theoretical result closely. In contrast, the performance of the per-element CPE estimation degrades rapidly due to the low SINR at each element.

Fig. 6 shows the SINR for the same single-user synchronous array, with frame length of 1ms (100 OFDM symbols) or 10ms (1000 symbols). Over several phase noise levels there is no evidence of channel aging. In particular, the SINR for the case of -90dBc/Hz phase noise is heavily phase noise limited, but still shows no evidence of channel aging.

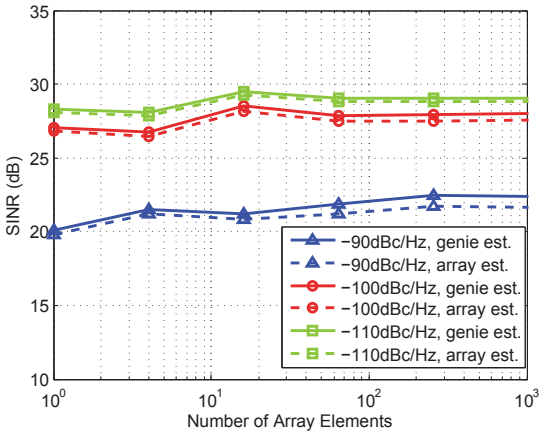

Fig. 7. Simulated SINR with multi-user, synchronous array. $B_{PLL} = 500kHz$, $B_{sc} = 624kHz$, and $N_{sc} = 64$

Fig. 7 shows the SINR for a multi-user synchronous array, using the same system parameters as above. The channel experiences Rayleigh fading, and the array utilizes zero-forcing beamforming to recover the user signals. The ratio M/K is fixed at 1/8 for $M \geq 16$. Each element's phase noise is scaled with the ratio M/K. As expected, the multi-user SINR is the same as the single-user case. Furthermore, the phase noise scaling is able to maintain constant performance across a wide range of array sizes.

VI. Conclusion

In this paper, we investigated the impact of phase noise on MIMO-OFDM arrays with realistic frequency generation. By exploiting the averaging of uncorrelated phase noise, the phase noise specification at each VCO can be relaxed proportionally to the array size. A key enabler of phase noise scaling is the introduction of correlated low-frequency phase noise at each element. This allows the use of a global pilot tracking loop, which is robust to phase noise scaling and maintains high performance across a wide range of array sizes. Additionally, this type of correlation eliminates the phase noise-induced channel aging effect. This level of phase noise correlation can be introduced by distributing a shared low-frequency reference and optimizing each element's PLL bandwidth.

The phase noise scaling breaks down when the oscillator's linewidth is so large that the energy in the CPE decreases. In this regime, it may be advantageous to estimate and correct for higher-order components of the phase noise process. This may favor the use of single-carrier transmissions with time-domain pilots to reduce the complexity of phase noise tracking.

Acknowledgment

This work was supported in part by NSF ECCS-1232318. The authors also acknowledge the students, staff, faculty, and sponsors of the Berkeley Wireless Research Center.

References

[1] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What will 5G be?" *IEEE J. Sel. Areas Commun.*, vol. 32, pp. 1065–1082, Jun. 2014.

[2] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 3590–3600, Nov. 2010.

[3] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up MIMO: Opportunities and challenges with very large arrays," *IEEE Signal Process. Mag.*, vol. 30, pp. 40–60, Jan. 2013.

[4] E. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, pp. 186–195, Feb. 2014.

[5] H. T. Friis, "A note on a simple transmission formula," *Proc. IRE*, vol. 34, pp. 254–256, May 1946.

[6] E. Björnson, M. Matthaiou, and M. Dèbbah, "Massive MIMO systems with hardware-constrained base stations," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 3142–3146.

[7] E. Björnson, J. Hoydis, M. Kountouris, and M. Dèbbah, "Massive MIMO systems with non-ideal hardware: Energy efficiency, estimation, and capacity limits," *IEEE Trans. Inf. Theory*, vol. 60, pp. 7112–7139, Sep. 2014.

[8] E. Björnson, M. Matthaiou, and M. Dèbbah, "Massive MIMO with non-ideal arbitrary arrays: Hardware scaling laws and circuit-aware design," *IEEE Trans. Wireless Commun.*, vol. 18, pp. 4353–4368, Apr. 2015.

[9] T. Pollet, M. van Bladel, and M. Moeneclaey, "BER sensitivity of OFDM systems to carrier frequency offset and Wiener phase noise," *IEEE Trans. Commun.*, vol. 43, pp. 191–193, Feb. 1995.

[10] P. Robertson and S. Kaiser, "Analysis of the effects of phase noise in orthogonal frequency division multiplex (OFDM) systems," in *Proc. 1995 IEEE Int. Conf. Communications (ICC)*, Jun. 1995.

[11] H. Steendam, M. Moeneclaey, and H. Sari, "The effect of carrier phase jitter on the performance of orthogonal frequency-division multiple-access systems," *IEEE Trans. Commun.*, vol. 46, pp. 456–459, Apr. 1998.

[12] L. Tomba, "On the effect of Wiener phase noise in OFDM systems," *IEEE Trans. Commun.*, vol. 46, pp. 580–583, May 1998.

[13] A. G. Armada, "Understanding the efects of phase noise in orthogonal frequency division multiplexing (OFDM)," *IEEE Trans. Broadcast.*, vol. 47, pp. 153–159, Jun. 2001.

[14] S. Wu and Y. Bar-Ness, "A phase noise suppression algorithm for OFDM-based WLANs," *IEEE Commun. Lett.*, vol. 6, pp. 535–537, Dec. 2002.

[15] ——, "OFDM systems in the presence of phase noise: Consequences and solutions," *IEEE Trans. Commun.*, vol. 52, pp. 1988–1996, Nov. 2004.

[16] D. Petrovic, W. Rave, and G. Fettweis, "Effects of phase noise on OFDM systems with and without PLL: Characterization and compensation," *IEEE Trans. Commun.*, vol. 55, pp. 1607–1616, Aug. 2007.

[17] T. Höhne and V. Ranki, "Phase noise in beamforming," *IEEE Trans. Wireless Commun.*, vol. 9, pp. 3682–3689, Dec. 2010.

[18] A. Pitarokoilis, S. K. Mohammed, and E. G. Larsson, "Uplink performance of time-reversal MRC in massive MIMO systems subject to phase noise," *IEEE Trans. Wireless Commun.*, vol. 14, pp. 711–723, Feb. 2015.

[19] R. Krishnan, M. R. Khanzadi, N. Krishnan, Y. Wu, A. Graell i Amat, T. Eriksson, and R. Schober, "Large-scale analysis of linear massive MIMO precoders in the presence of phase noise," in *Proc. 2015 Int. Conf. Communications Workshops (ICCW)*, May 2015, pp. 1172–1177.

[20] ——, "Linear massive MIMO precoders in the presence of phase noise – a large-scale analysis," *IEEE Trans. Veh. Technol.*, vol. PP, p. 1, Sep. 2015.

[21] M. R. Khanzadi, G. Durisi, and T. Eriksson, "Capacity of SIMO and MISO phase-noise channels with common/separate oscillators," *IEEE Trans. Commun.*, vol. 63, pp. 3218–3231, Sep. 2015.

[22] P. K. Hanumolu, M. Brownlee, K. Mayaram, and U.-K. Moon, "Analysis of charge-pump phase-locked loops," *IEEE Trans. Circuits Syst. I*, vol. 51, pp. 1665–1674, Sep. 2004.

[23] A. Demir, A. Mehrotra, and J. Roychowdhury, "Phase noise in oscillators: A unifying theory and numerical methods for characterization," *IEEE Trans. Circuits Syst. I*, vol. 47, pp. 655–674, May 2000.

[24] A. Mehrotra, "Noise analysis of phase-locked loops," *IEEE Trans. Circuits Syst. I*, vol. 49, pp. 1309–1316, Sep. 2002.

[25] H. Yang and T. L. Marzetta, "Performance of conjugate and zero-forcing beamforming in large-scale antenna systems," *IEEE J. Sel. Areas Commun.*, vol. 31, pp. 172–179, Feb. 2013.