PHASE NOISE AND SPUR REDUCTION IN AN ARRAY OF DIRECT DIGITAL SYNTHESIZERS

BY

THOMAS MATTHEW COMBERIATE

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2010

Urbana, Illinois

Adviser:

Professor Steven J. Franke

ABSTRACT

Many applications, including communications, test and measurement, and radar, require the generation of signals with a high degree of spectral purity. One method for producing tunable, low-noise source signals is to combine the outputs of multiple direct digital synthesizers (DDSs) arranged in a parallel configuration. In such an approach, if all noise is uncorrelated across channels, the noise will decrease relative to the combined signal power, resulting in a reduction of sideband noise and an increase in SNR. However, in any real array, the broadband noise and spurious components will be correlated to some degree, limiting the gains achieved by parallelization. This thesis examines the potential performance benefits that may arise from using an array of DDSs, with a focus on several types of common DDS errors, including phase noise, phase truncation spurs, quantization noise spurs, and quantizer nonlinearity spurs. Measurements to determine the level of correlation among DDS channels were made on a custom 14-channel DDS testbed.

The investigation of the phase noise of a DDS array indicates that the contribution to the phase noise from the DACs can be decreased to a desired level by using a large enough number of channels. In such a system, the phase noise qualities of the source clock and the system cost and complexity will be the main limitations on the phase noise of the DDS array.

The study of phase truncation spurs suggests that, at least in our system, the phase truncation spurs are uncorrelated, contrary to the theoretical prediction. We believe this decorrelation is due to the existence of an unidentified mechanism in our

DDS array that is unaccounted for in our current operational DDS model. This mechanism, likely due to some timing element in the FPGA, causes some randomness in the relative phases of the truncation spurs from channel to channel each time the DDS array is powered up. This randomness decorrelates the phase truncation spurs, opening the potential for SFDR gain from using a DDS array.

The analysis of the correlation of quantization noise spurs in an array of DDSs shows that the total quantization noise power of each DDS channel is uncorrelated for nearly all values of DAC output bits. This suggests that a near N gain in SQNR is possible for an N-channel array of DDSs. This gain will be most apparent for low-bit DACs in which quantization noise is notably higher than the thermal noise contribution.

Lastly, the measurements of the correlation of quantizer nonlinearity spurs demonstrate that the second and third harmonics are highly correlated across channels for all frequencies tested. This means that there is no benefit to using an array of DDSs for the problems of in-band quantizer nonlinearities. As a result, alternate methods of harmonic spur management must be employed.

To my family, for its love and support.

To my girlfriend, for her patience.

To my coworkers, for their gracious assistance.

ACKNOWLEDGMENTS

I would like first to thank Ken O'Haver, Salvador Talisa, and the rest of the Radar and Communication Systems Group at the Johns Hopkins University Applied Physics Laboratory. Specifically, I thank Keir Lauritzen for helping me with my experimental design and for always being willing to spare some time for any questions I had. I also thank Laura Ruppalt for patiently and thoroughly editing this manuscript and providing insightful feedback.

Lastly, I would like to acknowledge the financial support from the American Society for Engineering Education National Defense Science and Engineering Graduate Fellowship.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

1.1 Motivation

The direct digital synthesizer (DDS) is a commonly used device in modern radio frequency (RF) applications. The ability to quickly and directly modify the frequency control word (FCW), k, enables the DDS topology to offer the fastest frequency jumping and the finest frequency tuning resolution of any technology available today in a completely controlled digital environment. As a result, DDSs have found wide application in fields including communications and test and measurement equipment. However, the performance of a DDS can be limited by errors in the signal generation, most notably phase noise and periodic signal generation errors which manifest as spurs in the frequency domain. As has been the tendency in other technology industries, we look to improve the performance of a DDS by placing it in parallel with several other DDSs and aiming to take advantage of the decorrelation of noise across units to improve signal quality.

The purpose of this thesis is to examine the potential performance benefits that may arise from using an array of DDSs, with a focus on several types of common DDS errors. In doing so, this thesis offers and experimentally verifies a concise DDS phase noise model, which takes into account the possibility of combining the output of an array of N identical DDSs. It then goes on to define the mechanisms for three main sources of spurs seen in DDSs: phase truncation spurs caused by the limited size of the look-up table (LUT), quantization noise spurs caused by the limited

precision of the digital-to-analog converter (DAC), and spurs caused by nonlinearities in the DAC. It is our hope that this thesis will serve both as a guide to the advantages and limitations of using a DDS in a real system and as a survey of the potential benefits of using an array of DDSs to improve performance in signal generation systems limited by any one of a number of types of errors.

1.2 Outline

The thesis is organized as follows: Chapter 2 provides an introduction to the basic structure of a DDS and details the specifications of the particular DDSs used for the experiments reported in this thesis. It also includes a basic review of the theory of phase noise and phase noise measurements, along with a brief overview of correlated and uncorrelated power, that provides context for discussions in later chapters. Chapter 3 presents and provides experimental verification for a phase noise model for an array of N identical DDSs, showing that placing multiple DDSs in parallel can significantly reduce the total output signal phase noise. Chapter 4 identifies the generation mechanisms for phase truncation spurs, quantization noise, and quantizer nonlinearity spurs, and provides experimental data that demonstrates their level of correlation in a DDS array. Chapter 5 concludes with general discussion regarding the effectiveness of using an array of DDSs for improving system performance and provides some commentary on possibilities for future work.

CHAPTER 2

BACKGROUND

2.1 Structure of a DDS

The purpose of a direct digital synthesizer (DDS) is to use a single frequency oscillator to output a sinusoid whose frequency can be tuned rapidly and precisely over a wide frequency range. The topology of a conventional direct digital synthesizer (DDS) is shown in Figure 2.1 to consist of five primary elements: a driving clock, a phase accumulator (PA), a look-up table (LUT), a digital-to-analog converter (DAC), and a reconstruction filter [1, 2, 3]. At each clock cycle, the PA, which is effectively a counter, is incremented by an M-bit number k, the frequency control word (FCW). The phase stored by the PA is converted to a corresponding sine-wave amplitude by the phase-to-amplitude converter, often through the use of a simple sine LUT. The digital amplitude value from the LUT is then passed to the DAC, which converts it to an analog output. As the phase is increased by k each subsequent clock cycle, the amplitude output steps through the sine LUT, generating the desired analog sinusoidal signal. The amplitude of this generated signal is set by digitally scaling the input to the DAC or by a physical attenuator at the output, and the signal frequency is tuned by varying k. A larger k results in the PA moving through the period of the LUT more quickly, producing a higher-frequency sinusoid at the output, while a smaller k moves the accumulator through the LUT more slowly, yielding a lower frequency output sine wave. The maximum output frequency of a DDS is determined by its source clock, limited by

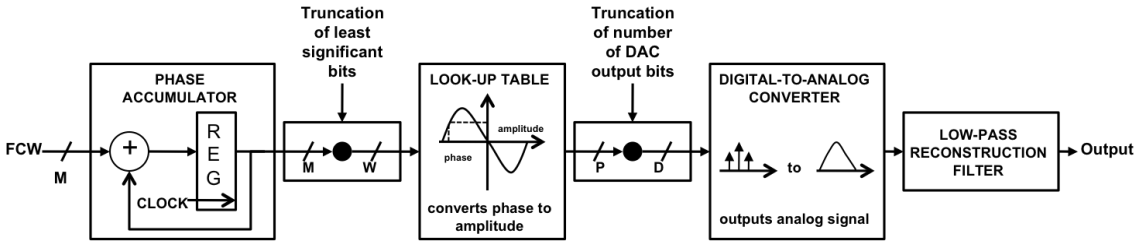

Figure 2.1: Basic structure of a generic direct digital synthesizer.

Nyquist constraints to one half of source frequency, and by the cutoff frequency of the lowpass reconstruction filter at the output.

A functional block diagram of the DDSs used in these experiments is shown in Fig. 2.2 [4]. This particular DDS design implements the PA and the LUT with an FPGA. Because the FPGA runs at half the rate of the source clock, there are two sets of PAs and LUTs whose outputs, even and odd, are alternated by a double data rate (DDR) buffer before being sent to the DAC. The following sections describe both the general purpose and particular specifications of the main components of our DDS units.

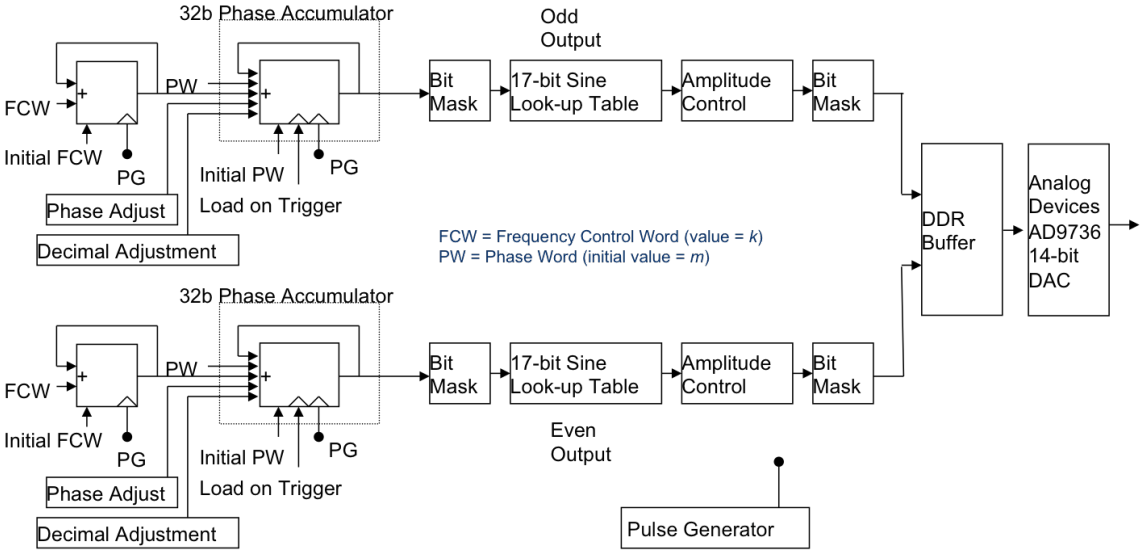

Figure 2.2: Functional block diagram of the DDSs used in this thesis.

2.1.1 Source Clock

The source clock of the DDS is generally a clock with relatively high spectral purity, similar in quality to that which would be used as the local oscillator of a receiver. The source clock used in this particular DDS system was based on a 100 MHz ultra low-phase-noise oven-controlled crystal oscillator (OCXO), which was amplified and multiplied up to 800 MHz using a frequency doubler and quadrupler. All of these components were made by Wenzel Associates, Inc.

2.1.2 Phase Accumulator

The PA is an M-bit overflowing counter that stores the digital phase of the output signal. Its value is added to k and then fed back into the PA on each clock impulse [2]. The 2^M possible values for the PA map to phase values of a sinusoid uniformly distributed from 0 to 2π. As a result, as the register overflows, a new period of the output sinusoid begins. The frequency of the sinusoid output for a given value of k is

$$f_{out} = \frac{k \cdot f_{Clk}}{2^M} \tag{2.1}$$

which holds whenever the Nyquist criterion

$$f_{out} \leq \frac{f_{Clk}}{2} \tag{2.2}$$

is satisfied [5]. As (2.1) suggests, the value of k directly determines the output frequency. Furthermore, by changing the FCW, k, this output frequency can be adjusted from one input clock impulse to the next while still maintaining phase continuity, enabling nearly instantaneous frequency tuning [2]. By setting $k = 1$ in

(2.1), we find the frequency resolution Δf to be [5]

$$\Delta f = \frac{f_{Clk}}{2^M} \tag{2.3}$$

The PAs, LUTs, and the DDR buffer are implemented using Xilinx Virtex 4 field programmable gate arrays (FPGAs). Each FPGA controls two DDS channels. Both the FCW and the PA are $M = 32$ bits long.

2.1.3 Look-Up Table

The LUT converts the phase value stored in the PA into an amplitude. Ideally this operation is given by

$$s(n) = \sin\left(2\pi \frac{k}{2^M} n\right) \tag{2.4}$$

where n is the sample number incremented each clock period ($t = nT_{Ck}$). In practice, the LUT is a read-only memory (ROM) that performs the function in (2.4) to a precision determined by the designer. The spectral error resulting from this finite precision, termed "phase truncation," has been explored [6, 7, 8] for a single DDS and will be further expounded upon for multiple DDS channels in Section 4.1. The output of the LUT then goes through a digital amplitude control. For coarse attenuation, the amplitude bits are right-shifted. For fine attenuation, the amplitude is multiplied by a fractional number.

In our DDSs, the function described in (2.4) is approximated by

$$\begin{aligned} \sin(\alpha+\beta+\chi) &= \sin(\alpha+\beta)\cos(\chi) + \cos(\alpha)\cos(\beta)\sin(\chi) - \sin(\alpha)\sin(\beta)\sin(\chi) \\ &\approx \sin(\alpha+\beta) + \cos(\alpha)\sin(\chi) \end{aligned} \tag{2.5}$$

where both the even and the odd LUTs perform both functions. The most significant bits (MSBs) are input to the $\sin(\alpha+\beta)$ portion of the LUT and the least

significant bits (LSBs) are input to the $\cos(\alpha)\sin(\chi)$ portion of the LUT [9]. This technique results in a huge compression ratio in the size of the LUT [2]. Each LUT is 17 bits long and the sinusoid amplitude entries in the LUT are 14 bits wide. A bit mask at the input of the LUT is included to allow the user to vary the number of bits extracted from the PA, denoted by W, from 1 to 17 [10], a capability utilized in the experiments of Section 4.1 to investigate the effect of LUT entry length on phase truncation errors.

2.1.4 Digital-to-Analog Converter

The output of the LUT is connected to a D-bit DAC, which generates an analog value corresponding to the D MSBs of the LUT output amplitude. Each of our DDSs uses a 14-bit Analog Devices 9736 DAC. Another bit mask at the input to the DAC allows the user to adjust D, the number of bits in the DAC, from 1 to 14 by truncating the appropriate number of LSBs. This added capability is utilized in the experiments of Section 4.2 to investigate the effect of number of DAC bits on quantization noise.

2.1.5 Reconstruction Filter

A lowpass reconstruction filter smooths the output of the DAC and limits the maximum output frequency to just under half of the frequency of the source clock in order to satisfy the Nyquist criterion. Typically most filters cut off at a frequency around 40 percent of the source clock frequency to allow for a transition band below the Nyquist frequency [3]. Our DDS systems have 7th order Chebyshev reconstruction filters, which limit the maximum output frequency of each DDS to 360 MHz.

2.1.6 Calibration and Alignment of an Array of DDSs

Many potential applications of parallel DDS arrays require precise phase and amplitude alignment of the individual DDS output signals in order to achieve the maximum array gain in signal quality. In our array, each identical DDS receives a common 800 MHz clock and trigger. The 3 ns rise time of the trigger is long compared to the 1.25 ns clock period, with the result being that the trigger does not reliably occur on the same clock period for every DDS channel. The phase error caused by this can be calibrated out by applying channel-specific increments to each PA. In practice, the DDS channels are powered on one at a time with an arbitrary number of clock cycles in between each channel's starting. All of the channels are then reinitialized to the same starting PA value, and each is connected to an identical analog-to-digital converter (ADC). Because of slight variations in the DDSs, the cables, and the ADCs among the channels, the outputs measured by the ADCs typically do not add coherently at this point. To correct this, we then measure the phases of each DDS's output frequency using the ADCs and increment the PA of each DDS until all of the channels are phase-aligned at the input of the ADC. The output signal amplitudes are measured with the ADCs in similar fashion, and each channel's output is digitally attenuated until all of the channels are amplitude-aligned. This calibration approach achieves excellent phase and amplitude alignment, with signals generally matched to within 170 μrad and 0.1 dB, respectively [10].

2.2 Phase Noise

A significant source of spectral broadening in DDSs is phase noise, a measure of phase instability which can be characterized by the spectral density of phase fluctuations in a signal [11]. Phase noise can degrade performance in systems using

common communication schemes such as orthogonal frequency division multiplexing (OFDM) [12]. The ability to model phase noise accurately allows designers to identify performance bounds on their communications systems.

2.2.1 Definition

In order to define phase noise mathematically, we start with the basic model for the instantaneous output of an oscillator,

$$v(t) = (V_0 + \epsilon(t)) \sin(2\pi\nu_0 t + \phi(t)) \tag{2.6}$$

where V_0 is the nominal amplitude, $\epsilon(t)$ is the deviation from the nominal amplitude, ν_0 is the nominal frequency, and $\phi(t)$ is the phase deviation from the nominal phase $2\pi\nu_0 t$ [11]. Deviations $\epsilon(t)$ and $\phi(t)$ are random variables. From this we define the phase spectrum as

$$S_\phi(f) = \frac{\phi^2_{rms}(f)}{BW} \tag{2.7}$$

where $S_\phi(f)$ has units rad^2/Hz and $\phi_{rms}(f)$ is a root mean square (rms) value in a specific frequency band offset from the nominal frequency. Most literature commonly uses the single sideband phase noise, which is defined as [11]

$$\mathcal{L}(f) = \frac{S_\phi}{2} \tag{2.8}$$

$\mathcal{L}(f)$ is usually expressed in decibels (dB) as $10\log_{10}\mathcal{L}(f)$. For small phase modulations, $\phi(t) \ll 1$ rad, and moderate to large frequency offsets, f, $\mathcal{L}(f)$ is equal to the ratio of the power density in one phase noise modulation sideband per 1 Hz bandwidth to the total signal power. As a result, $\mathcal{L}(f)$ is commonly expressed in units of decibels below the carrier in a 1 Hz bandwidth, abbreviated as dBc/Hz.

Figure 2.3 shows a graphical frequency domain representation of a sinusoid with phase noise.

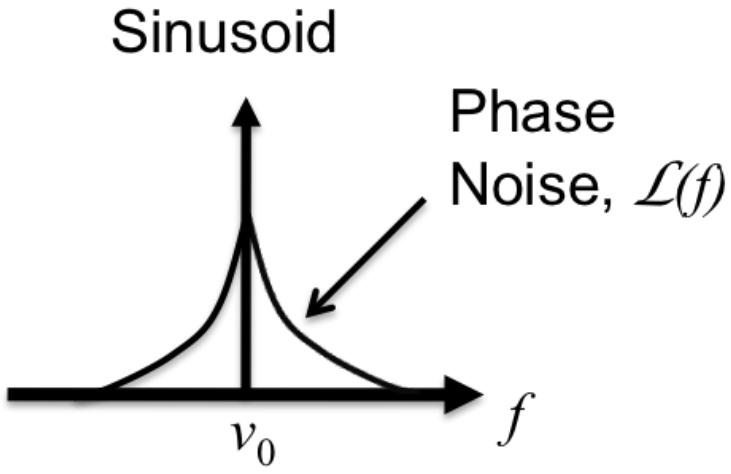

Figure 2.3: Graphical representation of a sinusoid with phase noise in the frequency domain. The phase noise manifests itself as pedestal for the sinusoid.

2.2.2 Effect of Frequency Multiplication

As suggested by the definition of phase noise, it is related to the second moment of the signal's phase variation $\phi(t)$. Therefore, running a sinusoid through an N-times frequency multiplier, as shown in Fig. 2.4, multiplies the phase noise of the input by N^2

$$\begin{aligned} \mathcal{L}_2(f) &= N^2 \mathcal{L}_1(f) \\ \mathcal{L}_2(f)_{[\mathrm{dB}]} &= \mathcal{L}_1(f)_{[\mathrm{dB}]} + 20\log(N)\,\mathrm{dB} \end{aligned} \tag{2.9}$$

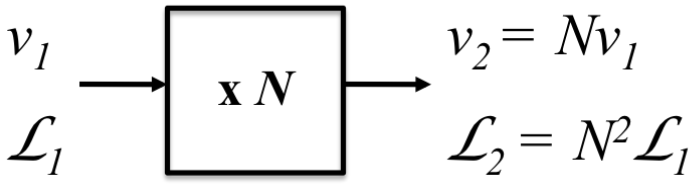

Figure 2.4: Frequency multiplier which outputs a frequency ν_2 equal to N-times its input frequency ν_1.

2.2.3 Phase Noise Measurement Process

A block diagram of the phase noise measurement system used in this thesis is shown in Fig. 2.5. This measurement system uses two sources at the same frequency. In our case, one of the sources has much lower phase noise than the other. We label this source as the reference, REF, and make the reasonable assumption that all of the phase noise measured at the output is contributed by the other source, the device-under-test (DUT). The measurement system relies on a superheterodyne technique in which the two sources are input to a double balanced mixer in quadrature. As a result, the mixer outputs a signal whose power is proportional to the phase noise of the DUT [13]. This output is low-pass filtered, amplified, and input into a spectrum analyzer which performs a fast Fourier transform (FFT) [14]. In order to determine the gain constant of the phase noise detector, we use a calibration approach in which a noise diode injects a well-known amount of broadband phase noise onto the DUT at a power level significantly greater than the phase noise of either source. The output value of this noise is measured with the spectrum analyzer and used to find the gain constant at each offset (or Fourier) frequency. The noise diode is then turned off and the actual phase noise of the DUT is measured [14]. The phase noise detector used was a Femtosecond Systems 1000E which enabled measurements up to 10 MHz offset frequency.

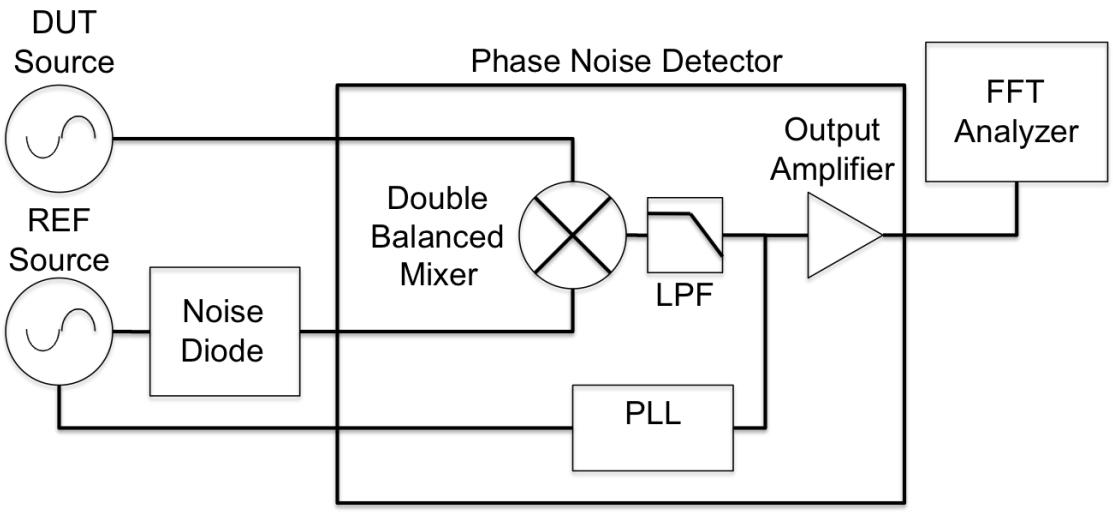

Figure 2.5: Block diagram of the phase noise measurement apparatus used in this thesis.

2.3 Correlated Power Versus Uncorrelated Power

Some of the greatest predicted benefits of DDS arrays are due to the expected decorrelation of noise across DDS channels. To assess the effect of decorrelation among multiple channels, we first consider the sum of two voltages from different channels

$$\begin{aligned} v(t) &= v_1(t) + v_2(t) \\ &= V_1 \cos(\omega t + \theta_0 + \phi_1) + V_2 \cos(\omega t + \theta_0 + \phi_2) \end{aligned} \tag{2.10}$$

applied across a 1 Ω resistor where θ_0 is an arbitrary starting phase common to all channels, and ϕ_i is a random phase error associated with the ith channel. The time-average power dissipated in this resistor is

$$P = \frac{1}{2}V_1^2 + \frac{1}{2}V_2^2 + V_1 V_2 \cos(\phi) \tag{2.11}$$

where $\phi = \phi_2 - \phi_1$. If ϕ_1 and ϕ_2 are independent random variables, uniformly distributed from 0 to 2π, then the signals $v_1(t)$ and $v_2(t)$ are uncorrelated. If $\phi = 0$, the two terms are always in phase. Assuming the channels have equal amplitudes, $V_1 = V_2 = V$, the ratio of the summed fully correlated power to the summed uncorrelated power is

$$\frac{P_{2\,\text{correlated}}}{P_{2\,\text{uncorrelated}}} = \frac{2V^2}{V^2} = 2 \tag{2.12}$$

This can be extended to N channel voltage signals. Assuming the magnitude is V for all signals,

$$\frac{P_{N\,\text{correlated}}}{P_{N\,\text{uncorrelated}}} = \frac{N^2V^2}{NV^2} = N \tag{2.13}$$

In our analysis of the level of correlation among different channels, we use the following definition for fully correlated power:

$$P_{N\,\text{correlated}} = \left(\sqrt{P_1} + ... + \sqrt{P_N}\right)^2 \tag{2.14}$$

which is equivalent to assuming that $\phi = 0$. For uncorrelated power, we use

$$P_{N\,\text{uncorrelated}} = (P_1 + ... + P_N) \tag{2.15}$$

which assumes that ϕ is uniformly distributed. Lastly, for the actual total summed power, we calculate

$$P_{N\,\text{actual}} = (V_1 + ... + V_N)^2 \tag{2.16}$$

which uses the measured values for ϕ to determine the relative level of correlation of the N signals. The correlated and uncorrelated powers provide theoretical bounds for $P_{N\,\text{actual}}$ such that

$$P_{N\,\text{uncorrelated}} \leq P_{N\,\text{actual}} \leq P_{N\,\text{correlated}} \tag{2.17}$$

where the closer $P_{N\,\text{actual}}$ is to $P_{N\,\text{correlated}}$, the more correlated the signals are.

CHAPTER 3

DDS PHASE NOISE

3.1 Complete Model

We present the following phase noise model for the combined output of an array of N identical DDSs, assembled from various reports in the literature [15, 16, 17, 18, 19]:

$$\mathcal{L}_{\mathrm{DDS}}(f, r) = \frac{1}{2} r^2 \cdot \mathcal{L}_{\mathrm{Ck}}(f) + \frac{1}{N} \left(\frac{r}{r_{\mathrm{R}}} \right)^2 \cdot \mathcal{L}_{1/f}(f, r_{\mathrm{R}}) + \frac{1}{N} \kappa(r) \cdot \mathcal{L}_{\mathrm{floor}} \tag{3.1}$$

In this model, the DDS phase noise, $\mathcal{L}_{\mathrm{DDS}}$, consists of contributions from the DDS source clock $\mathcal{L}_{\mathrm{Ck}}$, the internal DDS flicker noise $\mathcal{L}_{1/f}$, and the DDS's DAC noise floor $\mathcal{L}_{\mathrm{floor}}$. Ratio $r = \frac{f_{\mathrm{out}}}{f_{\mathrm{Ck}}}$ relates the DDS output frequency to its source clock frequency, and ratio $r_{\mathrm{R}} = \frac{f_{\mathrm{out,\,R}}}{f_{\mathrm{Ck,\,R}}}$ relates a particular reference output frequency to the source clock frequency. The source clock contributes a component equal to a scaled version of its own inherent phase noise [16], while the flicker and floor components are derived from the DDS circuitry itself and are referenced to a particular frequency defined by r_{R}. Knowledge of the flicker noise contribution at any given output and clock frequencies allows determination of the flicker noise for any other DDS output frequency by scaling by the appropriate r and r_{R} terms [18]. The DAC contributes to the overall white noise floor with a mild frequency-dependence, $\kappa(r)$, which is a weak function of r and is specific to a particular DAC [19]. By combining N identical DDSs in parallel, the uncorrelated

flicker noise and noise floor components are reduced by a factor of N relative to the carrier, while the common clock phase noise component is unchanged [15].

Figure 3.1 shows a typical DDS phase noise plot. The horizontal axis is the offset frequency from the carrier frequency of the source which is typically represented in Hz. The vertical axis is the signal's single-sideband phase noise, $\mathcal{L}(f)$, which is represented in dBc/Hz. At lower frequencies, the $1/f$ flicker noise of the DAC dominates. However, noise decreases 10 dB per frequency decade so that the source clock noise may dominate at higher frequencies. The DAC floor noise also contributes but is typically masked by the source clock phase noise.

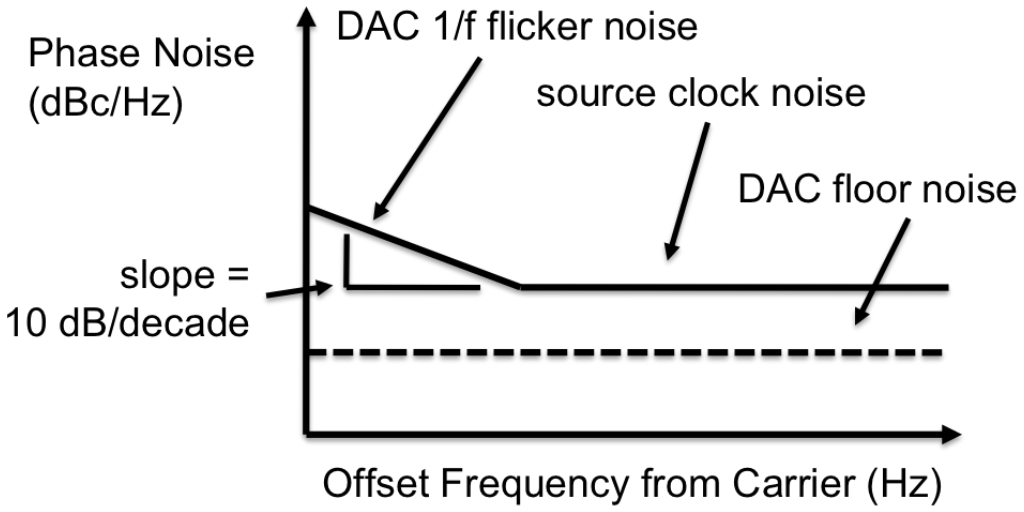

Figure 3.1: Typical Phase Noise Plot of a DDS.

3.1.1 Experimental Approach

In order to experimentally verify the phase noise model of (3.1), we used a test bed consisting of eight of the DDSs described in Section 2.1. Two, four, or eight DDS channels were combined using microwave power combiners after aligning the individual channel phases and amplitudes. The experimental arrangement is shown schematically in Fig. 3.2. Phase noise measurements were made using the phase

detection technique described previously, yielding a measure of the single-sideband phase noise, $\mathcal{L}(f)$, of the desired signal. For the measurements we present in this thesis, a variety of very low phase noise OCXOs were used as reference sources. In order to more easily distinguish the individual contributions of clock, flicker, and floor noise to the overall phase noise of the DDS output, separate cases where a single contributor dominates the DDS phase noise were examined.

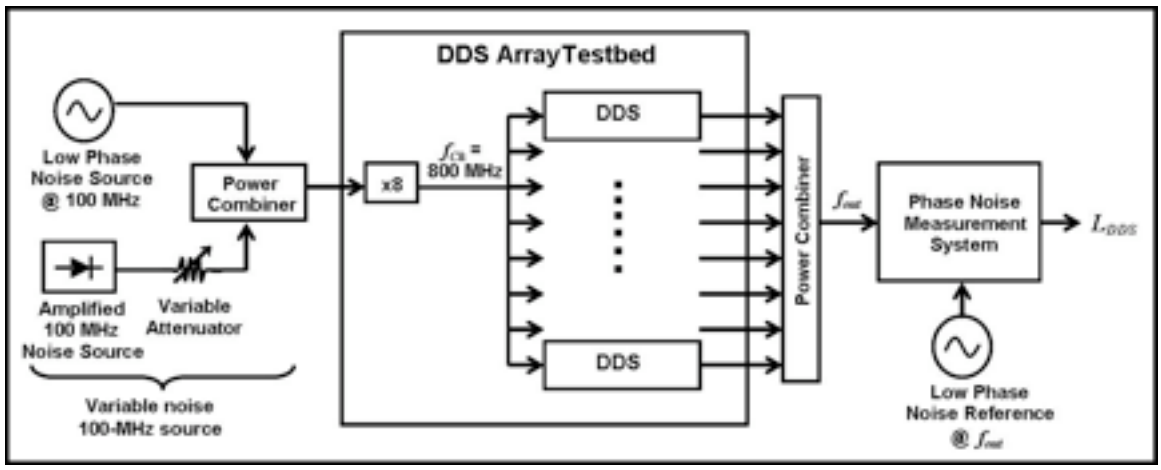

Figure 3.2: Experimental setup used to validate the model in (3.1).

3.2 Source Clock Noise Contribution

3.2.1 Background

In early designs, the phase noise of most DDSs was dominated by the clock oscillator phase noise, making contributions from the other components effectively negligible [19]. Since a DDS is essentially a frequency divider, it divides the phase noise from its source oscillator as described in Section 2.2.2. There is an additional factor of one half, from half of the phase noise of the source clock being converted into amplitude noise and limited in the diode doubler and quadrupler discussed in Section 2.1.1 [16]. Oscillator performance has improved significantly, though, and

the phase noise characteristics of current DDSs are generally dominated by internal flicker noise contributions, masking the effect of the clock noise. However, for systems with lower-quality clocks, clock phase noise can still have a considerable impact on DDS performance.

3.2.2 Verification

Considering the case where the clock phase noise dominates $\mathcal{L}_{\text{DDS}}$, (3.1) suggests that, when $\mathcal{L}_{\text{Ck}} \gg \mathcal{L}_{1/f}$ and $\mathcal{L}_{\text{Ck}} \gg \mathcal{L}_{\text{floor}}$, then

$$\mathcal{L}_{\text{DDS}} \approx \frac{1}{2} r^2 \cdot \mathcal{L}_{\text{Ck}} \tag{3.2}$$

Thus, when a low-stability clock is used, the DDS output phase noise should represent a scaled version of the clock phase noise, with the scaling factor related to the ratio of the output and clock frequencies. Taking advantage of the above simplification, we first aim to verify the relation

$$\mathcal{L}_{\text{DDS}} \propto \mathcal{L}_{\text{Ck}} \tag{3.3}$$

The DDS array test bed that was used is normally driven by a low noise 100 MHz OCXO. This 100 MHz source is multiplied by a factor of $N = 8$ up to the 800 MHz clock signal required to drive the DDSs, as shown in Figure 3.2. As a result, the 800 MHz clock possesses phase noise higher than that of the 100 MHz OCXO by a factor of $N^2 = 64$, or 18 dB [20]. The 100 MHz OCXO phase noise is specified to be -174 dBc/Hz at a 10 kHz offset from the carrier, and, even with the 18 dB increase, the DDS DAC flicker noise dominates the phase noise of the DDSs, preventing the observation of the clock noise dependence of $\mathcal{L}_{\text{DDS}}$. To overcome this limitation, a "noisy clock" was constructed by coupling the low-noise 100 MHz OCXO to an

amplified broadband noise source with a 100 MHz, 3 dB bandwidth. By varying the attenuation on the noise source, the magnitude of the clock phase noise after frequency multiplication, $\mathcal{L}_{\mathrm{Ck}}$, could be arbitrarily tuned and made to be much larger than the DDS flicker phase noise contribution. We first held f_{out}, the output frequency, and the source clock frequency, f_{Ck}, constant at 80 and 800 MHz, respectively, and varied $\mathcal{L}_{\mathrm{Ck}}$ by adjusting the attenuation of the broadband noise source, as described above. As no improvement in phase noise is expected from combining multiple DDS units with the same source clock, these phase noise measurements were taken on a single DDS output. As Fig. 3.3 shows, increasing the attenuation of $\mathcal{L}_{\mathrm{Ck}}$ in increments of 10 dB decreased $\mathcal{L}_{\mathrm{DDS}}$ by the same amounts, thus verifying (3.3). The sloped portions of the spectra at low offset frequencies ($f_{\mathrm{offset}} < 10$ kHz) and high levels of clock noise attenuation (attenuation ≥ 20 dB) indicate regimes in which $\mathcal{L}_{1/f}$ rather than $\mathcal{L}_{\mathrm{Ck}}$ begins to dominate DDS noise. As a result, further decreases in the clock noise do not affect the measured $\mathcal{L}_{\mathrm{DDS}}$ in these regions. The discontinuity seen at the 10 kHz offset frequency is a reproducible artifact generated by the phase noise measurement system.

Figure 3.4 compares the same DDS output phase noise data for the cases when clock noise is most dominant—the 0 dB, 10 dB, and 20 dB attenuation levels—against the measured phase noise of the noisy 100 MHz source used in those instances. As mentioned above, the 100 MHz source is first multiplied to $f_{\mathrm{Ck}} = 800$ MHz, resulting in an 18 dB phase noise increase:

$$\begin{aligned} \mathcal{L}_{\mathrm{Ck}}(f) &= 8^2 \cdot \mathcal{L}_{100\ \mathrm{MHz}}(f) \\ \mathcal{L}_{\mathrm{Ck,\ dB}}(f) &= \mathcal{L}_{100\ \mathrm{MHz,\ dB}}(f) + 18\ \mathrm{dB} \end{aligned} \tag{3.4}$$

with $\mathcal{L}_{\mathrm{dB}}$ defined as $\mathcal{L}_{\mathrm{dB}} = 10 \cdot \log_{10}(\mathcal{L})$. Next, that 800 MHz signal is used to clock the DDS. For an 80 MHz DDS output, the model from (3.2) predicts a DDS phase

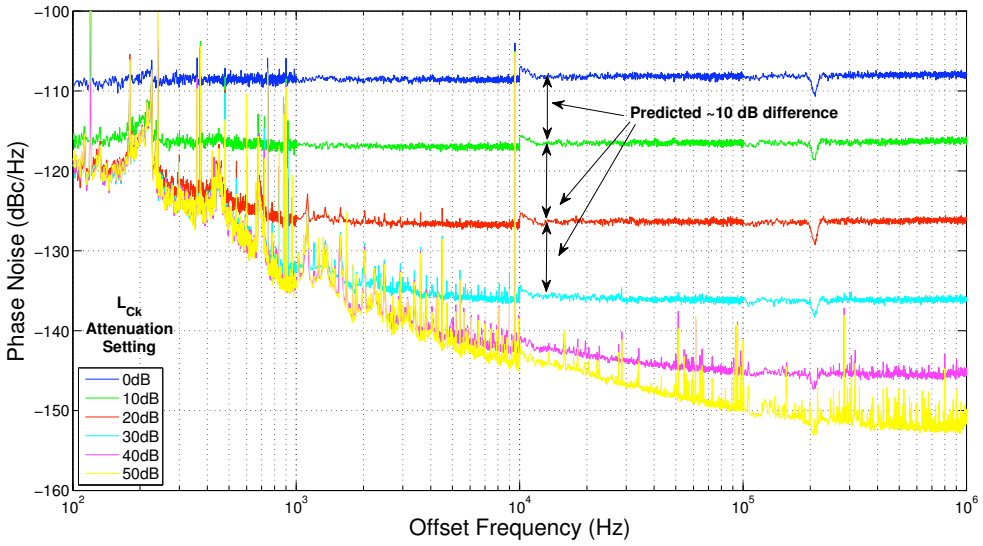

Figure 3.3: Phase noise of 80 MHz DDS outputs generated using the noisy clock configuration of Fig. 3.2 for various noise attenuation settings. The highest levels of clock phase noise are present for the 0 dB attenuation case, while the lowest levels of clock phase noise are present in the 50 dB attenuation case.

noise 23 dB lower than that of the 800 MHz clock:

$$\begin{aligned} \mathcal{L}_{\text{DDS, 80MHz}}(f) &= \frac{1}{2}\left(\frac{80}{800}\right)^2 \cdot \mathcal{L}_{\text{Ck}}(f) \\ \mathcal{L}_{\text{DDS, 80 MHz, dB}}(f) &= \mathcal{L}_{\text{Ck, dB}}(f) - 23\,\text{dB} \end{aligned} \tag{3.5}$$

As a result, the 80 MHz DDS output signal is expected to have phase noise 5 dB lower than that of the original 100 MHz source

$$\begin{aligned} \mathcal{L}_{\text{DDS, 80MHz, dB}}(f) &= \mathcal{L}_{\text{100MHz, dB}} + 18\text{ dB} - 23\text{ dB} \\ &= \mathcal{L}_{\text{100MHz, dB}} - 5\text{ dB} \end{aligned} \tag{3.6}$$

This predicted 5 dB decrease in phase noise is demonstrated clearly in Fig. 3.4 for all three source noise levels where $\mathcal{L}_{\text{Ck}}$ dominates DDS phase noise. Further validation is shown in Fig. 3.5 for the case where the output frequency of the DDS

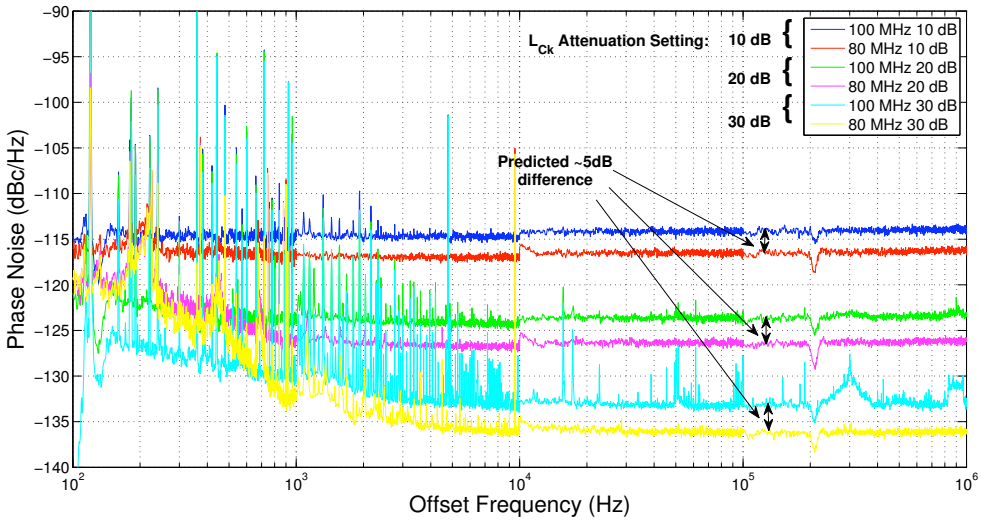

Figure 3.4: Phase noise of 80 MHz DDS outputs generated using the noisy clock configuration of Fig. 3.2 plotted alongside the corresponding variable noise 100 MHz source phase noise. For all three cases, the 80 MHz DDS outputs exhibit phase noise about 5 dB less than that of the 100 MHz source.

is set to 100 MHz. Using the same reasoning expressed in (3.4) through (3.6), but instead substituting 100 MHz for f_{out}, the expected 100 MHz DDS phase noise output is 3 dB lower than the variable noise 100 MHz source in regions where the source dominates. This decrease is clearly shown for the cases in Fig. 3.5, where the DDS phase noise is measured to be approximately 2 to 3 dB lower than the corresponding 100 MHz source noise.

Finally, the DDS output frequency was increased to 200 MHz and the resultant phase noise compared to that measured for the $f_{\text{out}} = 100$ MHz case. According to the model, the ratio of the phase noise at $f_{\text{out}} = 200$ MHz to that at $f_{\text{out}} = 100$ MHz is:

$$\frac{\mathcal{L}_{\text{DDS, 200 MHz}}}{\mathcal{L}_{\text{DDS, 100 MHz}}} = \frac{f^2_{\text{out, 200 MHz}}}{f^2_{\text{out, 100 MHz}}} = \frac{200^2}{100^2} = 4 \approx 6 \text{ dB} \tag{3.7}$$

The expected 6 dB difference is clearly seen in Fig. 3.6 for three different levels of

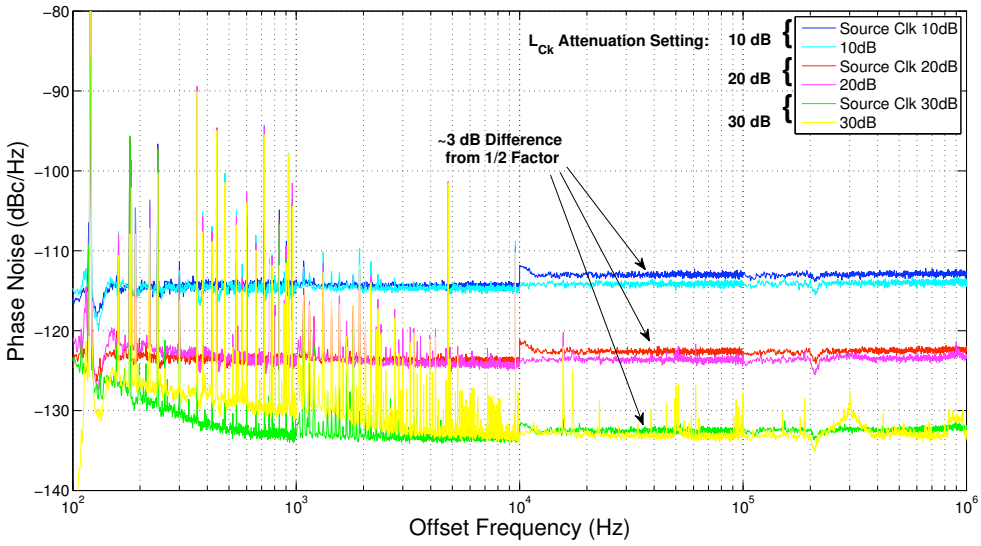

Figure 3.5: Phase noise of 100 MHz DDS outputs generated using the noisy clock configuration of Fig. 3.2 plotted alongside the corresponding variable noise 100 MHz source phase noise. For all three cases, the 100 MHz DDS outputs exhibit phase noise about 3 dB less than that of the 100 MHz source.

clock noise. Taken together, the results presented in Figs. 3.3-3.6 validate the $\mathcal{L}_{\mathrm{Ck}}$ dependence of DDS output phase noise for a single DDS unit captured by the proposed model.

In the case of multiple DDS units, if the DDS bank is driven by a common clock, the clock-noise contributions of each DDS unit are fully correlated and no improvement in overall output phase noise is expected by increasing the number of DDSs. As a result, the clock-dominated DDS phase noise measured for multiple combined outputs should be the same as that measured for a single DDS output. In order to confirm this, the phase noise of one, two, and four phase-aligned and combined channels was measured at $f_{\mathrm{out}} = 100$ MHz. Appropriate amplification and attenuation were used so as to normalize the input power to the phase noise measurement system to approximately 10 dBm in all cases for optimum phase noise measurement system performance. The results of these measurements are plotted in

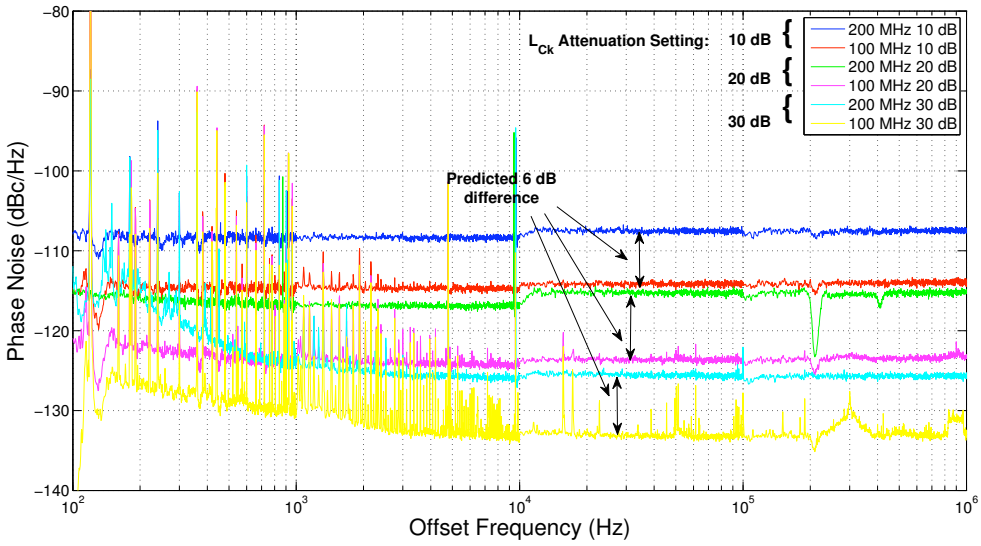

Figure 3.6: Comparison of the phase noise of the DDS outputs at 200 and 100 MHz. The 6 dB difference expected due to the r scaling factor is evident for all three levels of clock noise utilized.

Fig. 3.7 for two levels of source clock phase noise. As the figure shows, increasing the number of DDSs does not appreciably alter the DDS phase noise in regions dominated by the clock contribution.

3.3 Flicker Noise Contribution

3.3.1 Background

By definition, flicker noise has a $1/f$ character, which means that its magnitude drops 10 dB per decade of frequency [21]. Flicker noise in DDSs is the result of internal sources of noise in the biasing circuit and the switching transistors in the DAC [18]. The level of this flicker noise depends on the statistics of the noise processes in the particular DAC used in the DDS. Once the flicker noise is measured at a single reference frequency, the flicker noise for other frequencies can be calculated from this reference flicker noise by the frequency multiplication process

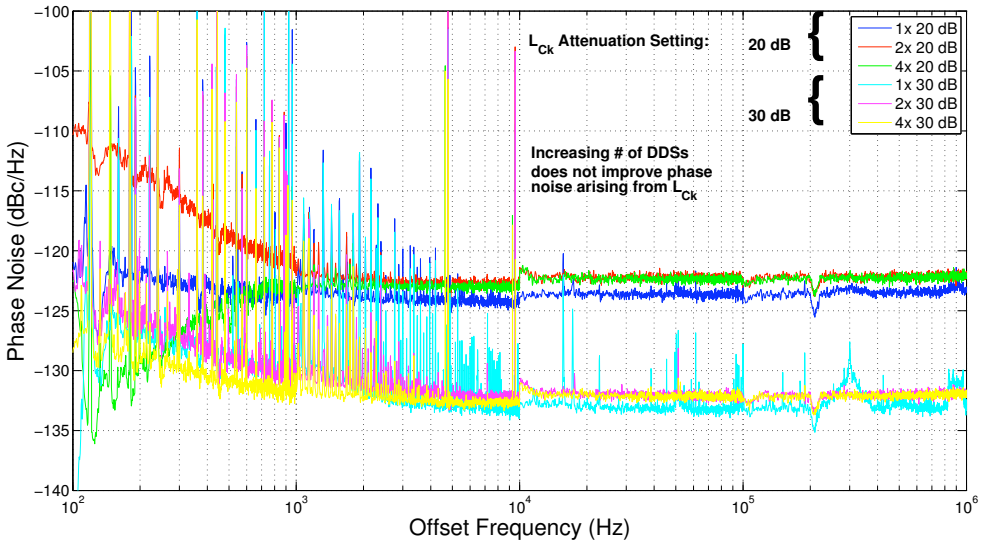

Figure 3.7: Comparison of the phase noise output of the DDS for various numbers of phase-aligned and combined channels generated using two different levels of attenuation of 100 MHz source phase noise. For a given clock phase noise, increasing the number of DDSs in parallel does not improve the combined output phase noise because the common clock noise dominates $\mathcal{L}_{\mathrm{DDS}}$, as predicted by (3.2).

described in Section 2.2.2.

3.3.2 Verification

This work next considers the case where flicker noise from the DAC dominates the total phase noise of the DDS array, that is $\mathcal{L}_{1/f} \gg \mathcal{L}_{\mathrm{Ck}}$ and $\mathcal{L}_{1/f} \gg \mathcal{L}_{\mathrm{floor}}$. Equation (3.1) becomes

$$\mathcal{L}_{\mathrm{DDS}} \approx \frac{1}{N}\left(\frac{r}{r_{\mathrm{R}}}\right)^2 \cdot \mathcal{L}_{1/f} \tag{3.8}$$

where $\mathcal{L}_{1/f}$ is the flicker noise measured at the reference frequency ratio r_R. Thus, for the case when DAC flicker noise dominates, the DDS output noise will be directly proportional to $\mathcal{L}_{1/f}$ with a scaling factor dependent on the number of DDSs present and the output and clock frequencies. Taking advantage of this simplification, the phase noise improvement expected by implementing an array of

parallel DDSs was verified first, specifically:

$$\mathcal{L}_{\mathrm{DDS}} \propto \frac{1}{N}\mathcal{L}_{1/f} \tag{3.9}$$

The variable noise source was replaced with the original low noise 100 MHz OCXO and the phase noise of N=1, 2, 4, and 8 phase-aligned and combined channels were measured at $f_{\mathrm{out}} = 100$ MHz. As Figure 3.8 shows, each doubling of the number of channels decreased the phase noise by 3 dB, in accordance with the expected $1/N$ scaling. The frequency-dependent scaling of the flicker noise term can be verified by

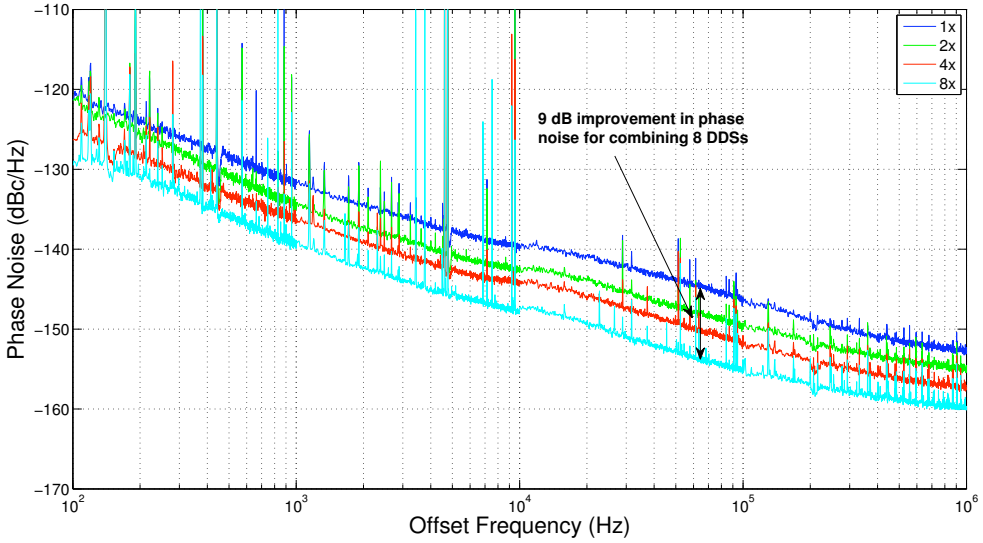

Figure 3.8: Comparison of phase noise output of the DDS at f_{out}=100 MHz for N=1, 2, 4, and 8 phase-aligned and combined channels. The expected $1/N$ decrease in phase noise for N-DDS arrays is clearly evidenced.

comparing the DDS phase noise at different output frequencies.

According to (3.8), for a given flicker-dominated DDS-array, the phase noise characteristics $\mathcal{L}_{\mathrm{DDS1}}$ and $\mathcal{L}_{\mathrm{DDS2}}$ recorded at output frequencies f_{out1} and f_{out2}, respectively, are related by

$$\mathcal{L}_{\mathrm{DDS1}} = \left(\frac{f_{\mathrm{out1}}}{f_{\mathrm{out2}}}\right)^2 \cdot \mathcal{L}_{\mathrm{DDS2}} \tag{3.10}$$

which assumes a common clock frequency for both measurements. Selecting frequencies $f_{\text{out1}} = 80$ MHz and $f_{\text{out2}} = 100$ MHz yields

$$\begin{aligned} \mathcal{L}_{\text{DDS, 80 MHz}} &= \left(\frac{80}{100}\right)^2 \cdot \mathcal{L}_{\text{DDS, 100 MHz}} \\ \mathcal{L}_{\text{DDS, 80 MHz, dB}} &= \mathcal{L}_{\text{DDS, 100 MHz, dB}} - 2 \text{ dB} \end{aligned} \tag{3.11}$$

for any N. Figure 3.9 shows $\mathcal{L}_{\text{DDS, 80 MHz, dB}}$ plotted alongside ($\mathcal{L}_{\text{DDS, 100 MHz, dB}} - 2$ dB) for varying values of N. The phase noise curves measured at 80 MHz overlap nearly exactly with the shifted 100 MHz phase noise curves, validating the predicted frequency scaling of the flicker noise contribution for all combinations of DDSs measured. Deviations in the measured phase noise at low offset frequencies ($f_{\text{offset}} < 1$ kHz) likely arise from differences in the phase lock-loops used in the phase noise measurement setup for the 80 MHz and 100 MHz oscillators.

Further verification of the flicker noise scaling is seen by comparing phase noise measurements taken at $f_{\text{out1}} = 200$ MHz to the reference measurements at $f_{\text{out2}} = 100$ MHz. According to (3.10), the phase noise relationship between these two frequencies is predicted to be

$$\begin{aligned} \mathcal{L}_{\text{DDS, 200 MHz}} &= \left(\frac{200}{100}\right)^2 \cdot \mathcal{L}_{\text{DDS, 100 MHz}} \\ \mathcal{L}_{\text{DDS, 200MHz, dB}} &= \mathcal{L}_{\text{DDS, 100MHz, dB}} + 6 \text{ dB} \end{aligned} \tag{3.12}$$

Similar to Figure 3.9, Figure 3.10 plots the measured DDS phase noise at $f_{\text{out1}} = 200$ MHz alongside ($\mathcal{L}_{\text{DDS, 100 MHz, dB}} + 6$ dB). While the overlap is not quite as good as that seen for the 100 MHz and 80 MHz outputs, the measured values still are within about 2-3 dB of the values predicted by frequency scaling, further confirming the validity of the proposed model. The source of the increased deviation is currently unknown, but may be related to the stability of the 200 MHz reference oscillator or the phase-lock technique implemented in the measurement system.

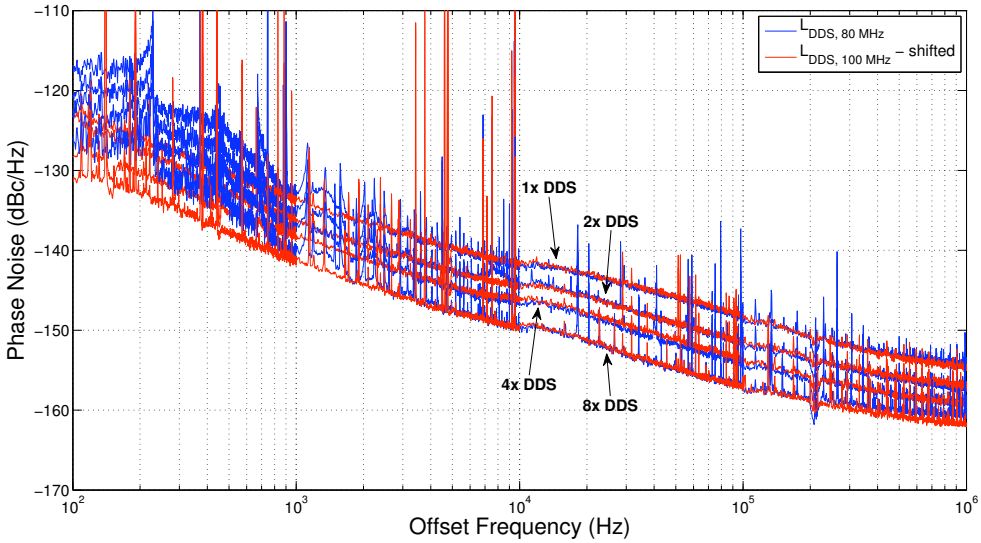

Figure 3.9: Plot of phase noise of DDS output at 80 MHz (blue curves) and DDS output at 100 MHz shifted downward by 2 dB (red curves). The two sets of measurements overlap very well, confirming the expected frequency scaling of the $\mathcal{L}_{1/f}$ contribution to DDS phase noise.

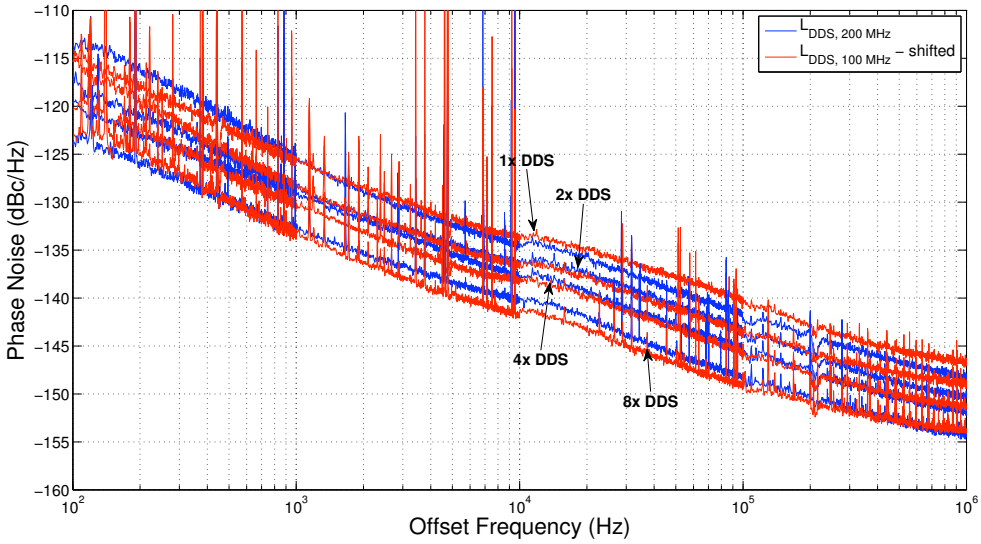

Figure 3.10: Plot of phase noise of DDS output at 200 MHz (blue curves) and DDS output at 100 MHz shifted upward by 6 dB (red curves).

3.4 Floor Noise

The floor noise arises from sources of thermal noise in the DAC, including resistors, transistors, and any other lossy components [19]. In the DDS array we used, the DAC floor noise contribution was negligible at the offset frequencies up to 10 MHz measured, and, unlike the clock noise, it could not easily be artificially increased to dominate the flicker noise. As a result, its contribution to the overall DDS noise could not be verified with the present measurement capabilities. However, $\mathcal{L}_{\text{floor}}$ is expected to be an identical but uncorrelated noise process among multiple DDS units, and we expect the noise floor contribution to $\mathcal{L}_{\text{DDS}}$, like the uncorrelated flicker-contribution, to scale as $1/N$ in a DDS array [15].

3.5 Conclusion

In summary, this chapter has presented a concise, usable model for the phase noise of an N-DDS parallel array. The expected dependence and frequency scaling of the output phase noise on clock and DAC flicker noise contributions has been experimentally verified using a custom-designed DDS test bed. Furthermore, we have demonstrated that combining multiple DDSs yields the predicted $1/N$ phase noise improvement in flicker noise contribution, but has no effect when the noise is dominated by the common clock contribution. The experimental validation of the proposed phase noise model suggests its utility in the design and analysis of systems requiring DDS waveform generation and confirms the effectiveness of parallel DDS arrays for high-performance agile frequency synthesis.

CHAPTER 4

SPURIOUS ERRORS

Because of many qualities inherent to its architecture, a DDS does not generate perfect sinusoids. Therefore, there is always some error between the actual output and an ideal output. According to Fourier theory, any periodic error appears as a sum of delta functions in the frequency domain. These delta functions are referred to as "spurs." In this section, we investigate the level of correlation of three different types of DDS output spurs in a DDS array: phase truncation spurs, quantization noise spurs, and quantizer nonlinearity spurs.

Phase truncation spurs are caused by the truncation of the LSBs of the PA, which is done to reduce the LUT to a manageable size. Quantization noise spurs are caused by the error incurred by rounding the DAC input signal to the limited number of output states in the DAC. Quantizer nonlinearity spurs arise from the error between the actual output levels in a DAC and the ideal output levels.

The mechanisms for generating all three of these error types are deterministic. Assuming that each DDS in an array is architecturally identical, is fed the same FCW, is initialized to the same starting PA value, and has its outputs phase-aligned, the spur-generating errors occur in an identical manner, producing identical spurs in each channel. Therefore, all of the spurs should sum coherently at the output, resulting in no net reduction in spur magnitude due to DDS parallelization. The following sections examine individually each of these spurs and their level of correlation in our DDS array. A more detailed introduction to each error type is provided, as well as the results from our testbed measurements.

4.1 Phase Truncation Spurs

Phase truncation spurs are the primary source of spectral impurity inherent to direct digital synthesis. The other significant errors are a result of the DAC and are not specific to the DDS architecture itself. As was noted, the mechanism that generates the phase truncation spurs is deterministic, and thus theoretically the spurs should be perfectly correlated at the output of a DDS array; consequently, there should be no spur-free dynamic range (SFDR) gain from combining multiple channels in an array.

4.1.1 Spur Origin

In practical DDS designs, the LUT described in (2.4) is implemented with a finite-sized ROM with limited output precision. Often several of the LSBs of the PA are truncated to prevent the size of the LUT from becoming unwieldy. For example, the DDS architecture used in this thesis has a 32-bit PA. Connecting this to an equally sized LUT would require a 32-bit ROM. Each entry in the ROM is 14 bits so as to use the entire dynamic range of the DAC. The resulting ROM would be 7 gigabytes in size, which is prohibitively expensive for most systems. Truncating the 15 LSBs of the PA reduces the size by a factor of 2^{15} to 224 kilobytes, a much more reasonable and affordable size. Further steps utilize the symmetry of a sinusoid to reduce the size of the LUT by a factor of four. Although this phase truncation is practical, it creates a periodic error that manifests itself as spurs in the frequency domain [9]. Figure 4.1 shows an example of a phase-truncated signal. The limited precision of the entries of the LUT also causes quantization spurs [22].

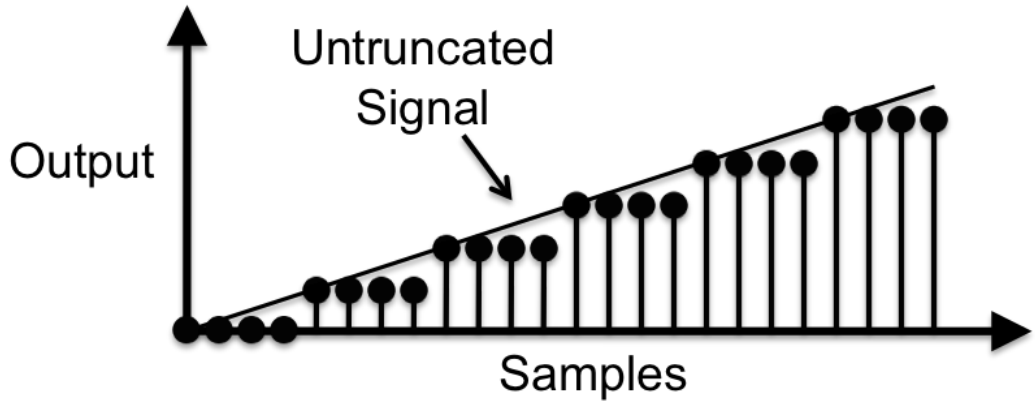

Figure 4.1: Example of the output of the LUT for a phase-truncated signal.

4.1.2 Phase Truncation Spectrum

Consider a DDS with an M-bit FCW having value k and a W-bit LUT as shown in Fig. 4.2. We then define the number of active bits, R, to be the number of bits M minus the number of trailing zeros in the FCW. The FCW and the number of truncated bits, $R - W$, determine the characteristics of the phase truncation spurs. Hence, an M-bit FCW whose $M - R$ LSBs are zeros is effectively the same as an R-bit FCW and will be treated as such in the following analysis [23]. In our initial analysis, we will also assume that the decimal value of the discarded bits is equal to one. That is, the discarded bits (in other words, the $(W + 1)$- to R-th bits) are all equal to zero except for the R-th MSB (see Fig. 4.2). This assumption, which does not have any effect on the magnitude of the signal or the spurs, is addressed in [9]. Taking the previous assumptions into account, (2.4) becomes

$$s(n) = \sin\left(2\pi \frac{k}{2^R} n\right) \tag{4.1}$$

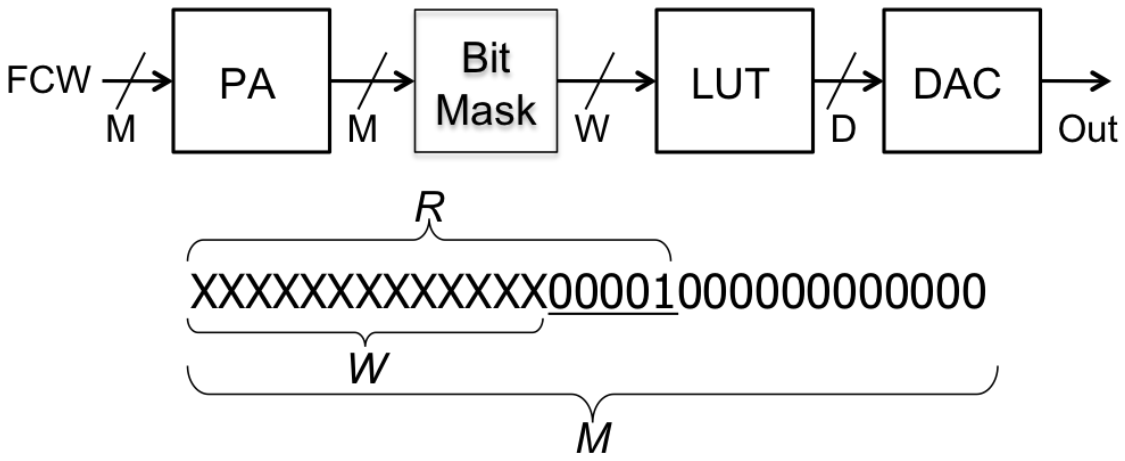

Figure 4.2: Example FCW in which M is the number of bits in the PA; R is the number of active bits; M minus the number of trailing zeroes; W is the number of bits in the LUT; and $R - W$ is the number of truncated bits.

where $k = FCW_{1:R}$, the value of the R MSBs of the original FCW. We define the R-point DFT of (4.1) as

$$S(k) = \sum_{n=0}^{2^R-1} s(n) e^{-j\frac{2\pi}{2^R}nk} \tag{4.2}$$

Since the $R - W$ LSBs of the condensed PA are truncated, the output $s(n)$ from (4.1) is constant for every 2^{R-W} samples. In order to analyze the spectrum of $s(n)$, we examine $r(n)$, a similar sequence without the repeated values, defined as

$$r(n) = s(2^{R-W}n) \tag{4.3}$$

We then define this sequence's W-point DFT as

$$R(k) = \sum_{n=0}^{2^W-1} r(n) e^{-j\frac{2\pi}{2^W}nk} \tag{4.4}$$

We can expand (4.2) into

$$S(k) = \sum_{n=0}^{2^W-1} s(2^{R-W}n)e^{-j\frac{2\pi}{2^R}(2^{R-W}n)k} + \sum_{n=0}^{2^W-1} s(2^{R-W}n+1)e^{-j\frac{2\pi}{2^R}(2^{R-W}n+1)k} + \\ ... + \sum_{n=0}^{2^W-1} s(2^{R-W}n+2^{R-W}-1)e^{-j\frac{2\pi}{2^R}(2^{R-W}n+2^{R-W}-1)k} \quad (4.5)$$

Returning to (4.3), we also note that

$$r(n) = s(2^{R-W}n) = s(2^{R-W}n+1) = ... = s(2^{R-W}n+2^{R-W}-1) \quad (4.6)$$

After placing all of the terms that do not depend on n outside of the summation, (4.5) becomes

$$S(k) = \underbrace{\left(1 + e^{-j\frac{2\pi}{2^R}k} + e^{-j\frac{2\pi}{2^R}(2)k} + ... + e^{-j\frac{2\pi}{2^R}(2^{R-W}-1)k}\right)}_{V(k)} \underbrace{\sum_{n=0}^{2^W-1} r(n)e^{-j\frac{2\pi}{2^W}nk}}_{R(k)} \quad (4.7)$$

where $V(k)$ is a finite geometric series that can be rewritten

$$V(k) = \frac{1 - e^{-j\frac{2\pi}{2^W}k}}{1 - e^{-j\frac{2\pi}{2^R}k}} \quad (4.8)$$

Hence, from (4.7) and (4.8), it is seen that the spectrum $S(k)$ is composed of 2^{R-W} replicas of $R(k)$ windowed by $V(k)$ [22, 23]. One of these is the desired output signal, while the other $2^{R-W}-1$ replicas appear as spurs in the positive frequencies (bins ranging from 0 to 2^{R-1}). From this analysis, we can establish several baseline facts about phase truncation spurs. First, the worst-case (i.e. largest) spur magnitude occurs when $R-W=1$, which means the R-th MSB in k is one and there is only a single phase truncation spur in the signal spectrum. Second, the magnitude of the worst-case phase truncation spurs decreases by 6 dB for each bit of increase in the size of W due to its dependence on $V(k)$.

4.1.3 Testing Phase Truncation Spur Correlation in an Array

To assess the level of phase truncation spur correlation in DDS arrays, we made use of 14 channels of our testbed. By evaluating the relative phases of truncation spurs on all channels, the level of spur correlation across channels was assessed. The LUT size, W, was varied by using a built-in feature of the JHU/APL DDS testbed, and output frequencies were carefully selected to result in exact truncated word sizes. To quantify the correlation, we define a meaure called "relative phase." The relative phase of a spur is equal to the phase of the phase truncation spur minus the phase of the fundamental signal. In a calibrated array of DDSs, the phases of the fundamentals are aligned, and so the statistics of this relative phase give us direct insight into the level of phase truncation spur correlation, as well as the potential for SFDR gain, in the combined output of an N-channel DDS array. In order to measure this relative phase, we took ten captures of the outputs of 14 DDSs for several different values of k. For each capture, we reinitialized the value of the PA to zero. Table 4.1 shows all of the W and k values measured, as well as their corresponding frequencies. We chose the k that produced the worst-case spur for each value of W, which was calculated using the algorithm from [22].

Table 4.1: W and k values (truncated portions after the decimal) measured along with their corresponding frequencies in MHz.

W	k	Frequency (MHz)
4	0011.1	87.5
6	001100.1	78.125
8	00110000.1	75.78125
10	0011000000.1	75.1953125

Figure 4.3 shows a histogram plotting the measured relative phase of the worst case spur for all 14 channels and 10 runs in 10°-wide bins. As is shown, the relative phases clump heavily together for each channel. We believe the 180° separation

between measurements on a given channel seen occasionally in the $W = 4$ case is an artifact arising from the dual PA architecture. For the $W = 10$ case, the noise power starts to degrade the relative phase measurement, causing some spreading in the phase. Most importantly, however, though the relative phase is constant across runs for a given channel, the relative phases of each channel vary widely, which suggests that there is some other variable determining the relative phase which varies from channel-to-channel.

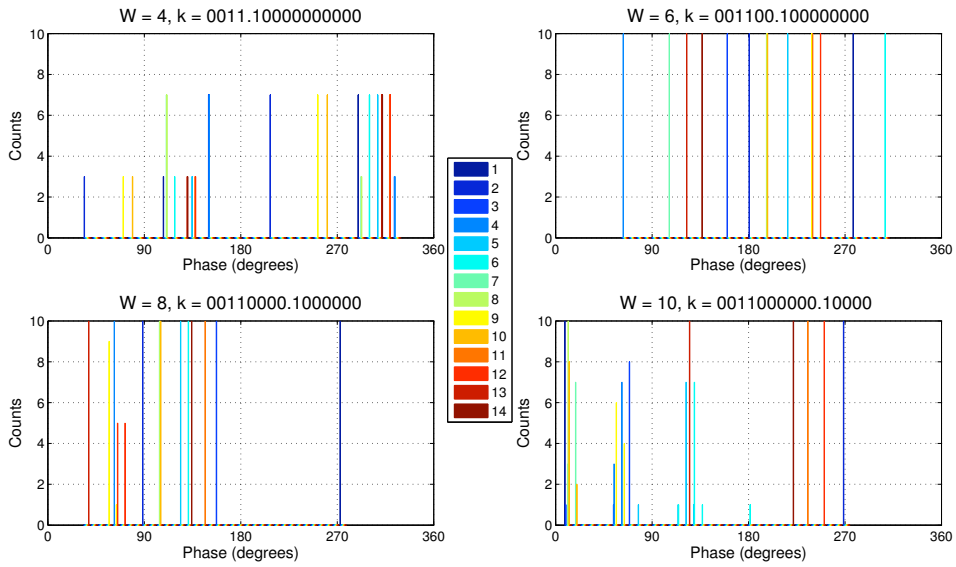

Figure 4.3: Plot of the relative phase measurements for 10 captures of a 14-channel array of DDSs for four different output frequencies as determined by the FCW, k.

4.1.4 Investigating the Starting Relative Phase

In the above measurements, subsequent runs were taken without powering down and powering up the DDSs in between runs. It is possible that some timing element in the FPGA during power up or initialization is responsible for setting the relative spur phase. This timing element could vary across channels, accounting for the variation in relative spur phase we observed. In order to examine this possibility, we

took an additional set of measurements in which we powered down and powered up the DDSs between each run. For this experiment, we took 32 captures of 14 DDS channels and chose $W = 4$ with the k value = 0011.1, corresponding to an output frequency of 87.5 MHz, in order to have the maximum spur-to-noise ratio possible. In between each capture, we turned off the power to each DDS, turned it back on, and then reinitialized it. Figure 4.4 shows a histogram plotting the measured relative phase in 10°-wide bins. As is shown, after power reset, the relative phase of the phase truncation spur is random even within a given channel with a near equal distribution for all possible phases.

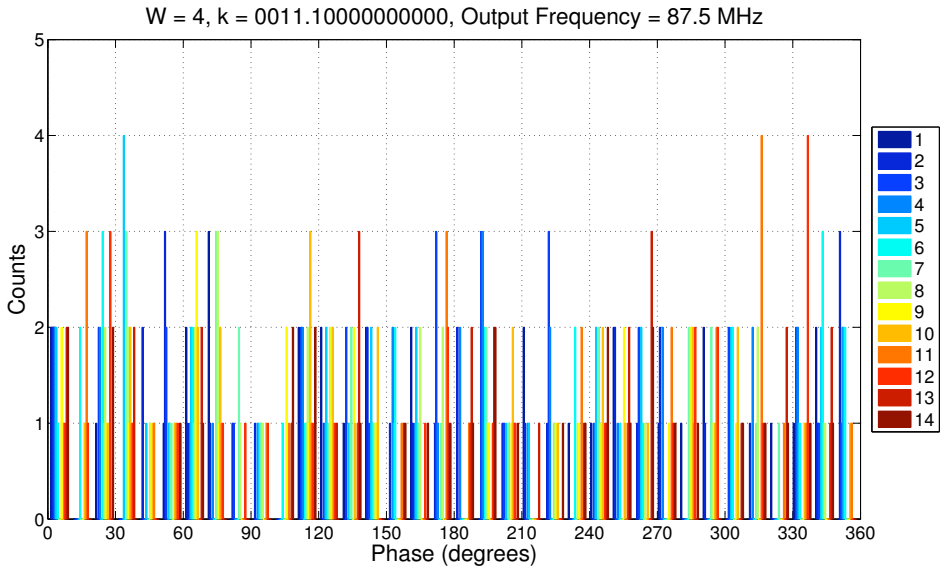

Figure 4.4: Plot of the relative phase measurements for 32 captures of a 14-channel array of DDSs for output frequency equal to 87.5 MHz.

4.1.5 Discussion

Analysis of the results from these experiments sheds some light on the correlation of phase truncation spurs in real DDS arrays. The initial experiment showed that, as expected, relative phase was constant for a given channel and did not change with LUT size. However, relative phase did vary across channels. This

suggests that spurs are not correlated across channels and that an increase in SFDR for phase truncation spurs is possible in this array. This differs from theory, which predicts total correlation and no SFDR gain. A constant relative phase for a given channel suggests there is some deterministic component which varies across channels but not across runs.

The second experiment eliminated this constant channel relative phase by resetting the power to the DDSs before each run. These results suggest that the primary determinant of phase truncation spur relative phase is some parameter that is established upon FPGA initialization. Our current model does not account for it, but Fig. 4.4 suggests that it is a uniform random variable. It is possible that other DDS arrays that implement their PAs, LUTs, and buffers in a different way, perhaps with an application-specific integrated circuit (ASIC), might not have the same randomness, but the result of the second experiment establishes the potential for achieving a SFDR gain for phase truncation spurs. Future work might be able to identify and perhaps control the mechanism in the FPGA that determines the relative phase. If identified, this mechanism could be used to the designers' advantage to purposely decorrelate phase truncation spurs and maximize SFDR gain.

4.2 Quantization Noise Spurs

4.2.1 Spur Origin

The analog output of a DAC is a quantized representation of the desired output signal; therefore, it does not have an ideal spectrum. Rather, its spectrum is composed of a fundamental tone at the desired output frequency, as well as harmonics due to signal distortion [24]. These harmonics alias with respect to the DDS's source clock, so they can be close to the fundamental frequency of the signal.

These aliased harmonics appear as a very large number of spurs, the collection of which is referred to as "quantization noise." Quantization noise is a form of spectral impurity that is inherent to a DAC, a fundamental component of the DDSs. In this section, we study the correlation of quantization noise among different channels in an array of DDSs by doing a comparison of the actual summed output power to the theoretical correlated and uncorrelated powers as described in Section 2.3. Determining the level of correlation among channels allows us to analyze the potential increase in signal-to-quantization noise ratio (SQNR) achieved by combining multiple channels in an array. Theoretically, the quantization noise for each DAC should be perfectly correlated, since the outputs of the aligned channels should have the same amplitude and suffer the same quantization error; however, the remaining calibration error and spectral differences across channels, such as phase truncation errors, might at least partially decorrelate the quantization noise. This partial decorrelation would allow for some SQNR gain from combining multiple DDS channels in an array.

4.2.2 Experimental Setup

We measured the quantization noise power in a 14-channel DDS array using the setup shown in Fig. 4.5. We varied the resolution of the DDSs by changing the bit mask before the buffer as described in Section 2.1.4. The DDSs were operated with no amplitude attenuation so as to use the entire dynamic range of the DAC. The DDSs' full-scale power is well-below the full-scale power of the ADCs so any potential ADC nonlinearities are minimized. No additional amplifiers or filters were introduced in order to reduce the impact of auxiliary electronics. In previous measurements, the DDSs were phase-aligned to about 170 μrad so as to fully correlate the fundamental outputs of the DDSs [10], and this same phase alignment was performed. The output frequencies of the DDSs were chosen to be in the second

Nyquist zone of the ADCs (60 MHz to 120 MHz). The exact experimental frequency (and the corresponding FCW) was an available variable which was chosen to fall in the center of an output FFT bin for computational ease.

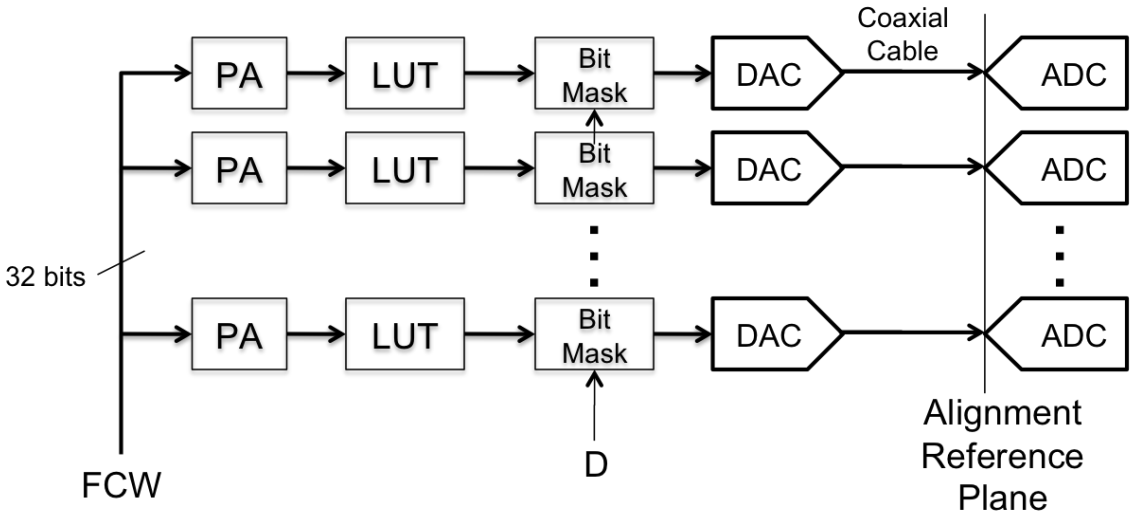

Figure 4.5: Quantization noise experimental setup includes 14 DDS channels phase-aligned at the input to the ADCs. In these measurements, the DDS DAC word length, D, was varied.

4.2.3 Calculation of Quantization Noise Power

We calculated the quantization noise power by first performing an FFT on the digital output of each channel's ADC and then squaring the FFT to yield the power spectrum. We used a flat top window on the digital output and compensated for the window loss in order to improve our amplitude accuracy. We next zeroed out the largest bin and its 15 surrounding bins on each side to remove the fundamental signal power. We then summed the linear (not decibel) values of the remaining bins to find the total noise power in each channel, and assumed the quantization noise was well above the other sources of noise. We then used the formulas in Section 2.3 to calculate the correlated and uncorrelated powers, $P_{N\,\text{correlated}}$ and $P_{N\,\text{uncorrelated}}$, respectively. To calculate the actual measured power, as in (2.16), we summed the

digital voltage outputs from all of the ADCs and then calculated the quantization noise power as we did for the other channels. Each output stream's variance was normalized so as to make the power of the fundamental signal equal for every stream of data. We performed this calculation for DAC word lengths $D = 2$ to $D = 13$ to generate a plot like the example in Fig. 4.6. The blue line shows the

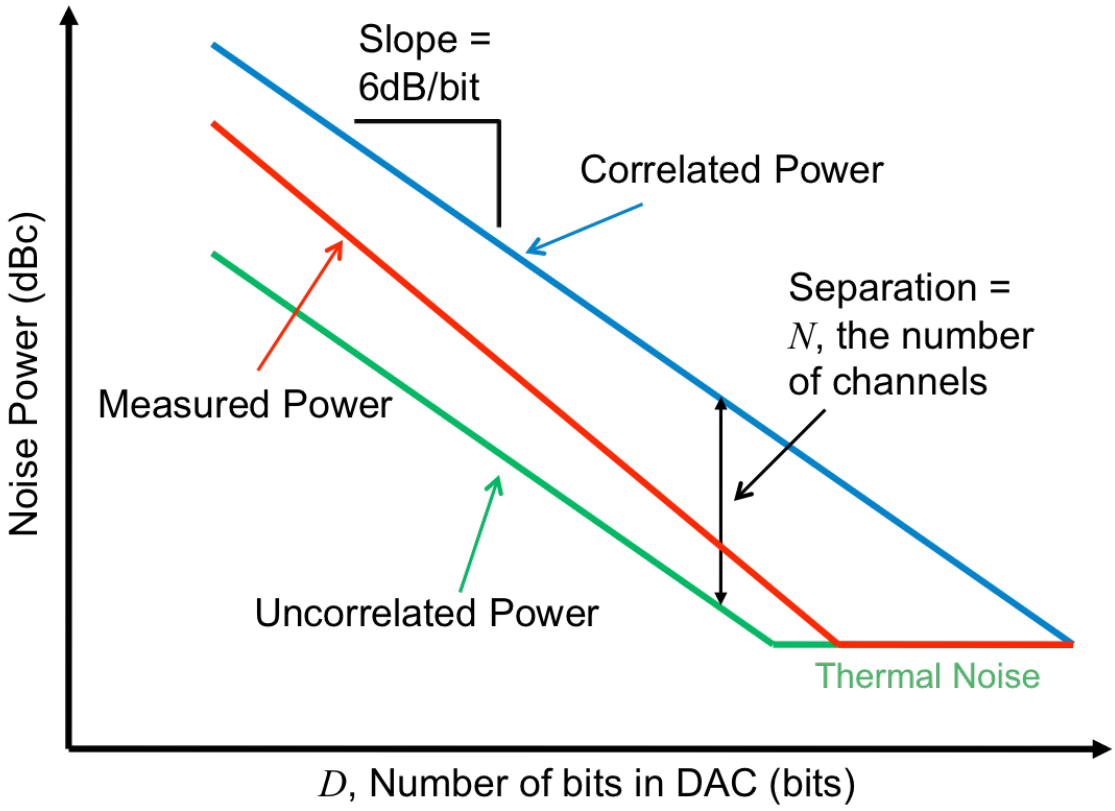

Figure 4.6: Example DAC power correlation plot. The blue line shows the theoretical correlated power, the green line shows the theoretical uncorrelated power, and the red line shows the actual measured power.

theoretical correlated power, the green line shows the theoretical uncorrelated power, and the red line shows the actual measured power. Both the theoretical correlated and uncorrelated powers decrease as D increases until the thermal noise floor is reached. The relative position of the red line with respect to the blue and green lines gives insight into the actual correlation of the quantization noise power

for any given D; the closer the measured values are to the blue line, the more correlated the noise; the closer they are to the green line, the more uncorrelated. Once the thermal noise limit is reached, any analysis of the correlation of the quantization noise power is inconclusive.

4.2.4 Results

The results of this measurement at four different output frequencies are shown in Fig. 4.7. For very high DAC resolutions, our results were inconclusive because uncorrelated thermal noise dominated our measurements; however, our analysis shows that the measured quantization noise power is fairly decorrelated among DDS channels for all of the lower DAC resolutions. This trend was independent of the output frequency of the DDS. The measured and uncorrelated powers never quite overlap after reaching the thermal noise limit, which we attribute to a small level of correlation in the ADC noise power between channels and not to the DDSs. These results suggest that it may be possible to construct a large array of low-bit DACs which generate a combined signal with very high SQNR.

4.2.5 Discussion

Despite theoretical predictions of fully correlated DDS quantization noise, the measured quantization noise power of the combined output signal suggests that the quantization noise of individual DDSs is mostly decorrelated for all values of D. The cause of the discrepancy between measurement and theory is unclear, but we note that the relative phases of the quantization noise spurs determine the extent of quantization noise power correlation among the elements of the array. It is possible that small discrepancies among the channels could serve to decorrelate the quantization noise as was observed for the phase truncation spurs. Nevertheless, the

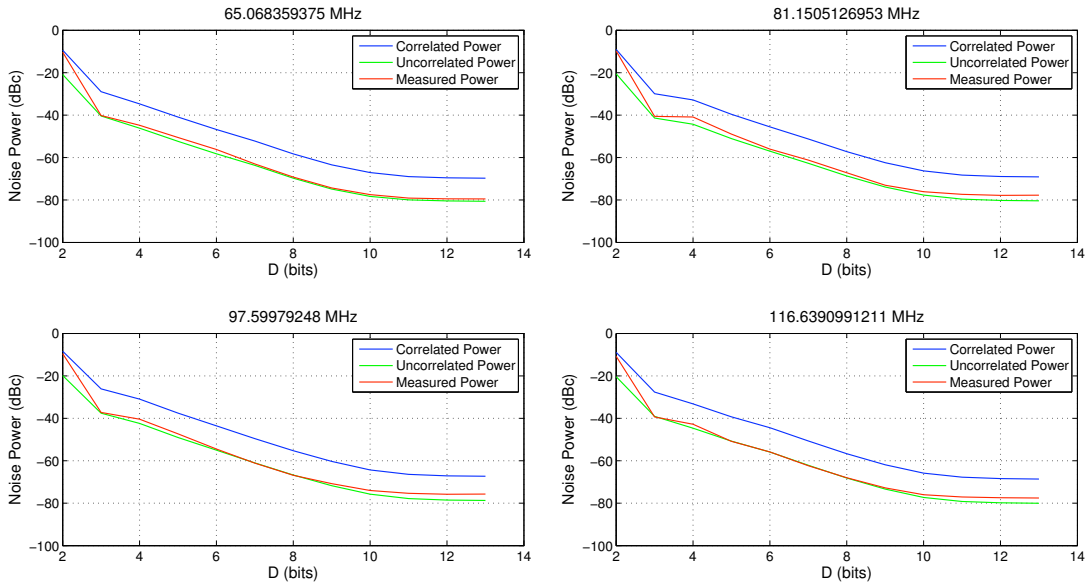

Figure 4.7: Quantization noise power results for four different frequencies.

results suggest that a large array of low-bit DACs that are amplitude and phase-aligned could yield high SQNR without a huge penalty from correlated quantization noise.

4.3 Quantizer Nonlinearity Spurs

4.3.1 Spur Origin

Every DAC has a transfer function that converts a digital code to an analog signal. The ideal DAC transfer function maps the digital codes to evenly spaced slices of the full-scale range of the DAC. Quantizer nonlinearities result from the imperfections in the design and fabrication of the DAC such that there is an error between the ideal output levels of a DAC and its actual output levels. This error is quantified with two metrics, differential nonlinearity (DNL) error and integral nonlinearity (INL) error. A DNL error is defined as the difference between the actual spacing between adjacent digital values (values differing by 1 LSB) of the

DAC and the ideal spacing. An INL is the deviation of an actual DAC transfer function from a straight line. This straight line can either be the best linear fit of the transfer function or it can be a straight line passing through the endpoints of the DAC's transfer function [25]. According to the data sheet for the Analog Devices AD9736 DAC used in our DDS array, the worst case DNL and INL errors are ± 2.1 LSBs and ± 5.6 LSBs, respectively [26]. The result of these errors is that the transfer function of the DAC is nonlinear, causing the output to contain both the desired fundamental output as well as harmonic distortion. Harmonic distortion manifests itself as harmonically related spurs in the output of the DAC whose amplitudes are not readily predictable. However, their location is predictable since they appear at multiples of the fundamental frequency. The harmonics whose frequencies are higher than half of the DDS clock frequency alias back onto the first Nyquist zone [27].

4.3.2 Experimental Setup

For the normal operating setting of the DDSs, $D = 14$ and $W = 17$, the phase truncation spurs and the quantization noise spurs are negligible compared to the quantizer nonlinearity spurs. The DACs' quantizer nonlinearities could not be measured directly using the ADCs as signal detectors because the ADC nonlinearities can overlap with the DAC harmonics in the frequency ranges of interest. Instead, we used an Agilent E4440A spectrum analyzer with a flat top window to measure the power of the fundamental and its harmonics for each channel in order to calculate the theoretical correlated and uncorrelated power. To calculate the actual correlated power, we then combined eight channels together with an eight-way Wilkinson power combiner, the largest we had available. We phase-aligned the channels by shutting all but the first channel and a second channel and then incrementing or decrementing the second channel's phase accumulator to achieve maximal destructive interference between the two channels. After finding this point,

we added 180 degrees to the second channel's phase accumulator to place the two channels in-phase. We then verified the expected gain, $20\log_{10}(2) = 6$ dB, increase in power over the single channel's power. After doing this for all eight channels, we combined all eight DDSs and verified the expected $20\log_{10}(8) = 18$ dB gain in power over the single channel's power. With this alignment complete, we measured the power of the fundamental and harmonics, adjusting for the insertion loss of the power combiner, which we measured for each frequency with a network analyzer.

4.3.3 Results

Results from these measurements for many harmonics of an output frequency of 20.5 MHz are plotted in Fig. 4.8. The plot shows the theoretical correlated and uncorrelated powers along with the actual measured power from the combiner for each harmonic measurement. As is shown in the plot, there is a strong level of correlation for the second and third harmonics. In the cases of the higher level harmonics, we see more erratic behavior. This is likely due to the higher order harmonics being more sensitive to noise because their magnitudes are lower and, therefore, closer to the thermal noise floor of the signal. Figure 4.9 shows the correlation of the second and third harmonics for multiple frequencies.

4.3.4 Discussion

As the measurements suggest, the lowest-order quantizer nonlinearity harmonics are highly correlated across channels for all tested frequencies. As a result, there is little benefit to using an array of DDSs to decrease the effect of in-band quantizer nonlinearities. Instead, system designers must address the problem of these spurs using alternate methods, such as careful frequency planning, to ensure that high-power harmonics (particularly the second and third) do not alias into the

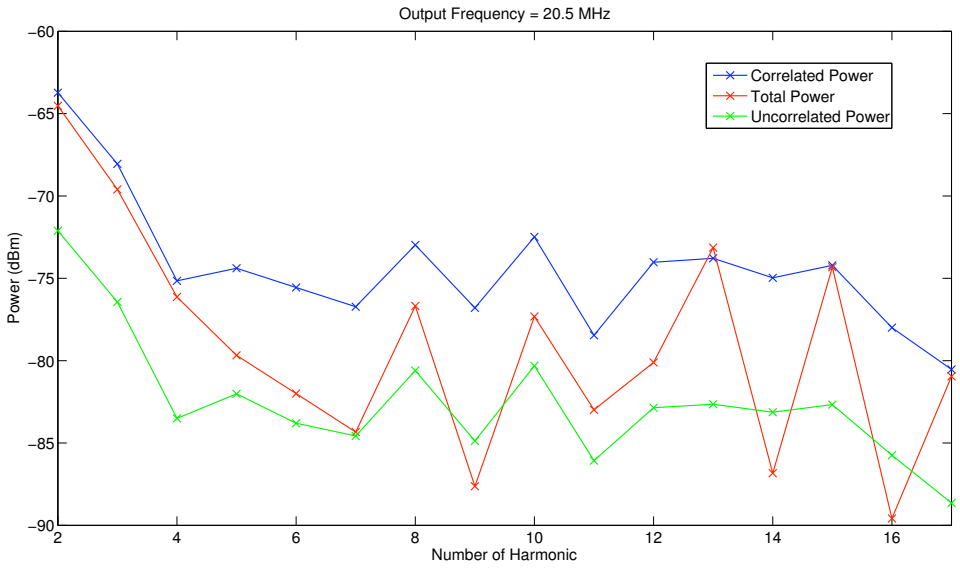

Figure 4.8: Plots of the total power of multiple harmonics for a DDS output frequency of 20.5 MHz.

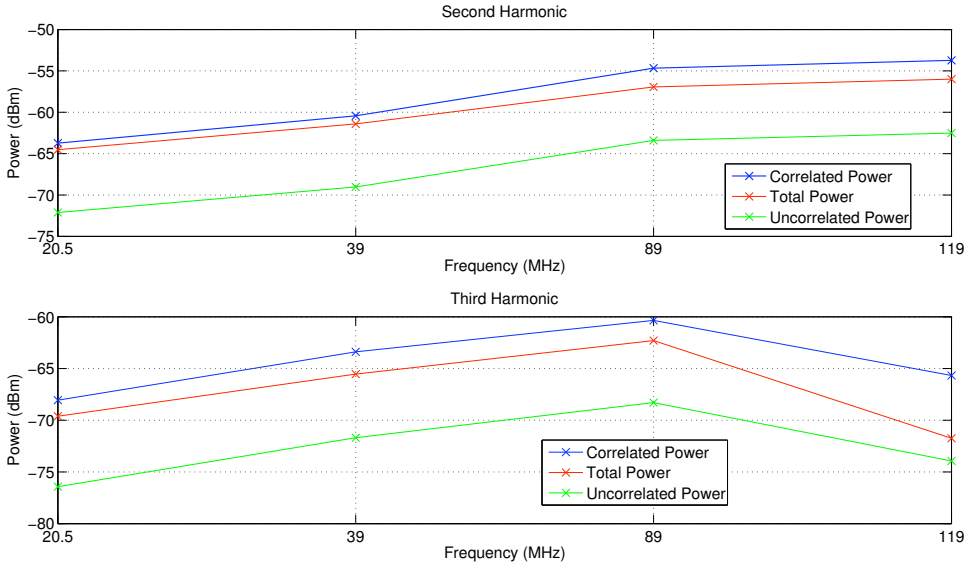

Figure 4.9: Plots of the total power for the second and third harmonics for multiple frequencies.

frequency band of interest and can be adequately filtered out [27]. Channels could also be phased in such a way to cancel specific higher order harmonics.

CHAPTER 5

DISCUSSION

5.1 Conclusion

The work reported in this thesis demonstrates that there are many potential performance benefits to using an array of DDSs. Our investigation of the phase noise of a DDS array indicates that the contribution to the phase noise from the DACs can be decreased to a desired level by using a large enough number of channels. In such a system, the phase noise qualities of the source clock and the system cost and complexity will be the main limitations on the phase noise of the DDS array.

Our study of phase truncation spurs suggests that, at least in our system, the phase truncation spurs are uncorrelated, contrary to the theoretical prediction. We believe this decorrelation is due to the existence of an unidentified mechanism in our DDS array that is unaccounted for in our current operational DDS model. This mechanism, likely due to some timing element in the FPGA, causes some randomness in the relative phases of the truncation spurs from channel to channel each time the DDS array is powered up. This randomness decorrelates the phase truncation spurs, opening the potential for SFDR gain from using a DDS array.

Our analysis of the correlation of quantization noise spurs in an array of DDSs shows that the total quantization noise power of each DDS channel is uncorrelated for nearly all values of DAC output bits. This suggests that a near N gain in SQNR is possible for an N-channel array of DDSs. This gain will be most apparent for low-bit DACs in which quantization noise is notably higher than the thermal noise

contribution.

Lastly, our measurements of the correlation of quantizer nonlinearity spurs demonstrate that the second and third harmonics are highly correlated across channels for all frequencies tested. As a result, alternate methods of harmonic spur management must be employed.

5.2 Future Work

Extensions to and modifications of the hardware used in this thesis work would open up several avenues for future research in the field of noise and spur correlation in DDS arrays. Increasing the maximum offset frequency of the phase noise measurement system may allow direct measurement of the DDS floor noise and verification of its contribution to the DDS array's phase noise. Developing the capability to directly observe the instantaneous value of the PA for each DDS channel would facilitate a more thorough analysis, possibly enabling the identification of the mechanism responsible for spur decorrelation. Developing a DDS testbed with more than 14 channels may provide more insight into the level of decorrelation of the total quantization noise from the DACs. Lastly, developing a way to sample the entire time series for the quantizer nonlinearity spur measurements of a DDS without being obscured by the ADC nonlinearities may allow for a more thorough analysis of the correlation of harmonics among the channels of a DDS array.

REFERENCES

[1] V. F. Kroupa, Ed., *Direct Digital Frequency Synthesizers*. New York, NY: Institute of Electrical and Electronics Engineers, 1999.

[2] L. Cordesses, "Direct digital synthesis: A tool for periodic wave generation (part 1)," *IEEE Signal Processing Magazine*, vol. 21, no. 4, pp. 50–54, July 2004.

[3] J. Tierney, C. Rader, and B. Gold, "A digital frequency synthesizer," *IEEE Trans. Audio Electroacoust.*, vol. AU-19, no. 1, pp. 48–57, March 1971.

[4] K. Lauritzen, "DDS Block Diagram," private communication, March 2010.

[5] J. Vankka, *Digital Synthesizers and Transmitters for Software Radio*. Dordrecht, The Netherlands: Springer, 2005.

[6] S. Mehrgardt, "Noise spectra of digital sine-generators using the table-lookup method," *IEEE Trans. Acoust. Speech Signal Process.*, vol. ASSP-31, no. 4, pp. 1037–1039, August 1983.

[7] H. Nicholas III, H. Samueli, and B. Kim, "Optimization of direct digital frequency synthesizer performance in the presence of finite word length effects," in *Proceedings of the Annual IEEE International Frequency Control Symposium*, June 1988, pp. 357–363.

[8] J. Vankka, "Methods of mapping from phase to sine amplitude in direct digital synthesis," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 44, no. 2, pp. 526–534, March 1997.

[9] H. Nicholas III and H. Samueli, "An analysis of the output spectrum of direct digital frequency synthesizers in the presence of phase-accumulator truncation," in *Proceedings of the 41st Annual Frequency Symposium 1987 (Cat. No.87CH2427-3)*, May 1987, pp. 495 – 502.

[10] K. Lauritzen, E. Richards, S. Talisa, and M. Peckerar, "Design of 16-channel analog-to-digital converter test bed to measure correlation," July 2010, unpublished.

[11] Ferre-Pikal, E.S. *et al.*, "Draft revision of IEEE STD 1139-1988 standard definitions of physical quantities for fundamental, frequency and time metrology-random instabilities," in *Proceedings of the 1997 IEEE International Frequency Control Symposium.*, May 1997.

[12] S. Wu and Y. Bar-Ness, "OFDM systems in the presence of phase noise: Consequences and solutions," *IEEE Transactions on Communications*, vol. 52, no. 11, pp. 1988–1996, November 2004.

[13] F. Walls, A. Clements, C. Felton, M. Lombardi, and M. Vanek, "Extending the range and accuracy of phase noise measurements," in *Proceedings of the Annual IEEE International Frequency Control Symposium*, vol. 42, June 1988, pp. 432–441.

[14] W. F. Walls, "Practical problems involving phase noise measurements," in *33rd Annual Precise Time and Time Interval (PTTI) Meeting*, November 2001, pp. 407–416.

[15] F. Walls, "PM and AM noise of combined signal sources," in *Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum, 2003.*, May 2003, pp. 532–540.

[16] Y.-C. Jenq, "Direct digital synthesizer with jittered clock," *IEEE Trans. Instrum. Meas.*, vol. 46, no. 3, pp. 653–655, June 1997.

[17] V. Kroupa, "Close to the carrier noise in DDS," in *Proceedings of the 1996 IEEE International Frequency Control Symposium (50th Anniversary) (Cat. No.96CH35935)*, June 1996, pp. 934–941.

[18] V. Kuleshov, H. Liu, and B. Kuleshov, "1/f fluctuations sources in direct digital frequency synthesizers and their contribution to the output oscillations power spectral density," in *Proceedings of the 1995 IEEE International Frequency Control Symposium (The 49th Annual Symposium)*, May 1995, pp. 282–287.

[19] V. Kuleshov and H. Liu, "Fundamental noise in direct digital frequency synthesizers," in *Proceedings of the 1995 IEEE International Frequency Control Symposium. (The 49th Annual Symposium) (Cat.No.95CH35752)*, May 1995, pp. 288–293.

[20] W. Robins, *Phase Noise in Signal Sources: Theory and Applications*. London, UK: Peter Peregrinus Ltd., 1982.

[21] A. Lance, W. Seal, and F. Labaar, *Phase Noise and AM Noise Measurements in the Frequency Domain*. Orlando, FL, USA: Academic Press, 1984, pp. 239–289.

[22] A. Torosyan and A. Willson Jr., "Exact analysis of DDS spurs and SNR due to phase truncation and arbitrary phase-to-amplitude errors," in *Proceedings of the IEEE International Frequency Control Symposium and Exposition*, 2005, pp. 50–58.

[23] S. Cheng, J. Jensen, R. Wallis, and G. Weaver, "Further enhancements to the analysis of spectral purity in the application of practical direct digital synthesis," in *Proceedings of the 2004 IEEE International Frequency Control Symposium And Exhibition (IEEE Cat. No.04CH37553C)*, 2004, pp. 462–470.

[24] R. Gray, "Quantization noise spectra," *IEEE Trans. Inf. Theory*, vol. 36, no. 6, pp. 1220–1244, November 1990.

[25] Maxim Integrated Products, "Application Note 283," September 2000. [Online]. Available: http://pdfserv.maxim-ic.com/en/an/AN283.pdf

[26] Analog Devices, "AD9734/AD9735/AD9736," Datasheet, 2006.

[27] Analog Devices, "A technical tutorial on digital signal synthesis," Application Note, 1999.