

Nereo Markulic
Kuba Raczkowski
Jan Craninckx
Piet Wambacq

Digital Subsampling Phase Lock Techniques for Frequency Synthesis and Polar Transmission

Analog Circuits and Signal Processing

Series Editors:

Mohammed Ismail, Dublin, USA

Mohamad Sawan, Montreal, Canada

The Analog Circuits and Signal Processing book series, formerly known as the Kluwer International Series in Engineering and Computer Science, is a high level academic and professional series publishing research on the design and applications of analog integrated circuits and signal processing circuits and systems. Typically per year we publish between 5–15 research monographs, professional books, handbooks, edited volumes and textbooks with worldwide distribution to engineers, researchers, educators, and libraries.

The book series promotes and expedites the dissemination of new research results and tutorial views in the analog field. There is an exciting and large volume of research activity in the field worldwide. Researchers are striving to bridge the gap between classical analog work and recent advances in very large scale integration (VLSI) technologies with improved analog capabilities. Analog VLSI has been recognized as a major technology for future information processing. Analog work is showing signs of dramatic changes with emphasis on interdisciplinary research efforts combining device/circuit/technology issues. Consequently, new design concepts, strategies and design tools are being unveiled.

Topics of interest include:

Analog Interface Circuits and Systems;
Data converters;
Active-RC, switched-capacitor and continuous-time integrated filters;
Mixed analog/digital VLSI;
Simulation and modeling, mixed-mode simulation;
Analog nonlinear and computational circuits and signal processing;
Analog Artificial Neural Networks/Artificial Intelligence;
Current-mode Signal Processing;
Computer-Aided Design (CAD) tools;
Analog Design in emerging technologies (Scalable CMOS, BiCMOS, GaAs, heterojunction and floating gate technologies, etc.);
Analog Design for Test;
Integrated sensors and actuators;
Analog Design Automation/Knowledge-based Systems;
Analog VLSI cell libraries;
Analog product development;
RF Front ends, Wireless communications and Microwave Circuits;
Analog behavioral modeling, Analog HDL.

More information about this series at <http://www.springer.com/series/7381>

Nereo Markulic • Kuba Raczkowski
Jan Craninckx • Piet Wambacq

Digital Subsampling Phase Lock Techniques for Frequency Synthesis and Polar Transmission

Nereo Markulic
IMEC
Leuven, Belgium

Kuba Raczkowski
IMEC
Leuven, Belgium

Jan Craninckx
IMEC
Leuven, Belgium

Piet Wambacq
IMEC
Leuven, Belgium

ISSN 1872-082X

ISSN 2197-1854 (electronic)

Analog Circuits and Signal Processing

ISBN 978-3-030-10957-8

ISBN 978-3-030-10958-5 (eBook)

<https://doi.org/10.1007/978-3-030-10958-5>

Library of Congress Control Number: 2018966811

© Springer Nature Switzerland AG 2019

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG.

The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

We are the music makers and we are the dreamers of dreams.

Arthur O'Shaughnessy, *Ode*, 1873

Willy Wonka, *Willy Wonka & the Chocolate Factory*, 1971

Aphex Twin, *Selected Ambient Works Volume I*, 1994

Preface

Wireless technology systems have intruded in almost every aspect of today's communication. Technology scaling and innovation in the field of integrated circuits (ICs) nurture this wireless revolution, while the need for higher data throughput continues to grow. These trends unfold a severe challenge: in today's over-allocated spectrum, its efficient use becomes absolutely essential.

In the heart of every transceiver lies a local oscillator (LO), typically implemented as a phase-locked loop (PLL). Crucial aspects of the LO synthesizer are its phase noise and spurious performance. These impose the fundamental limitation for efficient transmit and receive modes; hence, a considerable amount of energy and chip area are typically spent to minimize them. Moreover, in modern systems, the PLL is often used for phase modulation, within digitally intensive polar architectures. For spectrally pure and efficient operation, the digital-to-transmitted output conversion cannot be bandwidth limited or compromised by nonlinearities.

The book starts with an introductory overview of modern frequency synthesis techniques, delivering the basic operation theory in an intuitive fashion, with practical implementation in mind. A point of attention is in this context brought to recent subsampling architectures. These architectures overcome the performance boundaries typically encountered in classical implementations and have potential for redefining today's state of the art. The following chapters, built around three 28 nm bulk CMOS IC prototypes, explore this idea and present new, performance-leading PLL and polar transmitter designs.

The first presented prototype develops a new fractional PLL from a subsampling integer-N frequency multiplier, which in its original form could not be used for modern wireless standards. To enable fractional modes, while benefiting from extremely low-noise subsampling operation, we introduce the principle of *digital-to-time converter* (DTC)-based time domain signal processing. A DTC, in contrast to a time-to-digital converter (TDC), used in modern digital PLLs, easily reaches fine resolution that is crucial for spectral purity.

In the second prototype, we resolve the fundamental limitation of nonlinear phase-error detection within PLLs. We demonstrate the enhanced, background-calibrated subsampling PLL which operates with a record-breaking -247 dB figure

of merit, challenging the most advanced art in the field. The synthesizer is further expanded into a two-point digital phase modulator, used as a typical building block for wide-band polar transmission.

The third and final IC introduces the digital subsampling polar transmitter (SSPTX), a new transmitter architecture which combines the subsampling core and an amplitude modulating power amplifier within a single PLL. The system exhibits specific features that enable full background cancellation of phase/amplitude modulation-induced distortion. The 5.5 GHz polar transmitter achieves extremely accurate performance with -41 dB EVM at a 1024 QAM constellation allowing barrier breaking information throughput, essential for the upcoming wireless standards.

Leuven, Belgium

Nereo Markulic
Kuba Raczkowski
Piet Wambacq
Jan Craninckx

Contents

1	Introduction	1
1.1	A Transceiver with a Local Oscillator (LO) in Its Core	1
1.1.1	A Cartesian Transceiver	1
1.1.2	A Polar Transmitter	3
1.2	A Phase-Locked Loop (PLL) as an LO	4
1.2.1	From an Analog to a Mixed-Signal and Digital PLL	5
1.2.2	Small Signal Model of a PLL	9
1.2.3	Phase Noise in PLLs	11
1.3	Motivation and Research Objectives	14
1.3.1	A Subsampling PLL	14
1.3.2	Objectives of the Book	17
1.4	Book Outline	18
	References	19
2	A Digital-to-Time-Converter-Based Subsampling PLL for Fractional Synthesis	23
2.1	Introduction	23
2.2	Fractional-N Operation of a Subsampling PLL	25
2.2.1	Time-Domain Analysis of a Subsampling PLL	25
2.2.2	Enhancement of a Subsampling PLL to Enable Fraction-N Mode Operation	25
2.2.3	Digital Modulator for the Fractional-N Subsampling PLL	27
2.3	Implementation Limitations and Their Mitigation	28
2.3.1	DTC Quantization	28
2.3.2	A DTC Versus a TDC in Fractional Frequency Synthesis	30
2.3.3	DTC Offset and Gain Error	30
2.3.4	DTC Nonlinearity	31
2.3.5	DTC Phase Noise	32
2.4	Circuit Implementation	32
2.4.1	Implementation of the Subsampling Loop	34
2.4.2	Implementation of the Digital-to-Time Converter	36

2.4.3	Implementation of the VCO	41
2.4.4	Implementation of the Frequency-Acquisition Loop	43
2.5	Experimental Results	43
2.5.1	Measured Phase Noise Performance	46
2.5.2	Remaining Fractional Spur	49
2.5.3	DTC-Related Measurements	49
2.5.4	Performance Summary and Comparison to the State of the Art	52
2.6	Conclusion	54
	References	54
3	A Background-Calibrated Subsampling PLL for Phase/Frequency Modulation	57
3.1	Introduction	57
3.1.1	PLL-Based Phase Modulation	57
3.1.2	A DTC-Based Fractional-N Subsampling PLL for Phase Modulation	59
3.2	A Self-Calibrated DTC-Based FNSSPLL	60
3.2.1	Basic Operation of the FNSSPLL	60
3.2.2	The Random-Jump for DTC Quantization Noise Randomization	62
3.2.3	The Random-Jump for DTC Nonlinearity Randomization	64
3.2.4	Self-Calibration of the DTC Nonlinearity	64
3.2.5	Extraction of the Current's Sign and Comparator Offset Compensation	68
3.3	Two-Point Phase Modulator Based on the FNSSPLL	70
3.3.1	Modulating fDAC INL Calibration	71
3.3.2	Delay-Spread Cancellation	72
3.4	Experimental Results	74
3.5	Conclusion	81
	References	81
4	A Background-Calibrated Digital Subsampling Polar Transmitter	85
4.1	Introduction	85
4.2	System Overview	86
4.2.1	A Digital Fractional-N Subsampling PLL	86
4.2.2	Phase/Frequency and Amplitude Modulation	89
4.2.3	Prototype Targets and Building Block's Specifications	91
4.3	Digital Linearization Techniques	93
4.3.1	PM-to-PM Background Calibration	94
4.3.2	AM-to-AM Distortion Background Calibration	96
4.3.3	Phase-Domain Matlab Simulations of Background Calibration	100
4.4	Built-in AM-to-PM Distortion Filtering	101

4.5	Analog Building Blocks	104
4.5.1	Subsampling Path: From Sampler to Code	104
4.5.2	Inverse Class-D DPA with Harmonic Rejection Mixing (HRM)	109
4.6	Measured Results	113
4.6.1	Digital Subsampling PLL Measurements	113
4.6.2	Digital Subsampling Polar TX Measurements	116
4.7	Conclusion	122
	References	123
5	Conclusion and Future Outlook	127
5.1	Summary	127
5.2	Contributions	130
5.3	Future Outlook	131
5.3.1	Lock Time Optimization in a Subsampling PLL	131
5.3.2	Towards Higher Modulation Bandwidths and Better Out-of-Band Noise Rejection	132
5.3.3	Towards Higher TX Efficiency	133
5.3.4	Towards Other Modulation Schemes	133
	References	134
Index		135

Nomenclature

Q quality factor

ADC	Analog-to-Digital Converter
ADPLL	All-Digital Phase-Locked Loop
ACLR	Adjacent Channel Leakage Ratio
AM	Amplitude Modulation
CP	Charge Pump
DAC	Digital-to-Analog Converter
DCO	Digitally Controlled Oscillator
DPA	Digital Power Amplifier
DTC	Digital-to-Time Converter
EVM	Error Vector Amplitude
FMCW	Frequency-Modulated Continuous Wave
FNSSPLL	Fractional-N Subsampling Phase-Locked Loop
FOM	Figure of Merit
HRM	Harmonic Rejection Mixing
IF	Intermediate Frequency
INL	Integral Nonlinearity
IPN	Integrated Phase Noise
LNA	Low-Noise Amplifier
LO	Local Oscillator
LPF	Low-Pass Filter
LSB	Least Significant Bit
LUT	Look-Up Table
PA	Power Amplifier
PAPR	Peak-to-Average Power Ratio
PD	Phase-Error Detector
PFD	Phase/Frequency Detection
PLL	Phase-Locked Loop
PM	Phase Modulation
PVT	Process-Temperature-Voltage

RF	Radio Frequency
TDC	time-to-digital converter
TX	Transmitter
SAR	Successive Approximation Register
SSPTX	Subsampling Polar Transmitter
VCO	Voltage-Controlled Oscillator

List of Figures

Fig. 1.1	Zero-IF transceiver concept	2
Fig. 1.2	LO noise effect in a transmitter	2
Fig. 1.3	LO noise effect in a receiver	3
Fig. 1.4	Principle of polar transmission	4
Fig. 1.5	Classical analog PLL with integer- N frequency multiplication	5
Fig. 1.6	Classical mixed-signal PLL with fractional- N frequency multiplication in phase lock	6
Fig. 1.7	Simplified digital PLL in phase lock. A programmable $\Delta\Sigma$ driven divider is used in the feedback as in the mixed-signal PLL from Fig. 1.6	8
Fig. 1.8	Small phase-domain signal model of a classical PLL and illustration of the important transfer functions. Note that the axes are logarithmic	9
Fig. 1.9	Phase noise generation in the classical mixed-signal PLL. Equivalently, the phase noise generation in a typical digital PLL is depicted in 1.11 and analyzed later	11
Fig. 1.10	Output and input referred noise profiles (solid line) with corresponding PLL filtering (dashed line)	12
Fig. 1.11	Phase noise generation in a typical digital PLL	13
Fig. 1.12	Comparison of recent integer- N (black) and fractional- N (gray) PLLs in power-vs-RMS jitter trade-off FOM. Note that better results appear in bottom left of the figure. The black circles indicate the subsampling PLL architecture (see Sect. 1.3.1)	14
Fig. 1.13	Operation of a subsampling PLL	15
Fig. 1.14	Phase noise generation in a subsampling PLL	16
Fig. 2.1	General system of a subsampling PLL with example timing	24

Fig. 2.2	Implementation of fractional- N subsampling operation by delaying the sampling reference. The last sampling event aligns exactly with the beginning of the next cycle and therefore has no extra delay	26
Fig. 2.3	Digital computation flow of the DTC modulator for $N = 1.75$	27
Fig. 2.4	Digital DTC modulator including quantization, gain correction, and quantization noise shaping	28
Fig. 2.5	Background calibration method for correcting DTC gain error	31
Fig. 2.6	Simulation results of gain correction mechanism when a 10% error is applied to the DTC	31
Fig. 2.7	Architecture of the fractional- N subsampling PLL	33
Fig. 2.8	The subsampling PLL always locks into a state that guarantees zero output current, even in the presence of offset and mismatch	33
Fig. 2.9	Simplified schematic of the subsampling loop	34
Fig. 2.10	Schematic of the transconductor. The input pair is driven by differential sampled voltage. The output current (i_{out}) is duty-cycled and flows to the loop filter	35
Fig. 2.11	DTC architecture overview	37
Fig. 2.12	(left) Basic delay block; (right) delay block in the proposed DTC	37
Fig. 2.13	(a) Input slope-dependent comparator response; (b) typical nonlinear DTC delay transfer curve induced by slope-dependent comparator switching	40
Fig. 2.14	Internal DTC regulated supply for comparator and buffer	40
Fig. 2.15	Class-B VCO schematic and layout floorplan of the NMOS-only digital varactor unit cell of Fig. 2.16b	41
Fig. 2.16	Proposed and conventional switched capacitor structures. The proposed cell (a) is used to implement the digital varactor of the VCO. (a) Proposed switched capacitor cell. (b) Conventional cell of [Sjoland02]	42
Fig. 2.17	Architecture of the frequency-acquisition loop	43
Fig. 2.18	Chip microphotograph	44
Fig. 2.19	Measured VCO tuning range. Analog tuning (0–1.8 V) is used between digital words	44
Fig. 2.20	Measured VCO free-running phase noise for low-power ($V_{\text{DD}} = 0.9$ V) and high-power ($V_{\text{DD}} = 1.4$ V) mode	45
Fig. 2.21	Measured INL and DNL characteristics of the DTC	45
Fig. 2.22	Measured phase noise for a worst-case fractional- N scenario. For reference, the integer- N phase noise trace is shown as well	46
Fig. 2.23	Measured RMS jitter across fractional codes (integer part of $N = 250$) and integer- N jitter with respect to VCO tuning range	47

Fig. 2.24	Measured output spectrum of the PLL showing the worst-case fractional spur and the reference spur	48
Fig. 2.25	Measured effect of DTC gain mismatch. 1% error in gain was intentionally applied	48
Fig. 2.26	Measured phase noise as a function of the $\Delta\Sigma$ modulator order. The higher the order, the lower the spurs	49
Fig. 2.27	Micrograph of the DTC	50
Fig. 2.28	PLL output at 10 GHz for different DTC settings	51
Fig. 2.29	Measured PLL performance with DTC regulation ON (black dot) and OFF (open circle). (a) Worst fractional spur. (b) RMS jitter. (c) In-band phase noise	51
Fig. 2.30	Figure-of-merit comparison of recent fractional- N synthesizers	52
Fig. 3.1	Basic principle of single-point phase/frequency modulation with data predistortion. PLL_{TF} stands for the PLL phase signal transfer function from the input to the output. Note that modulation data propagates through the whole loop	58
Fig. 3.2	Basic principle of two-point phase/frequency modulation. Note that the modulation data propagates from the point-two to the output, but does not propagate through the loop (cancellation in point-one)	59
Fig. 3.3	Basic principle of two-point phase/frequency modulation in a DTC-based PLL. The modulation data propagates from the point-two to the output, but does disturb the loop (cancellation in point-one)	60
Fig. 3.4	A DTC-based fractional- N subsampling PLL. (a) Simplified schematic. (b) Time-domain operation	61
Fig. 3.5	(a) DTC input code calculation path in the fractional- N subsampling PLL with pseudo-random bit sequence (PRBS) generated random integer number (Random-Jump); (b) DTC input with and without Random-Jump calculation path in the time domain. (c) Generation of fractional spurs because of DTC INL	62
Fig. 3.6	A 10-bit 0.5 ps LSB DTC's quantization noise spectrum around a 10 GHz fractional carrier ($40\text{ MHz} \times 253 + 2^{-7}$) before the PLL filtering (in-band) without and with randomization (a) linear DTC (b) nonlinear DTC (± 1 LSB INL error). A randomized nonlinear DTC increases in-band noise floor masking the shaped ($\Delta\Sigma$) quantization noise. Note that a smooth, first-order nonlinearity is used in this example. In the presence of higher-order nonlinearities, spurs are only suppressed and not completely masked by noise	63

Fig. 3.7	(a) Concept of FNSSPLL DTC predistortion based on a lookup table (LUT); (b) predistortion principle: LUT stores a curve which mimics inverse of the DTC nonlinearity.....	65
Fig. 3.8	The sign of the G_m output current is correlated to the error of the DTC code	66
Fig. 3.9	(a) Digital self-calibration of the DTC with a 1024-entry LUT; (b) digital self-calibration of the DTC with a 32-entry LUT and piece-wise linear approximation; (c) time-domain simulation of the FNSSPLL with 32-entry LUT-based DTC calibration with a nonlinear DTC, and (d) final INL estimates.....	66
Fig. 3.10	G_m transconductor and current extraction mechanism schematic and operation time diagrams	68
Fig. 3.11	(a) Digital background offset calibration implementation. (b) Background offset calibration simulation with 0.15σ comparator input swing offset. When the loop settles, the overflow flag is activated approximately every 10th cycle.....	69
Fig. 3.12	Simplified two-point modulation schematic based on the FNSSPLL and time-domain modulation operation	70
Fig. 3.13	fDAC calibration implementation details	72
Fig. 3.14	Delay-spread background cancellation algorithm.....	73
Fig. 3.15	DTC-based FNSSPLL capable of self-calibrated fractional synthesis and two-point modulation	74
Fig. 3.16	Die microphotograph	75
Fig. 3.17	A comparison between PLL output phase noise profile without and with DTC random-jump. Spurs are indicated in dBc (the DTC INL calibration is not enabled). The higher-order spurs disappear with randomization; however, the dominant spur remains only substantially suppressed and is not completely masked by noise. This is due to a larger, higher-order DTC nonlinearity (note that the DTC uses no supply regulation in these measurements)	75
Fig. 3.18	Measured output phase noise profile: (a) low-power VCO, and (b) high-power VCO. The integration range for RMS jitter calculation is 10 kHz–40 MHz and includes all spurs (including worst fractional and integer)	76
Fig. 3.19	(a) Fractional spur with and without calibration at different fractional offsets; (b) calibrated spectrum analyzer output at a deep in-band fractional channel after calibration; (c) spectrum analyzer plot at out-of-band fractional channel after calibration	77
Fig. 3.20	GMSK spectrum and EVM with self-calibration enabled (10 Mb/s, close to a 10.24-GHz fractional carrier)	79

Fig. 3.21	Measured EVM at different fractional offsets with and without calibration	79
Fig. 4.1	(a) Simplified block diagram of the proposed digital SSPTX. (b) Small signal phase-domain model: q_{nDTC} , q_{nADC} , and q_{nDAC} represent quantization noise signals from the DTC, ADC, and the DAC, respectively. ϕ_{DTC} contains, both, fractional residue compensation information and phase modulation data. $\Delta\omega_{mod}$ is the scaled frequency modulation data that are perfectly cancelled in the phase-error comparison path by the DTC. N is the PLL multiplication number	87
Fig. 4.2	Basic SSPTX phase-error detection principle in PLL and TX modes. (a) DTC delay for fractional error compensation ensures near-to-zero crossing sampling. G_{PD} is linearly proportional to the DPA output slew rate. (b) DTC delay for fractional error and PM compensation ensures near-to-zero crossing sampling. Two point AM injection ensures constant G_{PD}	87
Fig. 4.3	Generic 1024 QAM constellation (left) and distribution of used amplitudes with uniform occurrence probability of any symbol (right)	90
Fig. 4.4	Detailed overview of the SSPTX. Gray blocks serve for background calibration	92
Fig. 4.5	(a) Simulated PLL performance—quantization noise only (green) and all noise sources included (blue) at a fractional 5 GHz output ($F_{ref} = 40$ MHz). (b) Simulated output constellation of a linear system as depicted in Fig. 4.4 at 2.5 MHz modulation bandwidth and 1024 QAM with all modeled noise sources enabled	93
Fig. 4.6	DAC INL background calibration algorithm (T and B stand for the DAC's number of thermometrically and binary coded bits, respectively): (a) implementation; (b) ideal (gray) and ideal interpolated (black) INL predistortion curve for $B = 4$, and $T = 7$	95
Fig. 4.7	Ideal and real sampling event in the presence of DTC quantization error in (a) PLL mode and (b) SSPTX mode at two different AM codes (A and B). G_{PD} is proportional to the slope around the zero-crossing (i.e., sinewave amplitude). Note that the same time-quantization excess leads to different sampled amplitude during AM	97
Fig. 4.8	Phase-error detection gain background calibration: (a) block diagram implementation indicated in gray within the SSPTX; (b) background calibration with 10% gain error	97

Fig. 4.9	(a) Simulated transfer curve of a 10 bit class D^{-1} DPA; (b) Constellation diagram of a transmitter using the nonlinear DPA	99
Fig. 4.10	(a) DPA INL background calibration algorithm implementation (T = thermometrically coded bits, B = binary coded bits); (b) Ideal and interpolated AM-to-AM predistortion curve	99
Fig. 4.11	Background SSPTX calibration. (a) DTC INL background predistortion estimation. (b) DAC INL background predistortion estimation. (c) DPA INL (AM to AM) background predistortion estimation	101
Fig. 4.12	SSPTX (a) before and (b) after the background calibration is enabled	102
Fig. 4.13	DPA induced phase deviations are high-pass filtered in transfer to the output of the SSPTX, similarly as the VCO phase noise. (a) PLL filtering profiles. PLL BW set to 2.5 MHz in this example. (b) TX signal composition and high-pass filtering profile of AM-to-PM distortion in linear scale	103
Fig. 4.14	AM-to-PM suppression in a SSTPTX in comparison to a typical polar TX with AM modulator out of the PLL. (a) Modeled static AM-to-PM distortion. (b) EVM in a Polar TX with AM to PM out of the PLL. (c) EVM in a SSPTX	103
Fig. 4.15	SSPTX (a) subsampling path block diagram and (b) timing diagram	104
Fig. 4.16	Sampler schematic. Note that the figure shows only single path sampling (with a dummy for load equalization) while in reality sampling is differential at both DPA sides	105
Fig. 4.17	Signal amplifier schematic	106
Fig. 4.18	DAC schematic and operation principle	108
Fig. 4.19	Inverse class-D (D^{-1}) DPA. (a) Simplified schematic. (b) Model	109
Fig. 4.20	Typical versus harmonic rejection mixed class D^{-1} DPA. (a) Operation principle of a typical inverse class-D DPA. (b) Operation principle of an inverse class-D DPA with harmonic rejection mixing	111
Fig. 4.21	HRM signals generation. (a) Harmonic rejection mixing implementation. (b) Harmonic rejection mixing waveforms	112
Fig. 4.22	HRM class D^{-1} overview	113
Fig. 4.23	Die micrograph	114
Fig. 4.24	Measured phase noise of the PLL. (a) High-power VCO mode. (b) Low-power VCO mode	114

Fig. 4.25	Measured spectrum of the PLL. (a) Spectrum with worst case fractional-N spur and integer-N spur. (b) Spectrum during synthesis of a channel with deep in-band spur	115
Fig. 4.26	SSPTX output at -2.7 dBm average power (1024 QAM). EVM = -41.3 dB	117
Fig. 4.27	SSPTX output at 1.2 dBm average power (1024 QAM). EVM = -40.1 dB	118
Fig. 4.28	Measured predistortion correction coefficients for (a) DTC INL; (b) DAC INL; (c) DPA INL calibration (after settling)	119
Fig. 4.29	Measured EVM at 1.2 dBm average output power as a function of PLL bandwidth	119
Fig. 4.30	Measured EVM with supply level variation: with and without background calibration	120
Fig. 4.31	SSPTX output at 10 MHz modulation BW	120
Fig. 4.32	HRM performance as a function of the AM scaling factor and relative phase delay between the 50% and 25% duty-cycled LO signal. (a) AM = 1023. (b) AM = 512	121

List of Tables

Table 2.1	Performance summary and comparison of the FNSSPLL to other low-jitter fractional- N CMOS PLLs	53
Table 3.1	Performance summary of the FNSSPLL2 and comparison to other fractional- N CMOS PLLs	78
Table 3.2	Performance summary of the phase/frequency modulator and comparison to the state of the art	80
Table 4.1	EVM verification at 2.5 MHz modulation BW with increasing number of background calibration algorithms	118
Table 4.2	SSPTX performance summary and comparison to other recent TX architectures.....	122

Chapter 1

Introduction



1.1 A Transceiver with a Local Oscillator (LO) in Its Core

1.1.1 A Cartesian Transceiver

To establish a reliable communication link, a typical wireless transceiver (Fig. 1.1)¹ necessitates a precise LO that accurately up-converts and down-converts the base-band data. The link can be briefly analyzed by observing the figure in both modes, as follows. In the transmit mode, the digital base-band data, described conventionally in a two-dimensional space with orthogonal *in-phase* (I) and *quadrature* (Q) vectors, are converted into an analog signal through a DAC, which is normally cascaded with a low-pass filter. The filtered analog data stream gets up-converted to the radio frequencies in the mixer, in the mixing process with the LO generated signal (carrier). Note that the LO produces two sinusoidal signals that are shifted in phase by 90° in this Cartesian transmitter example. The modulated carrier signal is then amplified to the desired level by a Power Amplifier (PA) that drives an antenna, which can be shared between the transmit and the receive modes. The modulated Radio Frequency (RF) carrier can thus be represented as:

$$v_{\text{out}}(t) = I(t) \cdot \sin(2\pi f_{\text{LO}} t) + Q(t) \cdot \cos(2\pi f_{\text{LO}} t), \quad (1.1)$$

where $\sin(2\pi f_{\text{LO}} t)$ and $\cos(2\pi f_{\text{LO}} t)$ are the two sinusoidal signals generated by the LO at the f_{LO} frequency and the $I(t)$ and $Q(t)$ are the in-phase and quadrature data.

In the receive mode, the process is reversed. The modulated carrier is received through the antenna and amplified by means of a Low Noise Amplifier (LNA). The

¹A direct conversion, i.e., zero-Intermediate Frequency (IF) transceiver is depicted in this example.

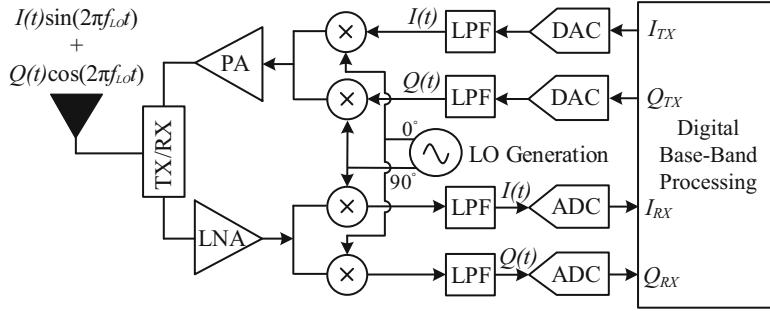
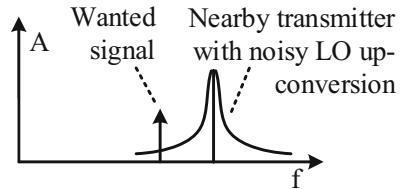


Fig. 1.1 Zero-IF transceiver concept

Fig. 1.2 LO noise effect in a transmitter

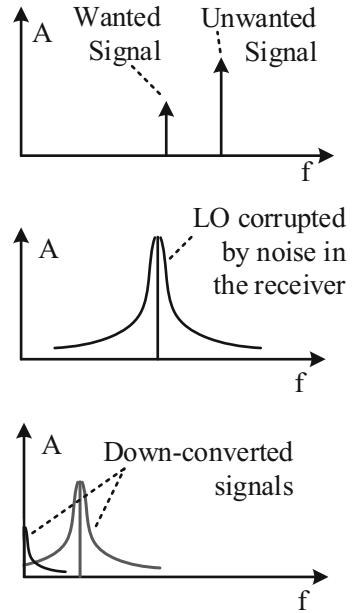


down-conversion is performed by mixing the received modulated signal with the LO (that operates at the modulated signal's central frequency), equivalently to the upper example in the transmit path. The down-conversion results with two quadrature signals, out of which it is possible to extract the original modulation data. This process is executed after low-pass filtering and analog-to-digital conversion, within a base-band digital processor.

Ideally, the LO output contains a sinusoidal tone in its output spectrum, exactly at the desired carrier frequency. In reality, the spectrum is corrupted by noise and spurs, which degrades the reception and transmission quality. For example, a transmitter that makes use of a noisy LO up-conversion has a profile as depicted in Fig. 1.2. Reception of a noiseless receiver detecting a weak nearby channel is corrupted by this transmitters' "tail" noise [Craninckx98b]. Similarly, if a receiver uses a noisy LO signal for the down-conversion, the detected signal is heavily corrupted by the noise of a nearby channel (Fig. 1.3), even if the transmitted signals are ideal. The situation is similarly problematic even if no interference (nearby blocker) is present. The LO imposes the fundamental signal-to-noise ratio limit in, both, the transmit and receive modes, reducing the communication quality and density with which the data can be packed within a limited bandwidth [Stauth08].

Besides the spectral mask of the LO, other relevant parameters are the granularity and range of the available output frequencies (note that a single LO could be used to accommodate multiple communication standards in wide frequency range that potentially necessitate nearby channel separation), LO settling time (wideband LOs are capable of faster channel hopping), area (cost), and power consumption

Fig. 1.3 LO noise effect in a receiver



(especially for mobile applications). Typically, these specifications need to be traded-off within the available LO design space to achieve the desired outcome.

1.1.2 A Polar Transmitter

In a polar transmitter, the LO is not used just for the sinusoidal signal synthesis as in a Cartesian TX. A polar TX operates with amplitude and phase modulation data that construct the output signal [Stauth08] as depicted in Fig. 1.4 and represented with:

$$v_{\text{out}}(t) = A(t) \cdot \sin(2\pi f_{\text{LO}} t + \Phi(t)), \quad (1.2)$$

where $A(t)$ is the carrier amplitude defined with:

$$A(t) = \sqrt{I(t)^2 + Q(t)^2}, \quad (1.3)$$

and $\Phi(t)$ is carrier phase defined with:

$$\Phi(t) = \arctan\left(\frac{Q(t)}{I(t)}\right). \quad (1.4)$$

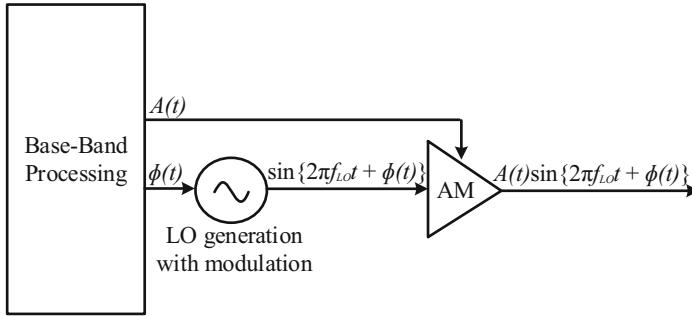


Fig. 1.4 Principle of polar transmission

In this case, the LO receives a modulation signal $\Phi(t)$ that represents the desired transmitted output phase, while an amplitude modulator handles transmitted amplitude information, $A(t)$. Note that the polar TX potentially offers some advantages over its Cartesian counterpart. It greatly simplifies the TX architecture and allows low-power implementation of continuous wave systems [Stauth08, Razavi98] since the classical mixers and DACs are eliminated. At the same time, a polar TX architecture reveals a set of design challenges that need to be addressed.

The purpose of the LO in a polar TX is not only to synthesize an accurate sinusoid but also to dynamically phase modulate its output. This process can be rather challenging in the limited LO bandwidth environment, especially in wideband modulation, and even more so since a fundamental conflict normally exists between the modulation bandwidth and generated noise. Moreover, not to degrade the quality of the transmitter signal, the phase and amplitude modulation need to be performed linearly with high resolution and with accurately matched timing, which is far from trivial in practice.

1.2 A Phase-Locked Loop (PLL) as an LO

A free-running RF oscillator (for example, a ring or an LC tank-based oscillator) is susceptible to phase noise and frequency drift with process-temperature-voltage (PVT) variations. These characteristics are obviously unacceptable for sophisticated transceivers that necessitate highly accurate LO generation. To overcome the problem, the LO is normally generated within a PLL. A PLL is a control loop that uses a high quality crystal reference to lock the (RF) oscillator, stabilize it (accurately determine the oscillation frequency), and suppress its close-in phase noise. The strict constraints in terms of its spectrum purity usually elevate a trade off with the offered functionality (modulation capabilities) and other distinctive features such as the chip area (in nanoscale CMOS) and power consumption (which is always critical in battery-limited mobile applications).

1.2.1 From an Analog to a Mixed-Signal and Digital PLL

PLLs were born as an independent research and design field back in the 1950s [Gardner66]. Enormous efforts have been directed into their development ever since, boosting and nurturing the wireless (r)evolution of the past decades. The initial solutions were purely analog and the first integer- N frequency multiplying PLLs that made use of feedback division appeared in the 1970s [Sepe70]. The low-cost, plain vanilla CMOS implementations started to appear in the late 1990s [Craninckx98a] and were mainly concentrated on achieving satisfactory phase noise performance using bulky, low-quality passives (inductors).

An example of a classical, frequency-multiplying analog PLL is shown in Fig. 1.5. At the input of the PLL, there is a Phase-error Detector (PD) that compares phase difference (time mismatch) between the input reference signal and the divided output signal coming from the VCO. The PD produces an UP or a DOWN pulse with duration proportional to the phase error detected. The UP pulse is generated if the reference leads the feedback signal and the DOWN pulse appears if the situation is reverse. The PD drives a charge pump, which pumps constant current into the loop filter during the pulse opening windows. The filter stabilizes voltage which controls the VCO, correcting for presence of phase errors between the input and the output. Note that the output frequency is defined with

$$f_{\text{out}} = N_{\text{int}} \cdot f_{\text{ref}}, \quad (1.5)$$

where N_{int} is the divider division factor and f_{ref} is the frequency of the reference signal. Two important notes with respect to this (overly) simplistic top-view on the

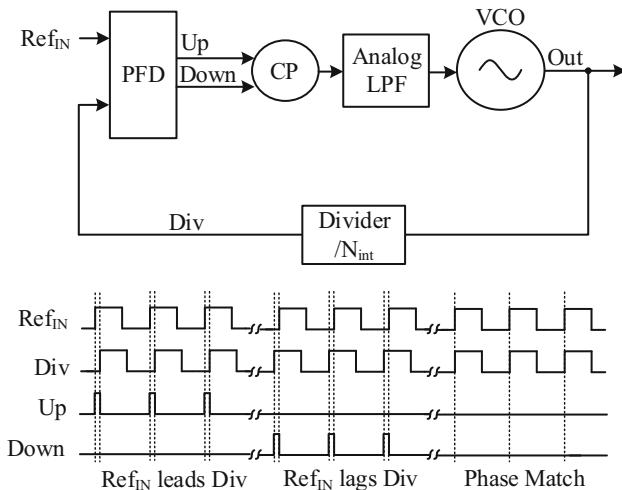


Fig. 1.5 Classical analog PLL with integer- N frequency multiplication

loop are the following: (1) for a fixed integer N , the loop can only provide output frequencies that are f_{ref} apart. For modern wireless standards, frequency granularity is essential. The only way to ensure dense channel selection with the presented loop is to use a *low* frequency reference signal. Unfortunately, this results with a need for large multiplication numbers which in turn yields to phase noise degradation. The input referred noise contribution is reduced (Sect. 1.2.3) with application of very narrow bandwidth PLL frequency filtering profiles (which also comes at the expense of oscillator noise increase). This leads to the second important note: (2) the presented loop necessitates an area consuming loop filter and large power consumption to ensure loop stability and optimal noise filtering profile (revisited in Sect. 1.2.3).

In context of modern PLL development, the introduction of mixed-signal environments that exploited principles of $\Delta\Sigma$ modulation for fractional- N frequency synthesis [Riley93] was a major breakthrough. This first digital signaling technique in PLLs enabled arbitrarily small frequency granularity at the output (independent of the reference frequency) through fast digital modulation of the division modulus.

The basic operation of a classical, mixed-signal PLL with fractional- N frequency multiplication is depicted in Fig. 1.6. The fundamental difference with the prior loop is in the programmable divider. The division number is continuously dithered between two or several integer values to achieve an average division by N_{frac} . The VCO output frequency appears thus exactly at:

$$f_{\text{out}} = N_{\text{frac}} \cdot f_{\text{ref}}. \quad (1.6)$$

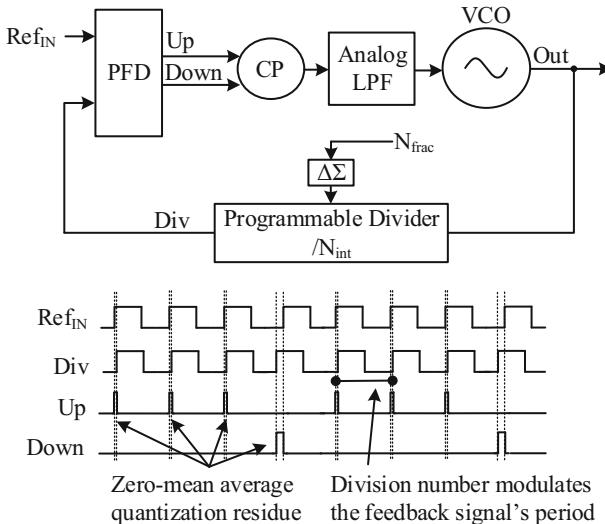


Fig. 1.6 Classical mixed-signal PLL with fractional- N frequency multiplication in phase lock

The pseudo-random dithering, applied through the $\Delta\Sigma$ modulator, shapes the division quantization error² to high frequencies. This residue (i.e., the instantaneously detected PD errors) is suppressed by the loop's low-pass filter.

The frequency granularity is independent of the reference frequency in $\Delta\Sigma$ PLLs, hence they can operate at higher reference clocks in comparison to the integer- N implementations with equal minimal output frequency separation. With smaller multiplication numbers, the reference contributes less to the phase noise (see Sect. 1.2.3), hence the loop bandwidth can be increased and the filtering capacitor size is decreased for area savings (and faster settling).

The availability of low-noise, fractional- N synthesis had a great impact on the performance of low-cost frequency synthesizers for use in consumer products. Immediately, a lot of attention has been brought to $\Delta\Sigma$ modulators in context of PLLs. They were already extensively used in oversampling data conversion systems and the history of expertise led to a fast progress.

Soon, the need for suppressing the quantization noise imposed a fundamental bandwidth versus phase noise trade-off in $\Delta\Sigma$ fractional- N PLLs. A related problem that was soon brought under attention was the mismatch, and hence nonlinearity in the charge pump-based phase detection. Any kind of nonlinearity presence in the phase comparison path leads to potential folding of the out-of-band noise. For example, high-pass-shaped $\Delta\Sigma$ spurs can be aliased to low frequency, thus degrading the PLL in-band noise [Lacaita07]. To address these problems, systems were devised with special care taken on the CP transfer function linearization and compensation for the $\Delta\Sigma$ quantization fractional residue. In [Pamarti04, Huh04, Meninger06, Gupta06] additional digital calibration circuitry is brought to address these problems. The techniques employ fractional residue DAC feed-forward compensation and/or PD/CP linearization, possibly with its adaptive calibration [Gupta06], or simply increase the modulation speed, hence reducing the quantization noise [Huh04]. These digitally intensive techniques have to be devised carefully with respect to time and amplitude mismatch (power trade-off), but can rather successfully compensate for the described problems.

The biggest breakthrough in recent PLL research is without a doubt the introduction of the All-Digital Phase Locked Loop (ADPLL) [Staszewski04] that uses time-mode signal processing instead of the analog (voltage) processing (as depicted in Fig. 1.7). In this environment, the phase error between the input and output is not measured by the classical Charge Pump (CP)-based phase/frequency detector but quantized by a time-to-digital converter (TDC). The digital information about the error is then processed in the digital loop, with a digital loop filter that drives a Digitally-Controlled Oscillator (DCO) instead of a VCO.

A digital PLL offers extreme area savings (since the loop filter can be implemented in the digital domain), ease of scaling, and compliance with digital, but

²The period of the divider's output is continuously modulated with accuracy of a single (or several) VCO periods. $\Delta\Sigma$ order determines the number of division factors used, i.e., with how many different VCO periods is the divider's output period modulated.

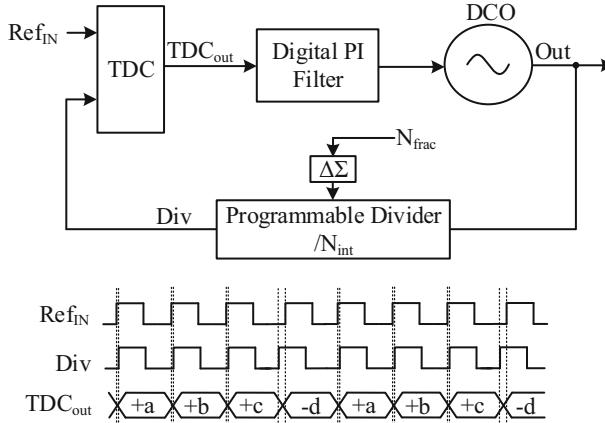


Fig. 1.7 Simplified digital PLL in phase lock. A programmable $\Delta\Sigma$ driven divider is used in the feedback as in the mixed-signal PLL from Fig. 1.6

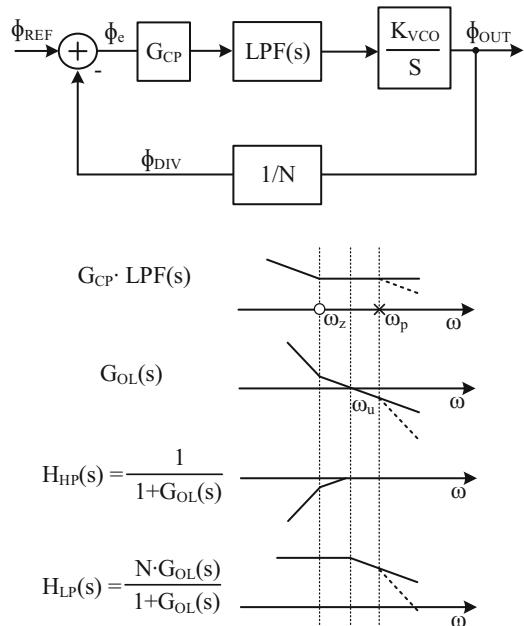
comes at the expense of design effort, since the design of time quantizers (TDCs) is far from trivial [Borremans10, Vengattara09]. Nevertheless, the ADPLL had an enormous influence on recent research [Hsu08, Temporiti10, Yao13, Kim13], revealing in its heritage space for many novel ideas and possible approaches for wideband frequency synthesis. The TDC-based phase detector, used in digitally intensive ADPLLs, has a finite resolution and introduces quantization in the phase-error information. The example from Fig. 1.7 shows a digital PLL that uses a programmable divider in the feedback, driven by a $\Delta\Sigma$ modulator (in fractional- N mode), similarly to the aforementioned mixed-signal PLL with an analog loop filter. In this architecture, the $\Delta\Sigma$ quantization noise (pseudo-periodically modulated divider period) still remains an issue that disturbs the loop. This disturbance can be compensated by a narrow-band loop, however, this comes at the expense of lessened oscillator noise suppression (Sect. 1.2.3). Similarly to the mixed-signal CP-based implementations described above, the $\Delta\Sigma$ residue can be compensated within the loop too, but here, the compensation is attained digitally. Intuitively, since the time-quantization error introduced by the divider and measured by the TDC is predictable, it can be simply subtracted as a digital number from the TDC output. For an accurate match during the compensation, TDC gain (LSB) must be precisely known, which can still represent a challenge, especially in PVT susceptible environments. A TDC is also prone to nonlinearities in the quantizing process. The consequence of such a behavior is in-loop phase noise and also, very often, spurious performance degradation. This effect is, in fact, very similar to the consequences of CP nonlinearity in classical, analog implementations. An assortment of powerful, complex digital randomization and calibration techniques has been introduced to compensate for these shortcomings and limitations.

1.2.2 Small Signal Model of a PLL

The classical CP PLL with a block diagram from Fig. 1.6 can be represented in its static operating point with a small signal phase-domain model from Fig. 1.8. In this section we briefly analyze several PLL transfer functions which prove to be relevant for understanding of the PLL's noise filtering capabilities (Sect. 1.2.3). The goal is to provide an intuitive understanding on how the oscillator phase is controlled and, importantly, how its phase noise is filtered by the loop. We will also show that the “clean-up” of the RF oscillator's noise is limited by the loop's imperfections, i.e., by the noise generated within the loop and by the noise in the reference signal.

Φ_{REF} , Φ_{VCO} , Φ_{DIV} , and Φ_e represent the small signal input and output phase of the PLL,³ divider output phase, and detected phase error, respectively. The CP-LPF cascade has typically a proportional-integral (PI) characteristic, as depicted in Fig. 1.8, below the model. In frequency domain, this means that the amplitude transfer experiences a -20 dB/decade roll-off, from the origin up to the ω_z frequency, where the zero is positioned. The dashed line in the plot represents amplitude transfer for an optionally added pole (at ω_p) in the loop filter. In this case, the cascade's amplitude transfer function continues with a -20 dB/decade roll-off after ω_p .

Fig. 1.8 Small phase-domain signal model of a classical PLL and illustration of the important transfer functions. Note that the axes are logarithmic



³Note that the loop settles to a condition where Φ_e is a zero-mean signal, i.e., where $\Phi_{\text{REF}} = \Phi_{\text{DIV}}$.

The open loop gain (G_{OL}) is defined as the gain from the PLL's input to the output of the divider (where the loop is cut for the sake of the analysis).

$$G_{OL} = G_{CP} \cdot LPF(s) \cdot \frac{K_{VCO}}{s} \cdot \frac{1}{N}. \quad (1.7)$$

G_{CP} in Eq. (1.7) represents the CP gain, $LPF(s)$ is the loop's filter transfer function, and K_{VCO} is the VCO's gain in rad/(sV). Bode plot (amplitude) of the G_{OL} transfer function (Fig. 1.8) is easily obtained by observing the plot for the $G_{CP} \cdot LPF(s)$ cascade: an additional pole in the origin (integrating behavior of the VCO) results with the G_{OL} 's initial -40 dB/decade slope that is partially straightened by the zero at ω_z (to a -20 dB/decade slope). In the presence of the aforementioned high frequency pole, the amplitude transfer drops again by -40 dB/decade after ω_p . Note that the total number of poles within the open loop determines the PLL's order (for example, with the added pole, the presented loop is of order three). Moreover, the PLL discussed here is of type-two, since in open loop it uses two poles at DC (i.e., two integrating stages). The unity gain frequency ω_u (and later the PLL's bandwidth) is defined at the position where the G_{OL} passes through the gain of one (0 dB). Note that for stability (and phase margin), ω_u needs to be at approximately four times higher frequency than ω_z and at approximately four times lower frequency than ω_p , if present [Craninckx98b].

Finally, two important transfer functions for the PLL's noise filtering capabilities (revisited in the following Sect. 1.2.3) are presented next. H_{HP} , the PLL's high-pass transfer function from the oscillator's output, to the output of the PLL is given with:

$$H_{HP} = \frac{1}{1 + G_{OL}}. \quad (1.8)$$

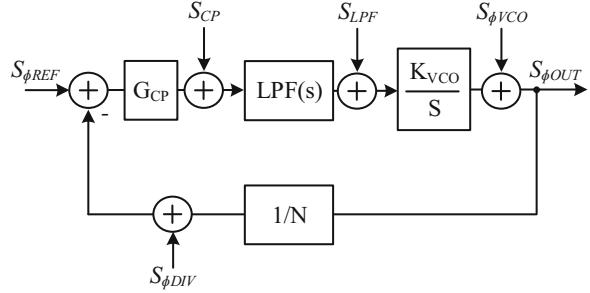
The amplitude Bode diagram of H_{HP} is depicted in Fig. 1.8. G_{OL} approaches infinity near the origin and drops with -40 dB/decade with frequency increase. For H_{HP} , the behavior is thus inverse, the amplitude is compressed to the negative infinity in the origin and it grows with an initial 40 dB/decade slope. At ω_z the slope changes to 20 dB/decade. As G_{OL} amplitude reduces at higher frequencies, H_{HP} approaches one (near ω_u).

H_{LP} is the PLL's low-pass transfer from input to the output, defined with:

$$H_{LP} = \frac{N \cdot G_{OL}}{1 + G_{OL}} \quad (1.9)$$

and depicted in Fig. 1.8. The graph can be obtained by multiplying the two previous Bode plots for G_{OL} and H_{HP} . The amplitude is simply constant (since the multiplied functions have opposing behavior) until ω_u . This frequency represents the PLL's bandwidth. At higher frequencies the amplitude transfer experiences a -20 dB/decade roll-off that increases to -40 dB/decade roll-off in the presence of the second pole at ω_p . In other words, the PLL is a low-pass filter for the input phase signal.

Fig. 1.9 Phase noise generation in the classical mixed-signal PLL. Equivalently, the phase noise generation in a typical digital PLL is depicted in 1.11 and analyzed later



1.2.3 Phase Noise in PLLs

The output spectral mask of the frequency synthesizer is probably the most important characteristic. In this section, we briefly review phase noise generation and filtering properties within a PLL that deteriorate the ideally pure sinusoidal oscillator output. The analog PLL is compared to a digital architecture in its phase noise filtering capabilities.

In phase lock (Fig. 1.9), there is ideally no input or output phase deviation over time. In reality, the VCO oscillation is imperfect, its frequency, i.e., phase accumulation fluctuates over time. This effect can be described by phase noise and $S_{\Phi_{VCO}}$ represents the phase noise spectrum at the VCO output (measured in rad^2/Hz). Other analog building blocks within the loop generate noise as well, S_{CP} and S_{LPF} are the charge pump and loop filter noise spectra, expressed in A^2/Hz and V^2/Hz , respectively [Lacaita07]. The divider is similarly imperfect, its output phase noise spectra is indicated by $S_{\Phi_{DIV}}$. Note that $S_{\Phi_{DIV}}$ can contain $\Delta\Sigma$ modulated quantization noise in fractional division. Not even the crystal oscillator reference is ideal, $S_{\Phi_{REF}}$ represents its phase noise spectrum (in rad^2/Hz).

For a better understanding of the noise filtering within the loop, it is useful to segregate the noise sources, referring them to the input or the output of the system in an open loop configuration [Lacaita07]. The open loop reference, divider, and charge pump noise can be referred to the input of the system as:

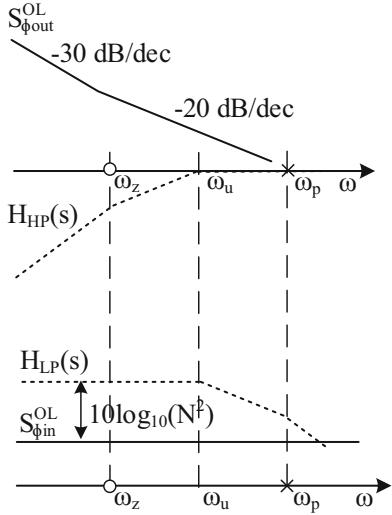
$$S_{\Phi_{in}}^{OL} = S_{\Phi_{REF}} + S_{\Phi_{DIV}} + \frac{S_{CP}}{(G_{CP})^2} \quad \left[\frac{\text{rad}^2}{\text{Hz}} \right], \quad (1.10)$$

where G_{CP} is the charge pump gain. Similarly, $S_{\Phi_{VCO}}$ and S_{LPF} can be referred to the output of the system:

$$S_{\Phi_{out}}^{OL} = S_{\Phi_{VCO}} + S_{LPF} \cdot \left(\frac{K_{VCO}}{2\pi f} \right)^2 \quad \left[\frac{\text{rad}^2}{\text{Hz}} \right], \quad (1.11)$$

where K_{VCO} is the gain of the VCO's varactor (in Hz/V) and f is the offset frequency.

Fig. 1.10 Output and input referred noise profiles (solid line) with corresponding PLL filtering (dashed line)



The PLL output phase noise spectrum can then be expressed as:

$$S_{\Phi_{\text{out}}}^{\text{PLL}} = S_{\Phi_{\text{out}}}^{\text{OL}} \cdot |H_{\text{HP}}|^2 + S_{\Phi_{\text{in}}}^{\text{OL}} \cdot |H_{\text{LP}}|^2. \quad (1.12)$$

From Eq. (1.12) it can be seen that the output referred noise is high-pass filtered. The case is completely opposite for the input referred noise which is low-pass filtered in transfer to the output. The art of optimal noise filtering (for a given PLL type and order) is in H_{HP} , i.e., H_{LP} bandwidth selection (ω_u). The goal is to minimize the overall noise contribution to the output.

Typically, the dominant output referred noise is the RF oscillator's phase noise ($S_{\Phi_{\text{VCO}}}$). A classical LC-based oscillator's phase noise profile is indicated in the top part of Fig. 1.10. Note that the phase noise increases towards the origin [Razavi96] with a 20 dB/decade slope (or a 30 dB/decade slope as the 1/f noise over-dominates the white noise contribution within the oscillator). Thanks to the closed loop, this noise is suppressed by H_{HP} below the PLL's bandwidth ω_u . Choosing a larger PLL bandwidth is thus beneficial for VCO phase noise filtering.

The dominant input referred noise is typically the CP ($S_{\Phi_{\text{CP}}}$) noise and the reference phase noise ($S_{\Phi_{\text{ref}}}$). In fractional multiplication, the added quantization noise in $S_{\Phi_{\text{DIV}}}$ imposes a large limitation as well. These sources are filtered by H_{LP} hence, choosing a smaller PLL bandwidth reduces their contribution to the overall output noise. Note that the input referred noise is multiplied by N^2 in the transfer to the output at low (in-band) frequency offsets. This, for example, means that operating at high multiplication numbers normally necessitates small loop bandwidth (which is unfortunate for the oscillator noise filtering).

By observing Eqs. (1.10) and (1.12), it can be noted that increasing the phase-error detection gain, i.e., CP gain (G_{CP}) in this example, helps with its in-band noise

suppression.⁴ To ensure the same phase margin, i.e., pole-zero separation, increase of the CP gain (phase-error detection gain) must be accompanied by increase of the filtering capacitor (area cost) or decrease of the K_{VCO} (frequency tracking range cost). This example serves as a proof that the PLL system properties (stability, phase margin, lock time) need to be co-optimized with its phase noise filtering properties—and the design process is normally iterative.

Phase noise generation can be analyzed for a digital PLL of the same type and order (Fig. 1.11). A digital PLL, as described in Sect. 1.2.1 uses no CP-based phase/frequency error detection but a TDC in the comparison path. All the information processing following the TDC up to the DCO is digital and thus “noiseless.” Even though a digital loop avoids addition of CP or analog filter noise, it introduces *quantization* noise. Namely, a TDC has finite resolution and it quantizes the input information with limited accuracy—producing just like any data converter, a quantization error. The quantization noise spectrum is indicated by $S_{\phi_{qTDC}}$, and can be assumed white. Similarly, the DCO which can be modeled as a VCO driven by a DAC operates in the presence of quantization residue, as well. The (otherwise noiseless) DAC then outputs a quantization noise spectrum indicated by S_{qDAC} . In comparison to an analog loop, a digital PLL can operate with a large error detection gain G_{TDC} , i.e., with a high resolution TDC, without any area cost, since the digital loop filter size is small (insignificant in comparison to a corresponding analog filter equivalent). Design of high resolution TDCs is far from trivial however, and they typically come at the cost of added power consumption and circuit complexity. Nevertheless, the trade-offs improve with technology scaling. This is similarly true for the digitally controlled oscillation that typically makes use of fine capacitive banks for frequency tuning (achievable metal density improves with technology scaling). Since the TDC noise is input referred, it is scaled by the PLL multiplication

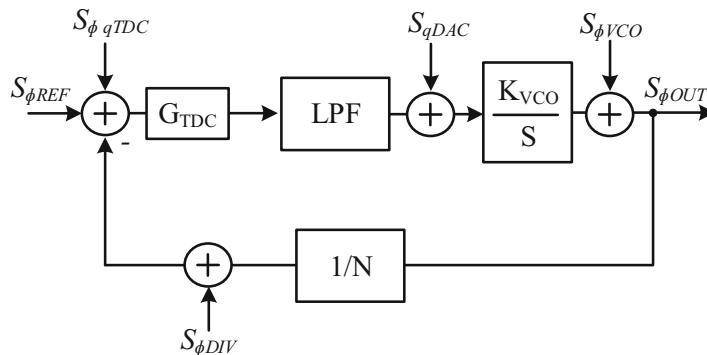


Fig. 1.11 Phase noise generation in a typical digital PLL

⁴The improvement is somewhat limited since for higher gain G_{CP} , the CP current must increase which also leads to a higher local noise generation of the CP.

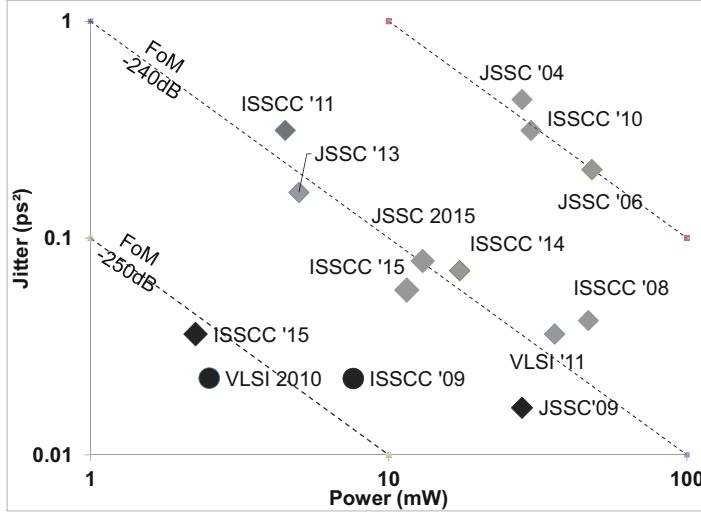


Fig. 1.12 Comparison of recent integer- N (black) and fractional- N (gray) PLLs in power-vs-RMS jitter trade-off FOM. Note that better results appear in bottom left of the figure. The black circles indicate the subsampling PLL architecture (see Sect. 1.3.1)

number in transfer to the output. It dominates the in-band noise (just like the CP noise for the analog PLL), while out-of-band it is filtered by H_{LP} .

The phase noise generation within the loop as described above complicates in the presence of nonlinear phase-error detection path, i.e., PFD/CP nonlinearities in analog PLLs [Meninger05] and TDC nonlinearity in digital PLLs [Straayer08], especially in fractional- N synthesis where the loop needs to handle the division quantization error residue (see Sect. 1.2.1), i.e., larger PD input range.

The $\Delta\Sigma$ modulated $S_{\Phi_{DIV}}$ noise injection into the system becomes in the presence of the nonlinear phase-error detection “colored” by the digital division factors used. This type of correlation between digital inputs and analog output errors leads to generation of spurious content [Lacaita07]. The recent state of the art offers a rich variety of solutions for LO generation, both in digital and in analog domain that deal with this issue in several ways (see Sect. 1.2.1). Nevertheless, the solemnity of the described problem is probably best visible by observing the recent art (Fig. 1.12) in frequency synthesis, where a significant performance gap still exists between the integer- N and fractional- N frequency synthesis.

1.3 Motivation and Research Objectives

1.3.1 A Subsampling PLL

In the early stage of the project this book was based on, central interest was brought to the subsampling PLL introduced in 2009 by Gao [Gao09]. A subsampling loop

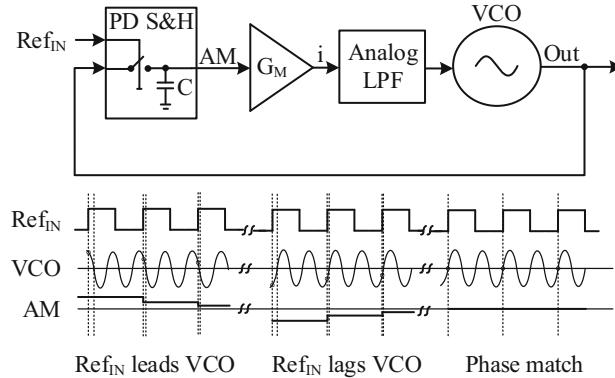


Fig. 1.13 Operation of a subsampling PLL

fundamentally differs from the classical analog PLL in the following characteristics: (1) the loop is divider-less; (2) the loop does not use a classical CP in the phase comparison path. The subsampling PLL operation is briefly analyzed starting from Fig. 1.13.

The loop operates by direct subsampling of the high frequency RF output waveform at the reference rate. The phase-error detection is performed by the PD sample and hold circuit, which transforms time information (phase error) into voltage, storing it across a capacitor C . The voltage amplitude is approximately linearly proportional to the detected phase error.⁵ If there is no mismatch between the input and the output phase, the output sinewave is sampled exactly at the zero-crossing and the stored voltage amplitude is zero. The sampled voltage biases a transconducting stage (G_M), controlling the amplitude of the current pumped into the loop filter. The loop filter can be implemented in the same fashion as in the classical analog PLL. Its purpose is to provide stable drive for the controlled oscillator that corrects for existing imbalance between the input and the output phase.

The attractive features of a subsampling PLL, compared to a classical analog PLL, become apparent in analysis of the phase noise generation within the loop (Fig. 1.14). The PLL operates without a divider in the feedback from the output. Besides the obvious benefit of divider noise elimination, the charge pump (transconductor in this case) noise in the subsampling environment becomes independent of the PLL multiplication factor N [Gao09]. In other words, the loop noise is not multiplied by N^2 in transfer to the output, as in the case of a classical analog PLL. Note that this is not the case for the reference phase noise, which is still virtually up-converted by N (as in a classical loop).

⁵Valid for small signal approximation of a sinusoid around zero.

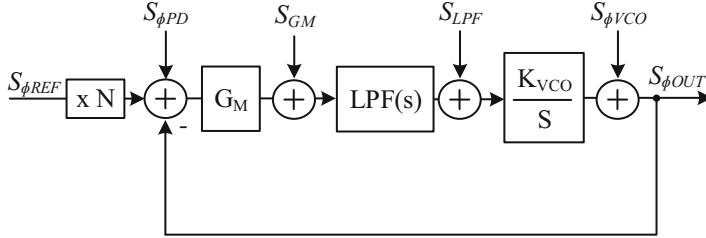


Fig. 1.14 Phase noise generation in a subsampling PLL

The noise in the loop is additionally suppressed, thanks to the very high phase-error detection gain (see Sect. 1.2.3). Namely, in the illustrative Fig. 1.13 the ratio between the input and output frequency is only two, but in practice the multiplication to RF necessitates $N > 100$. High frequency signals have high slew-rate around the zero-crossing, which directly translates to detection of high amplitudes even for very small deviations of phase near the sampling event.

Thanks to the loop noise compression features, the reference noise (and input buffering) prove to contribute with more than 70% of the overall in-band noise at the subsampling PLL output [Gao09]. The subsampling PLL is thus mostly limited by the reference noise (and the RF oscillator noise), which is in fact the desired outcome in high-performing frequency synthesis involving an architecture whose loop components do not degrade the output spectral purity. After all, the input crystal oscillator is the lowest reference noise source available.

As discussed in Sect. 1.2.3 to maintain the type-two PLL's zero-pole separation for optimal phase margin, the loop filtering needs to be adjusted for the large, subsampling PD gain. Since the dominant in-band noise source is the typically low reference noise, the subsampling PLL enables wideband loop operation (in MHz range instead in the range of several tens of kHz) which is beneficial for the oscillator noise filtering (See Sect. 1.2.3), but also for the filter area savings (since the filtering capacitor needs not to be excessively large). Still, to avoid a relatively large capacitor in the filter, the detection gain can be on purpose reduced by time-invariant pulsing of the transconductor (without a large compromise in in-band phase noise). The current is then not fed to the loop filter continuously over a complete reference cycle, but only pulsed during a fraction of the reference period [Gao09].

Thanks to the amplitude-based phase-error detection, the PD path becomes resilient to the typical nonlinearities in the charge pump. Mismatch between UP and DOWN pulses from the classical PLL is not applicable in a subsampling PLL [Gao10] since the loop settles inherently to a condition of zero-mean instantaneous output current (i) (Fig. 1.13). In integer- N mode, charge pump nonlinearity is not a fundamental issue, nevertheless, this feature lies in contrast to classical PLLs, which due to the CP nonlinearity often suffer from unwanted output spurs in fractional- N modes (this is similarly true for digital PLLs that make use of imperfect, and often nonlinear TDCs). These characteristics will be revisited soon.

Elimination of the divider leads to power-efficiency improvements, especially at high output RF frequencies that need large division for down-scaling and comparison to the reference period. For a fair comparison, one should note that the subsampling loop cannot discriminate between frequencies that are integer- N apart and therefore needs a frequency acquisition circuit (with a divider) in parallel to the main loop [Gao09]. The auxiliary loop can however be automatically disabled in phase lock for power savings.

All these features made the proposed subsampling architecture to stand out within the state of the art (see Fig. 1.12), which was mainly overrun by digital and digitally intensive solutions. Even though the loop is in essence analog, this relatively simple solution is compatible with modern digital CMOS technologies that typically offer good switches, and dense metal capacitors. Finally, the subsampling PLL achieves a very high PLL Figure of Merit (FOM) (that trades off jitter performance for power) and shows extreme potential⁶ for LO synthesis, with one fundamental limitation: it enables exclusively integer- N operation, making it in principle useless for the modern wireless standards that necessitate high frequency granularity. Moreover, since the loop has no modulation capabilities, it cannot be used in its original form for any kind of polar processing.

1.3.2 *Objectives of the Book*

The main objective of this book is to build upon the attractive characteristics of a high-performance subsampling PLL (as described in the previous subsection) and to extend its functionality with fractional- N operation and phase modulating capabilities. A subsampling PLL architecture offers inherently better integer- N environment, i.e., a better starting point for accurate frequency synthesis. The *fractionalization* needs to be carefully approached, not to degrade the original performance of the loop (during fractional synthesis or modulation). With this assumption in mind, it should be possible to enable low-noise fractional synthesis with state-of-the-art FOMs.

The history of modern PLL development (Sect. 1.2.1) shows a great amount of design effort pushed for similar purpose into the classical (both analog and digital) loop development. There is a large variety of techniques in the state of the art that can be investigated to reach this goal. The biggest opportunity in this context might fall in line with the latest design paradigm shift from the voltage-domain processing to time-domain processing that happened with the introduction of modern ADPLLs.

⁶Maybe the best proof for the *potential* this system has is in today's frequency synthesis state of the art. The current record holders in PLL FOM, presented at ISSCC 2018, are two (sub)sampling PLLs [Sharkia18, Sharma18]. Moreover, the subsampling architecture has been successfully and widely applied for high-performance integer- N LO synthesis since its introduction, in digitally intensive configurations [Ru13], ring-oscillator-based [Sogo12] loops, at millimeter frequencies [Szortyka14], and in automotive radar systems [Yi13].

The approach is built on the assumption that in modern CMOS, the available voltage headroom for quantization continuously reduces, which is not necessarily the case for the timing information, which is independent or even improves in modern digital nodes [Staszewski06]. This approach led to designs with best fractional- N PLL FOMs currently available.

Can a subsampling PLL enable fractional- N synthesis and phase modulation? How to ensure that the pristine output spectrum of an integer- N subsampling PLL remains uncorrupted by quantization noise in fractional- N frequencies? Is the application of highly efficient time-domain processing techniques (typically available in TDC-based digital PLLs) an option? Assuming that the loop becomes capable of fractional- N synthesis with high performance, could it then also be used for phase/frequency modulation in a Polar TX environment? Can accurate digital-to-modulated output conversion be achieved through a subsampling loop, despite challenges such as injection gain/time mismatch, nonlinearities, and cross-distortion mechanisms between the amplitude and phase modulation paths? And finally, can the loop maintain its performance across PVT variations, and operate reliably in a wide range of applications?

The goal is to explore systems which can challenge or outperform recent art in the field, extending successful principles of subsampling to new areas of LO generation and modulation. The questions above are answered in this book throughout the following chapters.

1.4 Book Outline

This book presents three 28 nm CMOS ICs that enable fractional frequency synthesis, self-calibrated fractional synthesis/wideband phase modulation, and highly spectrally efficient polar transmission, respectively, all based on a subsampling phase-error detection core that ensures low-phase noise operation and power efficiency.

In Chap. 2, we start by exploring principles of fractional synthesis using a *digital-to-time converter* (DTC) in a subsampling PLL. It will be shown that the benefits of the time-domain signal processing in frequency synthesis (well known from the area of digital PLLs) are applicable in the subsampling environment, without loss of the original loop's PLL properties. A DTC, as a core building block of the proposed PLL, transforms digital information into time delay—which is used for reference clock modulation, ensuring zero-crossing sampling even with non-integer ratio's between the input and the output frequency. The DTC should bring as little as noise as possible into the circuit; therefore, it is carefully optimized for low random and quantization noise operation. The presented PLL complies with the most advanced digital PLLs (at the moment of publication), but it still operates with a performance gap between integer and fractional operation mode. Confronting this limitation paved the road for innovation described in the following chapters.

In Chap. 3, we build upon the presented fractional synthesizer by extending its digital processing capabilities. These are used to deal with imperfections of the DTC (in terms of its nonlinearity) that proved to be the most sensitive block in the phase-error comparison path. To reduce fractional spurs and noise folding, we develop techniques for background self-calibration. These techniques are purely digital and based on single bit PD error extraction. The principles of predistortion are used for minimization of the performance gap between integer- N and fractional- N operation modes, in PVT variation susceptible environments. Moreover, the basic loop is enhanced for subsampling PLL-based frequency/phase modulation. We propose a background-calibrated, two-point injected phase modulator that efficiently operates with wide modulation bandwidths. Both, the frequency synthesizer and the phase modulator challenge recent art in the field in terms of measured jitter-power trade-off and measured EVM.

In the final technical Chap. 4 of this book, we present how to digitize the synthesizer and expand the background-calibrated phase/frequency modulator into a full polar TX. The digitization is carefully developed so that the added digital PLL quantization noise does not degrade the overall performance, and even more importantly, so that the loop still operates with virtually no performance gap between the integer and fractional synthesis modes. The phase modulator necessitates a block with amplitude modulating (AM) capabilities for polar signal processing (see Fig. 1.4). We implement the amplitude modulator using a nonlinear digital PA (DPA) class. In contrast to a classical polar TX, this DPA is implemented as a part of the phase-locked loop. Thanks to the subsampling phase-error detection core, it is possible to track amplitude of the output oscillation in the background and consequentially any nonlinearities in the AM-to-AM conversion. On top of this, the presented loop compresses undesired cross-correlation mechanisms of AM-to-PM distortion. To distinguish it from a typical polar TX architecture, we name the presented system as subsampling polar transmitter. This work achieves highly linear, low-noise digital-to-transmitted signal conversion which in turn enables complex modulation and high spectral efficiency with large data throughput.

In the final Chap. 5, we summarize the presented material, list the book contributions, and propose a future outlook with guidelines on how to improve the material.

References

- [Borremans10] J. Borremans, K. Vengattaramane, V. Giannini, B. Debaillie, W. Van Thillo, J. Craninckx, A 86 MHz-12 GHz digital-intensive PLL for software-defined radios, using a 6 fJ/step TDC in 40 nm digital CMOS. *IEEE J. Solid-State Circuits* **45**(10), 2116–2129 (2010)
- [Craninckx98a] J. Craninckx, M.S. Steyaert, A fully integrated CMOS DCS-1800 frequency synthesizer. *IEEE J. Solid-State Circuits* **33**(12), 2054–2065 (1998)
- [Craninckx98b] J. Craninckx, M. Steyaert, *Wireless CMOS Frequency Synthesizer Design* (Kluwer Academic Publishers, Dordrecht, 1998)

- [Gao09] X. Gao, E. Klumperink, M. Bohsali, B. Nauta, A low noise sub-sampling PLL in which divider noise is eliminated and PD/CP noise is not multiplied by N^2 . IEEE J. Solid-State Circuits **44**(12), 3253–3263 (2009)
- [Gao10] X. Gao, E. Klumperink, G. Soccia, M. Bohsali, B. Nauta, Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector. IEEE J. Solid-State Circuits **45**(9), 1809–1821 (2010)
- [Gardner66] F.M. Gardner, *Phaselock Techniques* (Wiley, London, 1966)
- [Gupta06] M. Gupta, B.-S. Song, A 1.8-GHz spur-cancelled fractional-N frequency synthesizer with LMS-based DAC gain calibration. IEEE J. Solid-State Circuits **41**(12), 2842–2851 (2006)
- [Hsu08] C.-M. Hsu, M.Z. Straayer, M.H. Perrott, A low-noise wide-BW 3.6-GHz digital $\Delta\Sigma$ fractional-N frequency synthesizer with a noise-shaping time-to-digital converter and quantization noise cancellation. IEEE J. Solid-State Circuits **43**(12), 2776–2786 (2008)
- [Huh04] H. Huh, Y. Koo, K.-Y. Lee, Y. Ok, S. Lee, D. Kwon, J. Lee, J. Park, K. Lee, D.-K. Jeong et al., A CMOS dual-band fractional-N synthesizer with reference doubler and compensated charge pump, in *2004 IEEE International Solid-State Circuits Conference, Digest of Technical Papers. ISSCC* (IEEE, Piscataway, 2004), pp. 100–101
- [Kim13] H.S. Kim, C. Ornelas, K. Chandrashekhar, D. Shi, P.-E. Su, P. Madoglio, W.Y. Li, A. Ravi, A digital fractional-N PLL with a PVT and mismatch insensitive TDC utilizing equivalent time sampling technique. IEEE J. Solid-State Circuits **48**(7), 1721–1729 (2013)
- [Lacaita07] A.L. Lacaita, S. Levantino, C. Samori, *Integrated Frequency Synthesizers for Wireless Systems* (Cambridge University Press, Cambridge, 2007)
- [Meninger05] S. Meninger, M. Perrott, Low Phase noise, High bandwidth frequency synthesizer techniques, Ph.D. dissertation, Massachusetts Institute of Technology, 2005
- [Meninger06] S.E. Meninger, M.H. Perrott, A 1-MHZ bandwidth 3.6-GHz 0.18- μ m CMOS fractional-N synthesizer utilizing a hybrid PFD/DAC structure for reduced broadband phase noise. IEEE J. Solid-State Circuits **41**(4), 966–980 (2006)
- [Pamarti04] S. Pamarti, L. Jansson, I. Galton, A wideband 2.4-GHz delta-sigma fractional-NPLL with 1-Mb/s in-loop modulation. IEEE J. Solid-State Circuits **39**(1), 49–62 (2004)
- [Razavi96] B. Razavi, A study of phase noise in CMOS oscillators. IEEE J. Solid-State Circuits **31**(3), 331–343 (1996)
- [Razavi98] B. Razavi, R. Behzad, *RF Microelectronics*, vol. 2 (Prentice Hall, New Jersey, 1998)
- [Riley93] T.A. Riley, M.A. Copeland, T.A. Kwasniewski, Delta-sigma modulation in fractional-N frequency synthesis. IEEE J. Solid-State Circuits **28**(5), 553–559 (1993)
- [Ru13] Z. Ru, P. Geraedts, E. Klumperink, X. He, B. Nauta, A 12GHz 210fs 6mW digital PLL with sub-sampling binary phase detector and voltage-time modulated DCO, in *2013 Symposium on VLSI Circuits (VLSIC)* (IEEE, Piscataway, 2013), pp. C194–C195
- [Sepe70] R.B. Sepe, Frequency multiplier and frequency waveform generator, U.S. Patent 3,551,826, 29 Dec 1970
- [Sharkia18] A. Sharkia, S. Mirabbasi, S. Shekhar, A 0.01 mm² 4.6-to-5.6GHz sub-sampling type-I frequency synthesizer with -254dB FOM, in *2018 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2018), pp. 256–257
- [Sharma18] J. Sharma, H. Krishnaswamy, A dividerless reference-sampling RF PLL with -253.5dB jitter FOM and <-67dBc reference spurs, in *2018 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2018), pp. 257–258

- [Sogo12] K. Sogo, A. Toya, T. Kikkawa, A ring-VCO-based sub-sampling PLL CMOS circuit with -119 dBc/Hz phase noise and 0.73 ps jitter, in *2012 Proceedings of the ESSCIRC (ESSCIRC)* (IEEE, Piscataway, 2012), pp. 253–256
- [Staszewski06] R.B. Staszewski, P.T. Balsara, *All-Digital Frequency Synthesizer in Deep-Submicron CMOS* (Wiley, London, 2006)
- [Staszewski04] R.B. Staszewski, K. Muhammad, D. Leipold, C.-M. Hung, Y.-C. Ho, J.L. Wallberg, C. Fernando, K. Maggio, R. Staszewski, T. Jung et al., All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS. *IEEE J. Solid-State Circuits* **39**(12), 2278–2291 (2004)
- [Stauth08] J.T. Stauth, *Energy Efficient Wireless Transmitters: Polar and Direct-Digital Modulation Architectures* (ProQuest, Ann Arbor, 2008)
- [Straayer08] M.A. Straayer, Noise shaping techniques for analog and time to digital converters using voltage controlled oscillators. Ph.D. dissertation, Massachusetts Institute of Technology, 2008
- [Szortyka14] V. Szortyka, Q. Shi, K. Raczkowski, B. Parvais, M. Kuijk, P. Wambacq, 21.4 A 42mW 230fs-jitter sub-sampling 60GHz PLL in 40nm CMOS, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, Piscataway, 2014), pp. 366–367
- [Temporiti10] E. Temporiti, C. Weltin-Wu, D. Baldi, M. Cusmai, F. Svelto, A 3.5 GHz wideband ADPLL with fractional spur suppression through TDC dithering and feedforward compensation. *IEEE J. Solid-State Circuits* **45**(12), 2723–2736 (2010)
- [Vengattara09] K. Vengattaramane, J. Borremans, M. Steyaert, J. Craninckx, A gated ring oscillator based parallel-TDC system with digital resolution enhancement, in *IEEE Asian Solid-State Circuits Conference, A-SSCC 2009* (IEEE, Piscataway, 2009), pp. 57–60
- [Yao13] C.-W. Yao, A.N. Willson, A 2.8–3.2-GHz fractional- N digital PLL with ADC-assisted TDC and inductively coupled fine-tuning DCO. *IEEE J. Solid-State Circuits* **48**(3), 698–710 (2013)
- [Yi13] X. Yi, C.C. Boon, J. Sun, N. Huang, W.M. Lim, A low phase noise 24/77 GHz dual-band sub-sampling PLL for automotive radar applications in 65 nm CMOS technology, in *2013 IEEE Asian Solid-State Circuits Conference (A-SSCC)* (IEEE, Piscataway, 2013), pp. 417–420

Chapter 2

A Digital-to-Time-Converter-Based Subsampling PLL for Fractional Synthesis



2.1 Introduction

The phase noise of the frequency synthesizer sets a limit to the achievable data rate or to the total radio power consumption, as one can often be traded for the other. This is equally true for high-throughput applications like LTE-Advanced or for sub-mW Internet-of-Things nodes. Traditionally, a frequency synthesizer has been an analog system, but analog performance, unfortunately, had degraded with last decade of CMOS scaling. This trend is only expected to continue with the development of logic-centric technologies.

The answer to this problem from the PLL community came in the form of the *loop digitization*. In other words, the PLL had to become as digital in nature as possible. The breakthrough and success of the first generation all-digital PLLs (ADPLL) [Staszewski04] has paved the way towards potentially scalable synthesizers. The intensive research of digital PLLs has led to systems such as [Tasca11] that deliver excellent performance at low power in an attractive, self-calibrating environment. And, although this trend of migrating functionality into the digital domain will only continue, the remaining analog blocks of a digital PLL, such as TDC and VCO, still form the fundamental phase noise performance bottleneck.

In this chapter, we build upon the subsampling PLL proposed in 2009 in [Gao09b], which stands in stark contrast to digital PLLs with its *analog* core. This loop still remains a leading design in recent state of the art in terms of phase noise as well as FOM. This competitive performance is achieved by removing the two classical contributors to phase noise in an analog PLL: the frequency divider and the charge pump. Additionally, thanks to a very high phase-error detection gain,

remaining noise generators within the loop are strongly suppressed. Interestingly, the advantage of an analog subsampling PLL is that it does not require high-performance analog building blocks. There is also no need for high-accuracy matching (like in classical charge pumps and phase/frequency error detectors) or for precise timing.

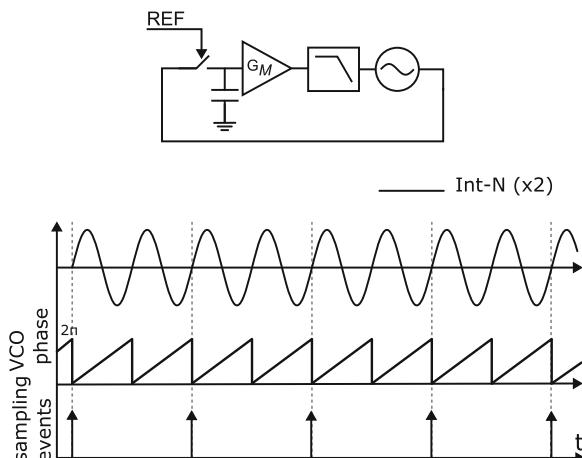
Unfortunately, the original subsampling PLL did not enjoy much popularity in its original form. Its inherent Integer- N operation is simply unacceptable for modern communication standards that necessitate high-frequency granularity.

In this chapter, we propose a solution that enables fractional- N operation of a subsampling PLL. This is done using a digital-to-time converter (DTC) in the system that cancels accumulated fractional residue during operation. Note that cancellation of the systematic errors is equivalent to phase-error measurements using a TDC in digital PLLs. We will show how DTC enables fractional- N lock, while retaining the key benefits of subsampling operation.

The subsampling PLL is in essence a system containing a VCO, loop filter, a switch, and a capacitor (Fig. 2.1). Note that even in deeply scaled technologies, where analog performance degrades, these components continue improving. Adding a DTC into the system does not change this, since it can be implemented as a few inverters and a capacitor bank. In this way, we arrive at a simple solution that can be truly technology independent.

This chapter is based on the material reported in [Markulic14, Raczkowski15] and is organized as follows. Section 2.2 explains the time-domain operation of a subsampling PLL and introduces a method for enhancing it to achieve a fractional- N lock. Section 2.3 examines the relevant system-level challenges that arise when using a practical, performance-limited DTC. Section 2.4 describes the circuit implementation of the fractional- N subsampling PLL and Sect. 2.5 presents the performance of the fabricated test chip. Finally, conclusions are drawn in Sect. 2.6.

Fig. 2.1 General system of a subsampling PLL with example timing



2.2 Fractional-N Operation of a Subsampling PLL

To understand the concept behind enhancing a subsampling PLL for fractional- N multiplication, we start with analysis of the basic, integer- N subsampling PLL.

2.2.1 Time-Domain Analysis of a Subsampling PLL

The starting point to the analysis is the basic subsampling PLL consisting of a VCO, a sampler that operates at a reference clock, a transconductor (G_M), and a Low Pass Filter (LPF) (Fig. 2.1). Compared to the classical mixed-signal PLL, there is no frequency divider and the PFD/charge pump are replaced by the sampler and the G_M . Phase detection happens by direct sampling of the VCO waveform with a rate dictated by the reference clock. The sampled voltage is converted into an error current, which is fed to the LPF. In the PLL's lock state, the phase error between the VCO and the reference is zero, and hence the sampled voltage is zero and no current flows to the LPF. No correction is necessary. In presence of a phase error, the error current is nonzero and a function of the phase error. The relation between the phase error $\Delta\phi_{VCO}$ the output current i_{G_M} is sinusoidal [Gao09b], though for small phase deviations it can be defined simply as $\beta_{SS} = \frac{\Delta i_{G_M}}{\Delta\phi_{VCO}} = A_{VCO} \cdot g_m$, where A_{VCO} is the amplitude of the VCO and g_m is the transconductance of the G_M . Contrary to a classical PLL, the phase detection circuits do not need to have high analog performance. Sampler nonlinearity or clipping can be tolerated, since the sampling point is always close to the zero-crossing of the sampled voltage. Furthermore, the output charge is produced by the G_M and the error information lies in the magnitude and sign of the resulting current and not in a variable current pulse duration (like in a conventional PLL). Finally, leakage of the sampler is corrected by the loop if the opening of the G_M output happens always with the same delay with respect to the sampling event [Gao10].

It is evident that the subsampling loop can only synthesize integer- N multiplications of the reference frequency. There is no phase modulation in the loop whatsoever, and the only stable point is when the zero-crossings of the VCO waveform match the timings of the edges of the reference (note that the loop cannot even distinguish between frequencies integer- N apart, which is also a problem addressed later). There is no divider in this loop, which means that the traditional method of applying $\Delta\Sigma$ modulation to the divider [Riley93] is out of reach.

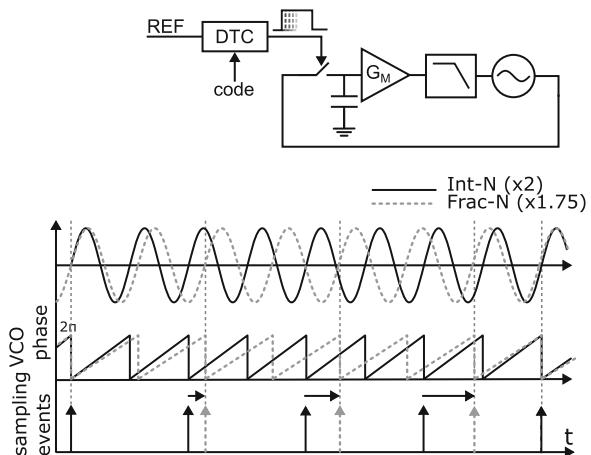
2.2.2 Enhancement of a Subsampling PLL to Enable Fraction-N Mode Operation

The basic subsampling PLL cannot synthesize fractional- N frequencies, because it lacks any phase modulation mechanism in the loop. There are in principle four nodes that can be considered in Fig. 2.1 to introduce a phase modulating element.

Addition of a divider in the feedback path would consequentially eliminate the PLL's detection gain advantage which is certainly undesirable. Moreover, a divider in the loop contributes to noise and power. Application of any costly operation at RF is not an attractive solution since it inevitably leads to power consumption and phase noise increase, as well. A residue DAC could be used, as in [Swaminatha07] to correct the phase errors, but the solution requires a DAC matched to the loop gain and it might be cumbersome considering the small linear phase-error detection range of a subsampling PLL. Final option is phase modulation of the reference clock, which is, in fact, equivalent to modulating a frequency divider in the feedback path of a classical PLL. Instead of adapting the phase of the divided fractional- N signal to match phase of the reference, we adapt the phase of the reference to match the phase of the fractional- N frequency of the VCO.

Assume, as an example, that the VCO works at a target fractional- N frequency that is different from an integer- N by 0.25 (e.g., in Fig. 2.2, $N = 1.75$). In the first cycle, the sampling event appears at the same time as in the integer- N mode. Then, in the second cycle, a timing error of $0.25 \cdot T_{VCO}$ is present. To still sample at the zero-crossing of the VCO waveform and to keep the loop in lock, the sampling event needs to be delayed by the same $0.25 \cdot T_{VCO}$. In the third cycle, the timing error increases by an additional $0.25 \cdot T_{VCO}$, up to $0.5 \cdot T_{VCO}$ and delay of the sampling event is readjusted accordingly. In the fourth cycle, the delay is $0.75 \cdot T_{VCO}$. Finally, in the fifth cycle, sampling should happen at $1 \cdot T_{VCO}$ after the reference edge. However, simply skipping a VCO cycle yields the same effect, since the sampler does not discriminate between multiples of the VCO periods. In other words, on the fifth cycle, sampling is aligned with the integer- N time again ($0 \cdot T_{VCO}$).¹

Fig. 2.2 Implementation of fractional- N subsampling operation by delaying the sampling reference. The last sampling event aligns exactly with the beginning of the next cycle and therefore has no extra delay



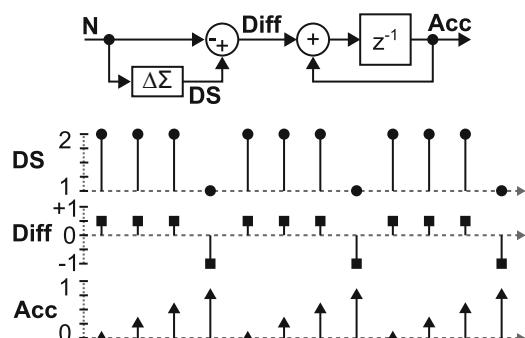
¹Note that the skipping operation is possible thanks to the sinusoidal detection gain of the subsampling PLL, which repeats every T_{VCO} .

Importantly, since the desired fractional- N frequency of the PLL and the reference frequency are known, it is always possible to calculate the position (ideal, targeted, in the absence of noise) of any of the following zero-crossings with absolute precision. This means that if an ideal delay generator were implemented, the PLL would be completely spur-less, unlike the traditional analog $\Delta\Sigma$ PLL. Additionally, the tuning range of the delay generator only needs to cover a single VCO period, since the calculations “wrap around” as in the aforementioned example.

2.2.3 Digital Modulator for the Fractional-N Subsampling PLL

The delay that needs to be inserted into the reference path can be calculated precisely based only on the multiplication factor N and the reference period T_{ref} . The digital computation of the necessary phase adjustment is depicted in Fig. 2.3. The difference of the targeted multiplication factor and its integer-quantized value is computed initially. The subtraction result represents the timing error that the phase detection is going to make in the following sampling event, scaled to the number of VCO periods. A first-order $\Delta\Sigma$ modulator is used to generate the integer quantization of N . The quantization error (Diff in Fig. 2.3) is in this way a zero-mean stream that is easily accumulated. The desired “phase wrapping” behavior is achieved without any additional circuitry. In the second operation point, the quantization error is accumulated, just as the PLL accumulates the phase difference between the VCO and the reference. At this point, it is possible to tell with absolute precision what the necessary delay will be in any of the following cycles. By observing Fig. 2.3, it is noticeable that the accumulated error is *reset* every time the $\Delta\Sigma$ modulator overflows towards the neighboring integer. Note that instead of calculating the *Diff* signal based on the $\Delta\Sigma$ modulator, an overflowing accumulator could just intake the fractional residue as an input. This however reduces the randomization capabilities in the DTC as explained later.

Fig. 2.3 Digital computation flow of the DTC modulator for $N = 1.75$



2.3 Implementation Limitations and Their Mitigation

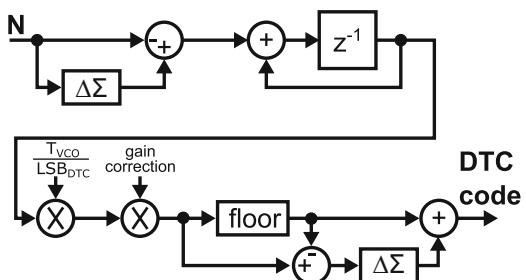
If a fractional- N subsampling PLL, as described in the previous section, were implemented with an ideal DTC, it would have the same performance as an integer- N subsampling PLL. This lies in stark contrast to the case of a traditional mixed-signal PLL, where there is an unavoidable penalty associated with the divider modulation noise (that needs to be filtered). Any practical implementation of the fractional- N subsampling PLL system will, however, be limited in a number of ways. The biggest contributor to these limitations is the DTC. We deal with the DTC implementation challenges one by one in the following sections, proposing adequate solutions.

2.3.1 DTC Quantization

A DTC, as any data converter, has a finite resolution. To scale the output of the accumulated phase error to a digital tuning code, the output of the accumulator in Fig. 2.4 needs to be multiplied by a factor $\frac{T_{\text{ref}}}{(\text{LSB}_{\text{DTC}}, N_{\text{frac}})}$. Even in a noiseless system, the sampling moments will occur with accuracy limited by the LSB of the DTC and the resulting error current will be fed into the LPF, thereby modulating the VCO and creating spurs.

One solution to the problem of limited DTC resolution is obvious: the quantization noise resulting from the DTC's limited resolution should be well below other noise sources. Additionally, adjusting the computed digital word to the available LSB steps is a standard modulation problem, where $\Delta\Sigma$ modulators are often used. As such, the purpose of the second $\Delta\Sigma$ modulator (Fig. 2.4) in this context is to shape the quantization noise beyond the PLL bandwidth. Thanks to the fact that the $\Delta\Sigma$ stream is perfectly accurate on average, the average PLL frequency is also accurate, with no visible modulation. Here, we propose to use an all-pass $\Delta\Sigma$ modulator [Schreier05], which shapes the quantization noise without affecting the DTC modulation signal.

Fig. 2.4 Digital DTC modulator including quantization, gain correction, and quantization noise shaping



The resulting DTC quantization noise power spectral density can be expressed as single sideband phase noise by:

$$\mathcal{L}(f) = 10 \log_{10} \left\{ \left(\frac{2\pi}{\sqrt{12} T_{\text{VCO}}} \frac{\text{LSB}}{T_{\text{VCO}}} \right)^2 \frac{1}{F_{\text{REF}}} [2 \sin(\pi f T_{\text{REF}})]^2 \right\} \left[\frac{\text{dBc}}{\text{Hz}} \right], \quad (2.1)$$

at in-band PLL output, where LSB represents the DTC LSB delay, T_{VCO} is the VCO period, F_{REF} and T_{REF} are the crystal reference frequency and period, respectively, and f is the frequency offset. The term $[2 \sin(\pi f T_{\text{REF}})]^2$ comes from first-order $\Delta\Sigma$ shaping [Schreier05], at reference rate. If the DTC quantization error is set to 0.5 ps, then the in-band (500 kHz) phase noise at 3.6 GHz output (40 MHz reference) appears at -148 dBc Hz^{-1} . This value suggests that quantization noise can easily be kept under the thermal noise floor in a DTC-based subsampling PLL.

Equation 2.1 assumes a uniform Gaussian quantization noise distribution. This assumption does not strictly hold for a periodic saw-tooth input quantized by a first-order $\Delta\Sigma$ and results in additional spurs. This is most pronounced for fractional residue offsets that are in-band near the loop cut-off, where the DTC sees a fast varying input without the benefit of PLL filtering. The quantization error of such a signal has a significant amount of energy stored within spurious tones, typically at harmonics of the fractional residue $.f$. Such tones, induced by “colored,” i.e., repetitive DTC quantization errors, can appear as fundamental (fractional) spurs in the output spectrum.

Another modification to the basic system that helps to mitigate the problem of limited DTC resolution is to use a MASH modulator [Schreier05] in the beginning of the computation path (Fig. 2.4). A MASH modulator provides better randomization of the generated code, which helps in reducing spurious content. Compared to a first-order $\Delta\Sigma$, the generated codes have a larger range, which results in a larger delay range of the DTC.² Looking at the randomization in time domain, it is possible to note that by generating delays larger than one T_{VCO} and sampling one of the neighboring VCO zero-crossing to compensate for the same targeted fractional residue, the sampling data becomes decolorized. This happens because the VCO accumulates slightly different fractional residue at the new moment of sampling, which means that the instantaneous quantization error produced by the DTC becomes different, too. Another way of looking at this effect can be the following: randomizing DTC codes provides an effect similar to dynamic element matching, since, e.g., four DTC codes are used in MASH 1-1-1 mode to generate the same effective sampling phase, their average timing is effective and the apparent DTC (quantization) errors are randomized.

²A first-order $\Delta\Sigma$ generates modulation of only 1. This means that the DTC operates with delays of up to a single VCO period. A popular MASH 1-1-1 modulator has an output range of 7, which is reduced after some filtering in the phase accumulator.

2.3.2 A DTC Versus a TDC in Fractional Frequency Synthesis

A DTC exploits *time-domain* signal processing rather than voltage-domain signal processing which is a trend often seen in recent digital PLLs [Staszewski04] that exploit time-to-digital conversion. Introduction of the digital PLLs did open up new opportunities in the area of fractional frequency synthesis. A TDC is used in these systems to measure the existing phase error and the error is digitally (area efficient) filtered. A DCO is used instead of a VCO for phase-error control. Notably, the quantization error of the TDC used for phase-error detection limits the spectral purity of such PLLs [Borremans10, Temporiti10]. The minimum time-quantization step of the TDC is typically related to the minimum gate delay defined by the physical technology used. Even in advanced 28-nm CMOS it is approximately 10 ps, which is enough to limit the PLL phase noise performance.

A DTC, on the other hand, can easily operate with a resolution on the order of 100 fs [Ru15, Markulic14] as presented in Sect. 2.4.2, and as suggested in the section above, which is fundamentally superior to a TDC in the same technology. In general, the phase modulating capabilities of a DTC can be used for digital PLLs, too, in the feedback path of the PLL [Marzin12, Tasca11], or in the reference path [Pavlovic11, Marucci14]. In either case, the DTC is used to cancel the phase error induced by the fractional- N residue. It enforces a *near-to-zero* phase-error regime of the PD, which relaxes the PD design in terms of its range. Recent designs that use a DTC–TDC combination for phase-error comparison exploit this DTC favoring trade-off [Pavlovic11, Chillara14, Zhuang12]. An extreme case of this are the 1-bit, Bang-Bang TDC PLLs [Marzin14, Tasca11], and MDLLs [Marucci14, Marucci15].

2.3.3 DTC Offset and Gain Error

If the DTC is placed in the path of the reference, any fixed delay (offset) it introduces will propagate towards the output of the PLL. However, this offset is rarely an issue and can be made small by proper design of the DTC.

DTC gain can be defined as the amount of delay in seconds per LSB code. Because the DTC is analog in nature and susceptible to PVT variations, the absolute gain will be unknown and varying with time and temperature. Gain error in the delay steps will introduce incorrect (not near-zero) sampling and hence noise and spurs in the spectrum of the PLL. It is critical to enable automatic background calibration, which will track the gain variations and compensate in either analog or digital domain.

An automatic DTC gain calibration (Fig. 2.5) can be designed similarly to the popular least-mean-square (LMS)-based mechanisms used in digital PLLs [Tasca11, Levantino13]. Simply stated, the sign of the sampled voltage error needs to be extracted and correlated with the change in direction of the DTC word. An intuitive explanation of the process can be given by considering a situation where

Fig. 2.5 Background calibration method for correcting DTC gain error

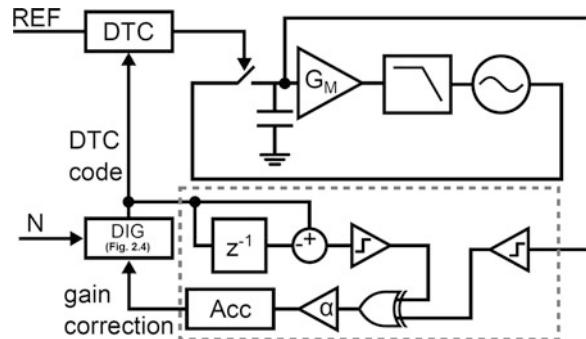
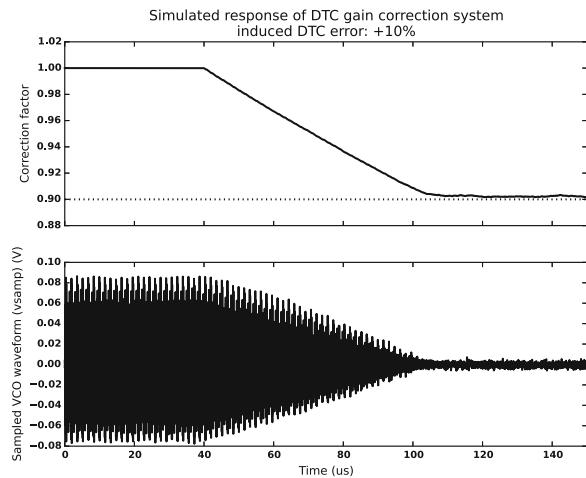


Fig. 2.6 Simulation results of gain correction mechanism when a 10% error is applied to the DTC



the modulator “tells” the DTC to sample later, but due to a DTC gain error the samples are consecutively produced “early”—from this process, it is possible to detect that the DTC gain is too low. After accumulation, the correction word can be applied as a scaling factor to the computation path of Fig. 2.4. After the correction loop converges, there is no penalty on phase noise. Figure 2.6 shows a simulation result where a 10% gain error was applied to the DTC. This error introduces a large ripple in the sampled voltage, which in turn results in large spurs at the output of the PLL. After the DTC gain is corrected, the sampled voltage converges back to zero.

2.3.4 DTC Nonlinearity

Just as in any data converter, the DTC suffers from nonlinearity. This nonlinearity will naturally increase noise and create spurious content at the output of the PLL. For example, simulations show that a 0.5-ps/LSB DTC with a systematic INL of ± 1 LSB results with in-band spurs of approximately -47 dBc, when a multiplication

factor of 250 is used from a 40-MHz reference. This value is purely indicative, since the spur level depends also on the particular nonlinearity shape (i.e., order), delta-sigma modulation order, range of the DTC used, etc. Many techniques for improving linearity which are present for DACs are also applicable in the DTC design. For example, careful layout of the tuning element is of highest priority. Advanced nanometer-scale technologies offer a significant advantage in this regard, thanks to the ever-improving lithography resolution. Matching improves with technology for the same area of a capacitor array. Dynamic element matching (DEM) can be also used to improve the linearity of the array [Chang14]. In addition, it is worth noticing that the third-order MASH modulator, generating codes for the DTC, effectively introduces averaging to the DTC nonlinearity, since there are multiple codes spaced all along the range of the DTC that can be used to sample the same phase offset.

2.3.5 DTC Phase Noise

In this section, we propose a solution to enhance an integer- N subsampling PLL by placing a phase modulator (DTC) in the path of the reference. Unfortunately, the phase noise contribution of the DTC adds directly to the phase noise of the reference. Ultimately, the in-band phase noise of the subsampling PLL is limited by the phase noise of both the reference and the DTC, since both pass the system in the same way. Therefore, great care must be taken to minimize the DTC's contribution to phase noise, otherwise the unique phase noise advantages of the subsampling architecture will be lost. Here, scaling of CMOS technology is again on our side, since transistors are getting faster with every node, reducing jitter and phase noise.

For example, simulations show that with a DTC that operates from a 40-MHz reference with a -155 -dBc/Hz noise floor, induces at 10 GHz³ a PLL in-band noise floor of approximately -107 dBc/Hz.

2.4 Circuit Implementation

The subsampling phase locked loop can only detect phase errors, which makes it susceptible to false locking at any N . Therefore, a frequency-acquisition loop is required in addition to the subsampling loop [Gao09b] (see Fig. 2.7). A simple conventional PLL easily fulfills this requirement. It can be disabled once frequency has been acquired in order to save power. Another option is to put the loop in a *low-power mode* (not available in this design). This increases robustness of the system in nonstandalone environments.

³PLL multiplication factor N is 250 in this example, and phase noise is multiplied by N^2 in transfer to the output (in PLL band).

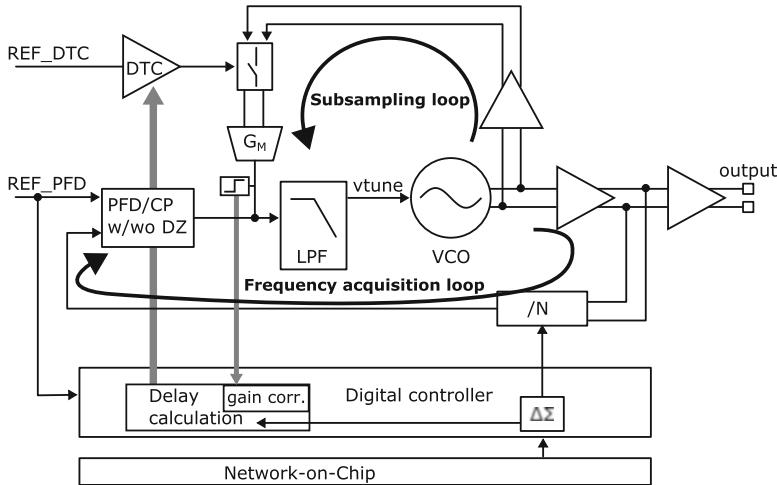


Fig. 2.7 Architecture of the fractional- N subsampling PLL

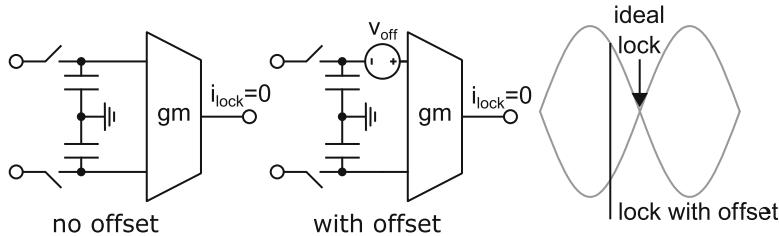


Fig. 2.8 The subsampling PLL always locks into a state that guarantees zero output current, even in the presence of offset and mismatch

Common to both, the frequency and phase acquisition loop are the low-pass filter (LPF) and the VCO. For the purpose of demonstrating the concept of the fractional- N subsampling PLL, we opt for the simplest LPF design—a passive third-order lead-lag filter. Such a simple filter can cause increase in reference spurs and is often avoided in classical charge-pump-based PLLs. Spurious content can increase because the varying level of tuning voltage can introduce mismatches between the currents of the charge pump. In this design, however, any offset in currents of the G_M is compensated by a slight modification of the locking point (Fig. 2.8). A locked condition always means zero output current of the G_M . If changes to the output level cause an input referred offset of the G_M , the PLL will adapt its locking phase to compensate for this offset. Tunable resistance in the LPF has been implemented to be able to change the bandwidth of the PLL.

2.4.1 Implementation of the Subsampling Loop

The subsampling loop consists of a VCO buffer, a sampler, and a G_M . Additionally, the DTC provides the required phase modulation. Figure 2.9 shows the circuits along the subsampling path.

A VCO buffer is inserted in order to reduce the kickback effect from the sampler to the VCO [Gao10] and to interface the signal levels between the blocks. In this test chip, to accommodate for changing phase noise requirements of a software-defined radio, we have implemented a low-noise VCO that can be operated from a variable supply as high as 1.8 V. Therefore, the input buffer needs to convert the level between the high-voltage VCO domain (max. 1.8 V) and the core domain (0.9 V). Additionally, the signal processed by the buffer needs to remain roughly sinusoidal in shape, so that the detection gain (and hence loop gain) can be controlled. The buffer is implemented with a tunable capacitive attenuator and a source follower pair (see Fig. 2.9). The tunable attenuator is built with metal–oxide–metal (MOM) capacitors and provides additional tuning of loop gain. The buffer is also the largest contributor to power consumption in this loop, as it needs to process a GHz-range signal.

The sampler is built around an NMOS switch and a small MOM capacitor. In total, taking into account the input capacitance of the G_M , the sampling capacitance is 20 fF. Thermal $\frac{K_T}{C}$ noise can be neglected because it is already suppressed by the large detection gain. The implemented sampler uses an auxiliary sampler operating

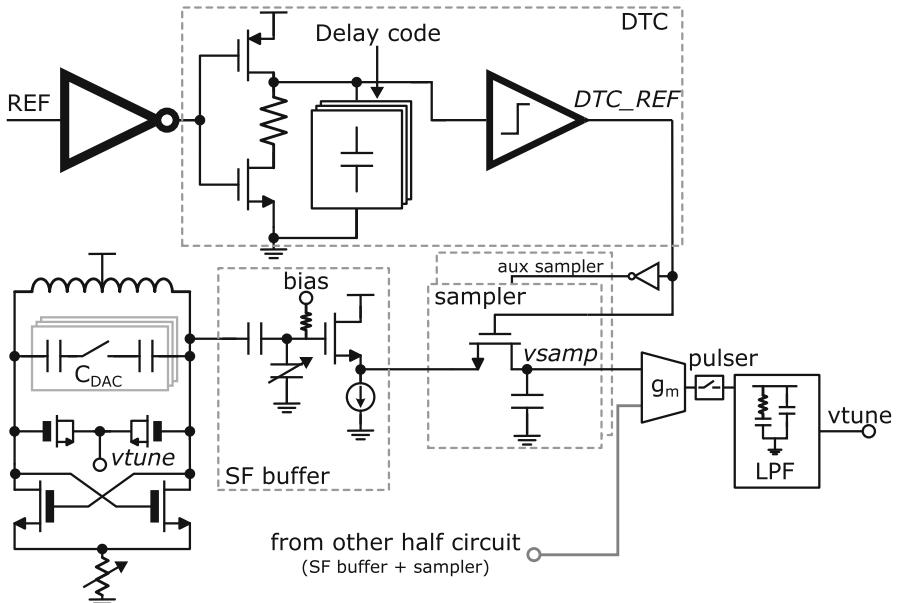


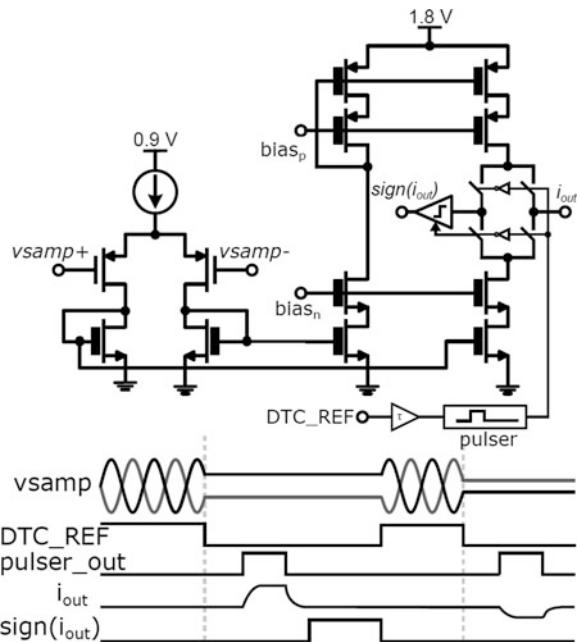
Fig. 2.9 Simplified schematic of the subsampling loop

in inverted phase to the primary sampler in order to reduce load variability of the VCO.

Since the implemented VCO can operate from the IO voltage, the tuning voltage also has a range larger than the core voltage. Therefore, the output stage of the G_M needs to provide translation from the low-voltage domain of the sampler to the high-voltage domain of the LPF and the VCO. This translation is done in current domain between the first and the second stage of the G_M (Fig. 2.10). The first stage is a simple differential pair providing the necessary transconductance, whereas the second stage implements a charge pump-like output. Identically to [Gao09b], the detection gain is so large that duty-cycling is required in the output stage of the G_M . Pulsing is done with a simple digital pulse generator that opens the output switches of the G_M . Importantly, variations in the pulse width merely change the loop gain. A solution to varying loop gain and loop bandwidth would be to implement a loop bandwidth tracking. The solution of [Marzin14] could be attractive here, since it uses the same error signal as the DTC gain compensation.

An important part of the system is the background correction of DTC gain. As said earlier, the error signal from within the PLL is present in the sign of the sampled voltage. However, this is true only if no mismatches are present in the system. If there are any mismatches in the phase detection circuitry (VCO buffers, samplers, and G_M), the PLL will adjust the locking phase (and sampled voltage) so that the output current of the G_M is zeroed (Fig. 2.8). Therefore, the gain correction mechanism requires detection of the sign of output current. This is why the output stage of the G_M is realized using cascodes. The slightest imbalance of current in the

Fig. 2.10 Schematic of the transconductor. The input pair is driven by differential sampled voltage. The output current (i_{out}) is duty-cycled and flows to the loop filter



output branch results in a large swing of voltage at the output node. Using a simple clocked comparator to detect the sign of this swing in relation to $vtune$ voltage is sufficient to obtain information about the sign of the output current.

2.4.2 Implementation of the Digital-to-Time Converter

A real DTC [Markulic14] limits the performance of a fractional- N SSPLL in several ways, as discussed in Sect. 2.2.3. Any deviation of the actual sampling moment w.r.t. the ideal one introduces noise in the system. Obviously, the first limit is the DTC resolution. The linearity of an implemented DTC is equally important as INL and DNL in the DTC transfer function lead to an increased, code-dependent quantization error. The DTC is in fact part of the PLL's phase-error comparison path, where any nonlinearity induces potential noise folding and spurs. Furthermore, the DTC phase noise sets the in-band phase noise of the FNSSPLL: since the DTC is at the input of the system, its phase noise is multiplied by N^2 when transferred to the PLL output, where N is the PLL multiplication number (here: 48 dB as $N = 250$).

Another key performance metric in the design is noise in the supply. Care must be taken that any noise (coming from, e.g., digital switching in the PLL) does not modulate delay generated by DTC. Since the DTC is implemented as a digital delay line, as discussed here (an inverter with a tunable load), any supply noise becomes potentially part of the delay signal. This holds in the other direction as well, i.e., it is equally important that the DTC itself does not disturb the power supply (there should be no code-dependent power consumption) to avoid influence on the sensitive analog subsampling path.

We target in this design a 10-bit DTC, that is based on linear RC settling that accurately determines a 0.55-ps delay step. The designed DTC does not limit the FNSSPLL in-band phase noise performance because of its low phase noise floor and suppression of the 1/f noise component. To increase power supply immunity, high driving strength buffers are put on a regulated supply. Thanks to the high DTC performance at low power, as presented in the following sections, the in-band phase noise of the PLL challenges recent DPLLs that typically invest a lot power in the TDC design for the same goal.

2.4.2.1 Delay Control Block

The architectural overview of the DTC is given in Fig. 2.11 [Markulic14]. The crystal oscillator signal is brought to the REF node, where it is transformed into a fast rise/fall time square wave signal by input buffers. The DTC output is at the node *sampling clock*, where the falling edge controls the sampler's hold state, providing phase-error information used in the loop correction mechanism. Note that the rising DTC output edge need not to be “noise free” since it generates no information in phase comparison.

Fig. 2.11 DTC architecture overview

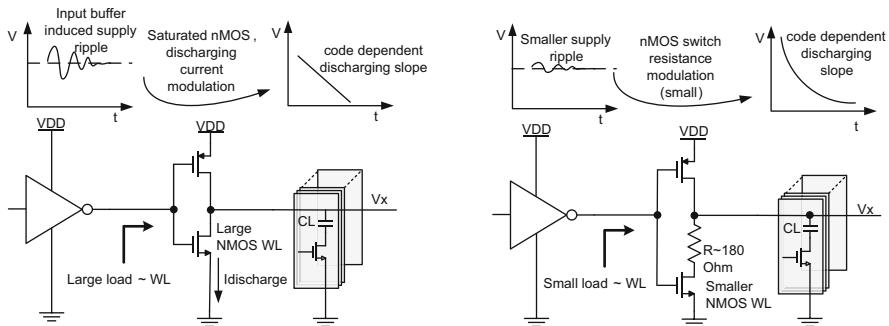
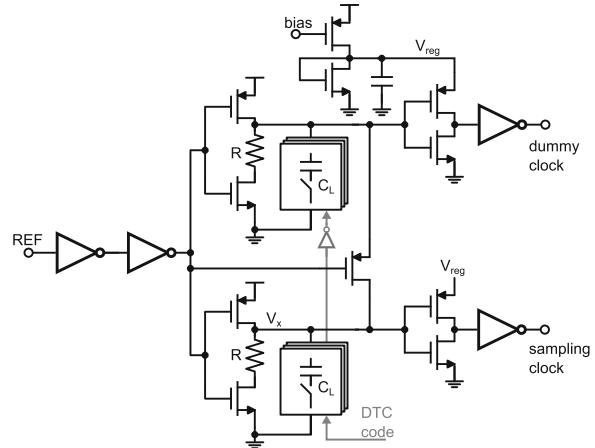


Fig. 2.12 (left) Basic delay block; (right) delay block in the proposed DTC

The input buffers drive the delay control block which could be realized as an inverter with controllable load (see Fig. 2.12—on the left). The output node (V_x) transition speed is dependent on the load capacitance and the inverter strength (ignoring parasitics and assuming a steep input slope). Transitions are ideally noiseless, i.e., always equal for a given code. We concentrate on the relevant high-to-low (HTL) output (V_x) transition. According to the theory presented in [Levantino04], white and flicker phase noise at such a delay element output are proportional to:

$$\mathcal{L}_{\text{white}} \propto 10 \log \left(f_{\text{out}} \frac{kTC_L}{I_{\text{dis}}^2} \right) \propto 10 \log \left(f_{\text{out}} \frac{kT\tau_{\text{delay}}}{I_{\text{dis}}} \right), \quad (2.2)$$

$$\mathcal{L}_{\text{flicker}} \propto 10 \log \left(f_{\text{out}}^2 \frac{C_L^2 2K}{I_{\text{dis}}^2 WLf} \right) \propto 10 \log \left(f_{\text{out}}^2 \tau_{\text{delay}}^2 \frac{2K}{WLf} \right), \quad (2.3)$$

where I_{dis} is the discharging current set by the inverter driving strength, f_{out} is the signal frequency, K is a technology constant, W and L are the transistor width, and length and f is the frequency offset from the reference. τ_{delay} , the targeted delay, is proportional to $\frac{C_L}{I_{\text{dis}}}$.

The white noise contribution at the highest delay can be suppressed below, e.g., -160 dBc/Hz by increasing the discharging current. To keep the same delay step, the load is increased accordingly, and hence the dynamic power consumption $C_L V_{DD}^2 f_{\text{out}}$ increases.

From Eq. (2.3), it is visible that the flicker noise contribution grows quadratically with frequency. Furthermore, the 1/f corner is pushed to high frequencies in newer technologies (such as 28-nm CMOS). The phase noise profile of a delay element as depicted left in Fig. 2.12 is, therefore, limited by the flicker noise (at offsets corresponding to the PLL loop bandwidth). One can try to suppress the flicker contribution by (heavily) over-sizing the inverter both in length and width. This increases load for the input buffers, which then need to have a higher driving strength to preserve fast switching and low phase noise. That leads to a larger supply ripple generated by the buffers and hence larger modulation of the VGS of the discharging nMOS. This in turn leads to linearity degradation.

To reduce this effect, we introduce a resistor above the NMOS which now acts as a switch, as depicted on the right in Fig. 2.12. The exponential discharging is thus determined by the corresponding RC time constant. The delay is, however, a linear function of capacitance. A resistor sets the discharging slope and hence contributes to the output phase noise; however, it generates no 1/f noise. The nMOS transistor can have minimal length and operate merely as a switch, immediately after the relevant transition. The switch resistance is an order of magnitude lower compared to the discharging resistor, hence it has marginal influence, white and 1/f noise wise, on the discharging slope. Since the load is now smaller, the input buffers can be reduced in size, too. Furthermore, the supply ripple only modulates the nMOS switch resistance (small compared to the main resistor) which does not significantly affect the delay. Following the same phase noise analysis for this improved delay element leads to:

$$\mathcal{L}_{\text{white}} \propto 10 \log \left(f_{\text{out}} \frac{kTC_L R^2}{V_{DD}^2} \right) \propto 10 \log \left(f_{\text{out}} \frac{kT \tau_{\text{delay}} R}{V_{DD}^2} \right), \quad (2.4)$$

where R is the resistor value and V_{DD} is the supply voltage of the delay element and f_{out} is the signal frequency. Based on (2.4) and the targeted minimal delay step of 0.55 ps (10-bit DTC), we size the $R = 180 \Omega$ and unit $C = 3 \text{ fF}$ to lower the noise of this stage to -160 dBc/Hz for maximal delay, leaving margin for noise contributions from other blocks.

To further improve the robustness against generation of supply ripples, a complementary delay element similar to the one described is added. This twin circuit is driven by the same input buffer and fed with codes complementary to the main delay element as shown in Fig. 2.11. In this way, we equalize the amount of

charge pulled from the supply in every transition, avoiding any kind of dependence between code and supply ripple. Furthermore, during the pre-charging transition of the delay element the main and the complementary branch are connected via a pMOS switch. This helps equalizing the voltage levels at V_x before the relevant HTL transition.

To suppress mismatch-based errors for the chosen unit size, the capacitor array employs 5+5 bit segmentation. For cancellation of processing gradient-based errors, we implement a common-centroid scheme. For a compact layout, we further exploit the arrays property of a shared top plate. The bottom plate is devised in three levels of metal fingers completely surrounded by the top plate. The inner fingers are interconnected with the minimal size nMOS switch only on the metal level below.

2.4.2.2 Comparator and Output Buffer

After the delay control block, we employ a simple CMOS inverter as a comparator to restore steep slopes ensuring low phase noise. Ideally, the comparator instantaneously toggles after its input signal reaches the threshold level. This toggling moment is unfortunately dependent on the input slope shape [Auvergne00] which degrades the linearity of a high dynamic range DTC as described next.

Namely, as the input slope of the comparator becomes flatter, the comparator experiences short-circuiting operation range before the PMOS starts taking over as the only conducting transistor in the low-to-high output transition (see Fig. 2.13a). Consequentially, the comparator output gets parasitically pre-charged before its input crosses the threshold level. The effect is ever more pronounced for higher DTC input codes, i.e., for slower transitions at the comparator input. Notably, the time needed for the comparator output to reach threshold of the following stage (with respect to the input threshold crossing event) becomes nonlinearly shorter with DTC input code increase. The resulting transfer curve nonlinearity then typically has a shape depicted in Fig. 2.13b.

Another important parasitic effect is that the switching of the comparator produces supply ripples that appear at *delay (code)-dependent instants*. Injection of this code-dependent ripple in the signal path, or parasitic coupling towards the rest of PLL, can degrade the spectral purity of the overall system. Both of the problems, the nonlinear DTC transfer curve and the code-dependent power consumption, are addressed through introduction of the regulated supply in the system.

2.4.2.3 Regulated Supply

To protect the overall supply, the tunable regulated supply shown in Fig. 2.14 is used for the comparator and the first inverter in the chain. The regulated supply consists of a constant current source biasing a diode-connected transistor. A capacitor of 4 pF is used for additional decoupling of the regulated supply node. The regulated supply is set at 720 mV to ensure saturation of the current source while V_{DD} is 0.9 V.

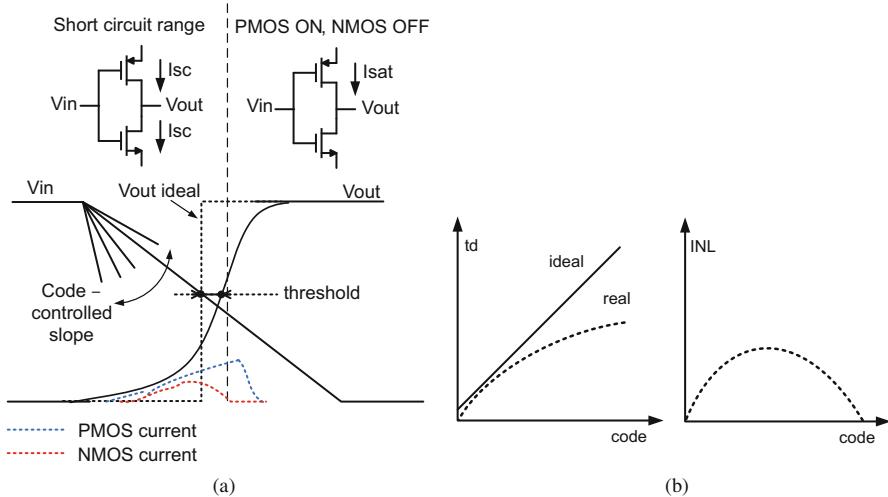


Fig. 2.13 (a) Input slope-dependent comparator response; (b) typical nonlinear DTC delay transfer curve induced by slope-dependent comparator switching

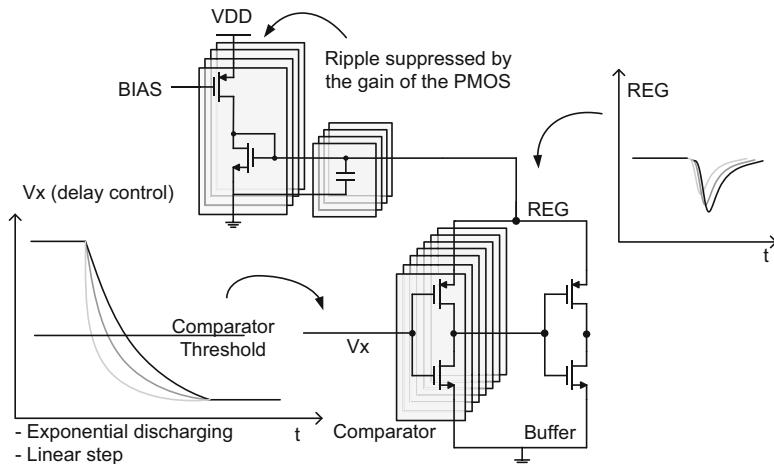


Fig. 2.14 Internal DTC regulated supply for comparator and buffer

At the moment of toggling, charge is instantaneously pulled from the capacitor, and not from V_{DD} . The dip in the regulated supply node is suppressed by the gain of the current source before reaching the top supply. The constant current restores the regulated voltage level. The restoration speed is set by the time constant defined by the diode transistor transconductance and the decoupling capacitance size. The dynamic charge flow is in this way kept within the structure itself.

The regulated supply can partially compensate the DTC's code vs. delay transfer function corrupted by the comparator's input slope dependency. Since the comparator's short-circuiting duration frame becomes larger with higher codes, the instantaneous drop of regulated voltage at the node *REG* increases, too (see Fig. 2.14). The regulated supply restores this drop, as mentioned above, with a time constant defined by the capacitor and the diode-connected NMOS. The comparator exits the short-circuiting region before the regulated supply voltage fully restores. This means that the PMOS operates as a current source (when the NMOS is off) with a code-dependent overdrive voltage. The overdrive is lower for the higher input codes (that induce a larger *REG* dip—as seen in Fig. 2.14). This in turn *slows the output transitions down* for higher input codes—which opposes the nonlinear behavior induced by the parasitic output pre-charging (see the subsection above). By careful regulated supply–comparator codesign, the INL error of the DTC can be kept within few LSBs. The regulated supply is fully tunable and can be switched off by shorting the regulated node to VDD so that its influence can be verified in measurements.

2.4.3 Implementation of the VCO

The VCO used in this PLL is the class-B structure of Fig. 2.15 taken over from [Hershberg14]. A digital varactor utilizing ultra-low V_T thin oxide transistors provides 6-bit digital coarse frequency tuning, and an analog thick-gate oxide varactor provides fine tuning. The cross-coupled $-g_m$ transistors, M_C , see the full VCO swing and are implemented as thick oxide devices. A digitally tunable tail resistor can be used to trade power consumption for phase noise performance.

Fig. 2.15 Class-B VCO schematic and layout floorplan of the NMOS-only digital varactor unit cell of Fig. 2.16b

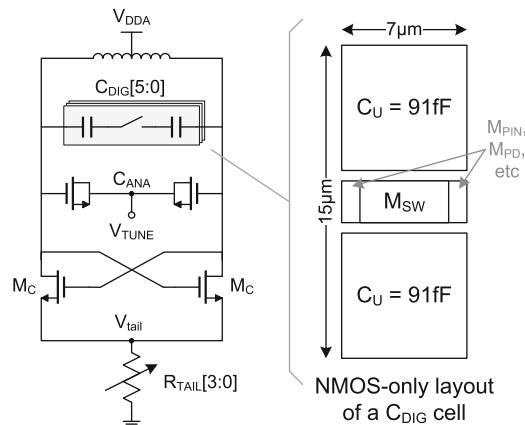
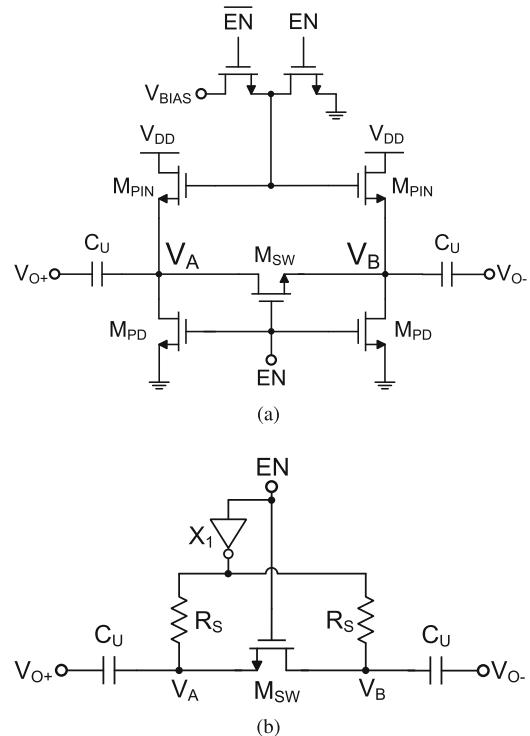


Fig. 2.16 Proposed and conventional switched capacitor structures. The proposed cell (a) is used to implement the digital varactor of the VCO. (a) Proposed switched capacitor cell. (b) Conventional cell of [Sjoland02]



The digital varactor unit cell with “bottom-pinning bias technique” used in this VCO is illustrated in Fig. 2.16a. Compared to the widely used conventional cell of [Sjoland02], shown in Fig. 2.16b, the proposed structure has a number of advantages, particularly in the context of nanoscale CMOS technology[Hershberg14].

In the on-state ($EN = 1$), the proposed switched capacitor cell operates very similarly to the conventional cell: the switch M_{SW} differentially shorts nodes V_A and V_B together and the linear (MOM) capacitors C_U add to the overall tank capacitance of the VCO. However, in the off-state, the M_{PIN} transistors provide a “bottom-pinning” functionality, setting the DC bias levels of the cell such that voltage stress on the devices can be minimized [Hershberg14]. Additionally, this structure naturally produces the highest off-state Q possible, since all leakage is dynamically compensated by the pinning transistors. The proposed cell also benefits from its compact, NMOS-only implementation. As seen in the simplified cell layout of Fig. 2.15, it can be realized as a single composite NMOS block placed between the two unit capacitors. By comparison, the conventional cell uses both PMOS and NMOS transistors and a polysilicon resistor, which cannot be abutted. In this design, there are 15 thermometrically switched capacitor cells, together with a half and a quarter cell.

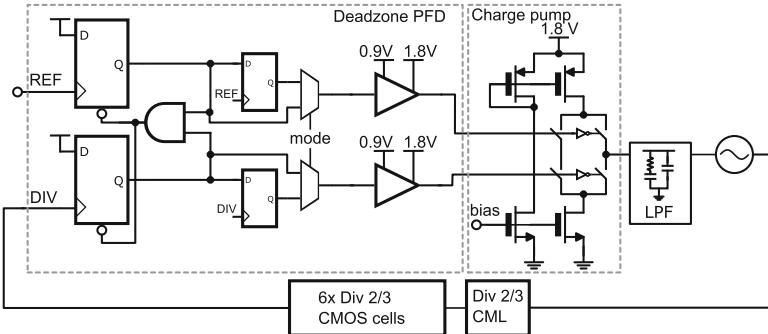


Fig. 2.17 Architecture of the frequency-acquisition loop

2.4.4 Implementation of the Frequency-Acquisition Loop

The frequency-acquisition loop (Fig. 2.17) has been implemented with a chain of divide-by-2/3 circuits [Vaucher00], a traditional 3-state PFD, enhanced with a large deadzone following [Gao09b] and a very simple charge pump. The first stage of the divider is made with CML logic [Szortyka14], since the VCO frequency can reach >12 GHz, but the following stages of the divider are standard CMOS gates. The divider itself is driven by the same MASH $\Delta\Sigma$ modulator used in the digital block of the subsampling path (see Fig. 2.3). The charge pump does not require any mismatch correction, since during frequency acquisition its phase noise performance is irrelevant. Once the frequency acquisition is complete, the loop automatically becomes inactive thanks to the increased dead zone in the PFD and can be completely shut down, saving power.

In general, the loop components for both the phase- and the frequency-acquisition loop can be made very simple and do not require neither good precision, nor good matching, nor low noise. This makes the system suitable for deeply scaled technologies where analog performance is low and also for very-high-frequency applications, where accuracy may be a problem.

2.5 Experimental Results

Prototype chip of the fractional- N subsampling PLL was manufactured in 1P9M 28-nm bulk digital CMOS technology and occupies an area of 1 mm^2 (Fig. 2.18). The active area of the PLL is obviously smaller, dominated by the low-noise VCO which occupies $500 \times 250\text{ }\mu\text{m}$. The chip is powered by 0.9 V and 1.8 V supplies. The 1.8 V supply is used for the IO interface, the charge pump, and the G_M stage. The VCO is designed to work with a low-dropout regulator (LDO) operating at 1.8 V. This LDO, however, is not present on chip and for the results shown below the VCO runs from an unregulated 0.9-V supply. Power consumption (excluding the 50-ohm output drivers and powering down the PFD-based loop) is 13 mW, where

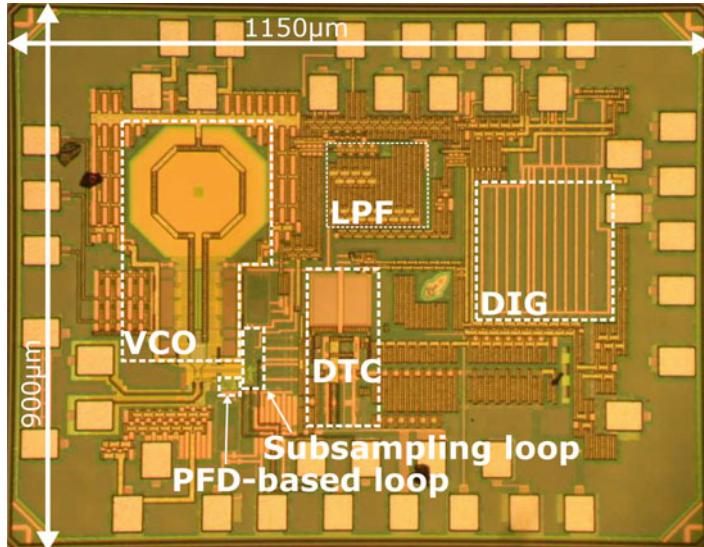


Fig. 2.18 Chip microphotograph

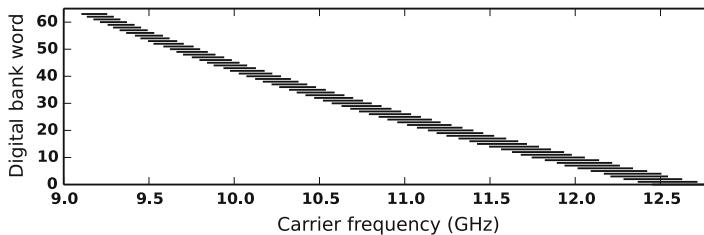


Fig. 2.19 Measured VCO tuning range. Analog tuning (0–1.8 V) is used between digital words

the DTC and G_M consume 0.5 mW and 0.6 mW, respectively, the VCO 8 mW, the source-follower VCO buffer 1 mW, and the digital controller 2.5 mW. The digital controller was neither optimized for power nor for area and includes additional testing circuitry that cannot be clock-gated.⁴ The VCO frequency tuning spans from 9.2 GHz to 12.7 GHz [Hershberg14] with sensitivity to analog voltage that reaches 200 MHz/V around 10 GHz (Fig. 2.19). The out-of-band phase noise can be lowered by 2 dB at the cost of additional 10 mW [Hershberg14] if the VCO is running at 1.4 V (Fig. 2.20).

Oscilloscope measurements of the DTC show INL and DNL of less than 1.5 LSB and 0.8 LSB, respectively (Fig. 2.21). The nominal time resolution is 550 fs, which was confirmed via output of the DTC gain estimation algorithm.

⁴For instance, the digital controller features a full lookup table of the DTC, which is built with 10k flip-flops. The LUT was programmed with a perfectly linear mapping. It was, therefore, not necessary.

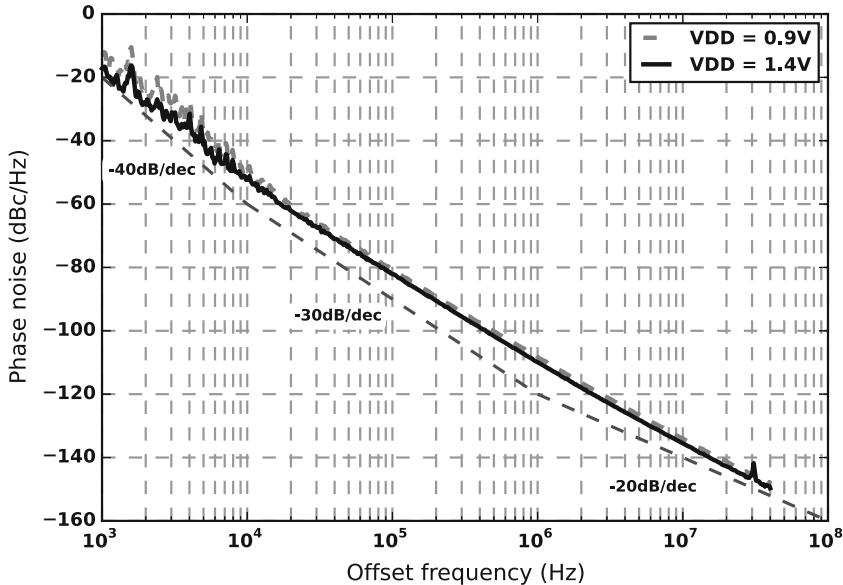


Fig. 2.20 Measured VCO free-running phase noise for low-power ($V_{DD} = 0.9$ V) and high-power ($V_{DD} = 1.4$ V) mode

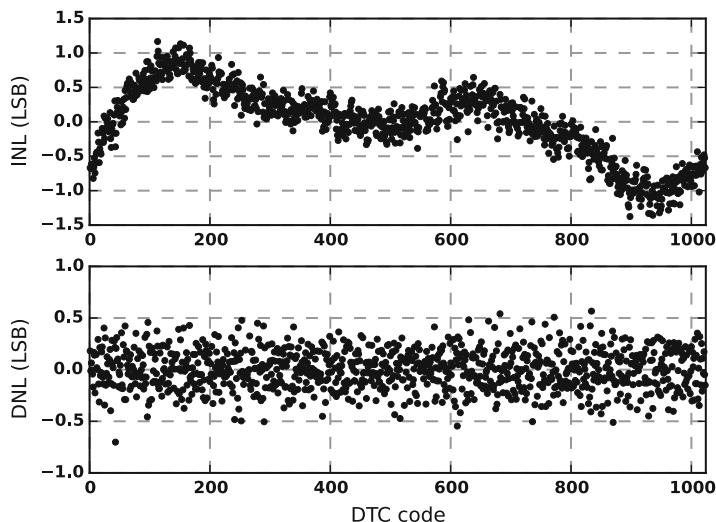


Fig. 2.21 Measured INL and DNL characteristics of the DTC

2.5.1 Measured Phase Noise Performance

Phase noise was measured using an Agilent E5052B signal analyzer with an external 7-GHz downconverter. A sample phase noise result around a carrier frequency of 10 GHz showing the fractional- N spectrum with the worst-case spur (880 kHz) is shown in Fig. 2.22. For comparison, the integer- N phase noise is visible as a memory trace, showing little degradation in the fractional- N mode. The in-band (200 kHz) phase noise reaches -104 dBc/Hz in the fractional- N mode. If bandwidth of the PLL is extended beyond optimum of RMS jitter, the in-band phase noise level can drop to -108 dBc/Hz. The noise at low offset frequencies is a $\frac{1}{f}$ noise of the reference chain and the DTC. Also, the regulated supply of the DTC is adding some filtered $\frac{1}{f}$ noise.

The integrated phase noise in fractional- N mode spans between -40 dBc and -38 dBc depending on the fractional number. In integer- N mode, it reaches -41 dBc. Phase noise integration was done from 10 kHz to 60 MHz and includes all spurs. No compensation or correction was applied to the system, apart from the online DTC gain correction. The PLL is working in a MASH 1-1-1 mode. Jitter was extracted from the integrated phase noise and is shown versus fractional codes in Fig. 2.23.

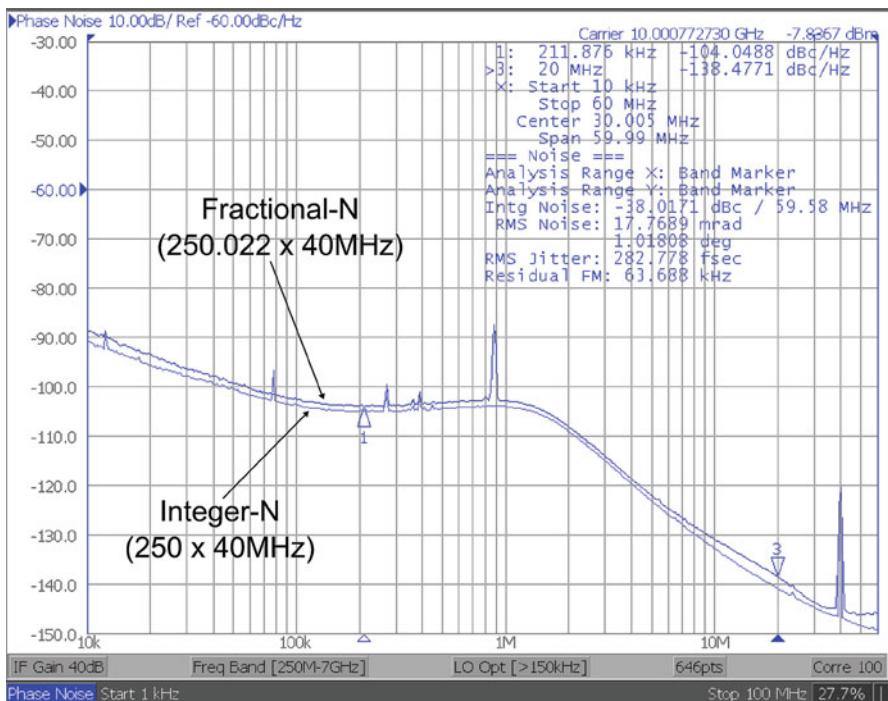


Fig. 2.22 Measured phase noise for a worst-case fractional- N scenario. For reference, the integer- N phase noise trace is shown as well

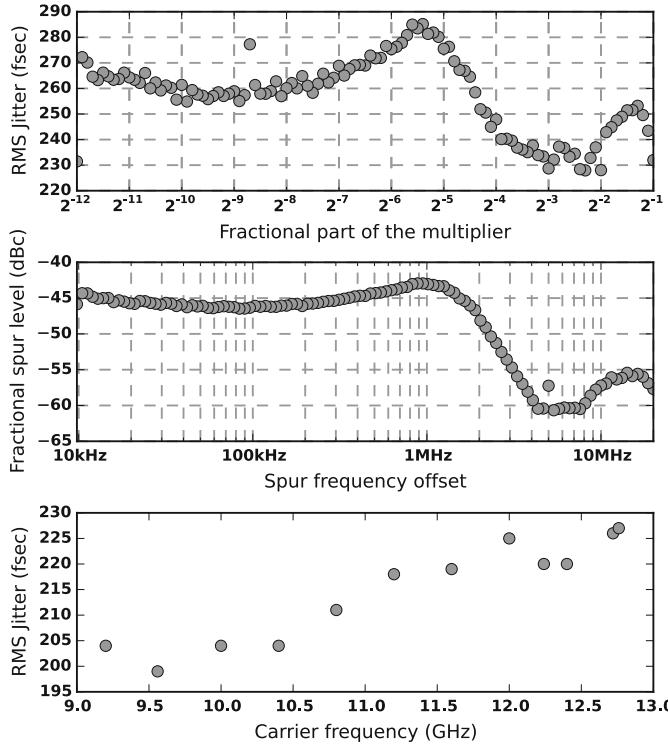


Fig. 2.23 Measured RMS jitter across fractional codes (integer part of $N = 250$) and integer- N jitter with respect to VCO tuning range

With out-of-band fractional multiplication, the RMS jitter reaches 230 fs. When working in integer- N mode, the synthesizer achieves RMS jitter of only 204 fs. Integer- N RMS jitter is reported against VCO tuning range in Fig. 2.23. Settling time (with a fractional step of 20 MHz) is below 2 μ s. Locking time from a free-running VCO (with preselected band) is below 12 μ s. No automatic band selection mechanism is present on chip. Once in subsampling phase lock, the PLL keeps its stable operating point (hold in range) with slow variations of the output frequency in the approximate range of ± 40 MHz (determined with the KVCO). Since the locking process of this PLL had not been investigated, we lack additional investigations on the lock-in and pull-in range. The PLL successfully locks for any starting point with < 40 -MHz difference to the targeted frequency.

Spurious response was measured using a Rohde & Schwartz FSQ26 spectrum analyzer and is shown in Fig. 2.24. The worst-case in-band fractional spur is -43 dBc. The reference spur is -60 dBc. The low measured integer- N spur proves that the relatively large sensitivity of the varactor (200 MHz/V) did not lead to any unexpected spectrum degradation.

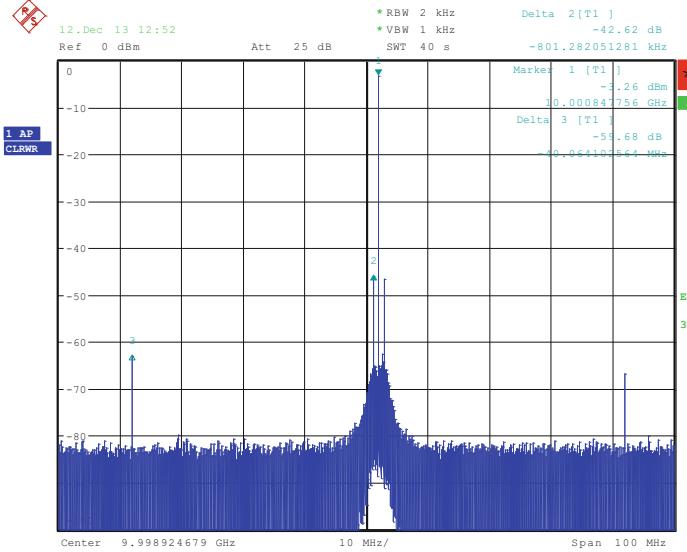


Fig. 2.24 Measured output spectrum of the PLL showing the worst-case fractional spur and the reference spur

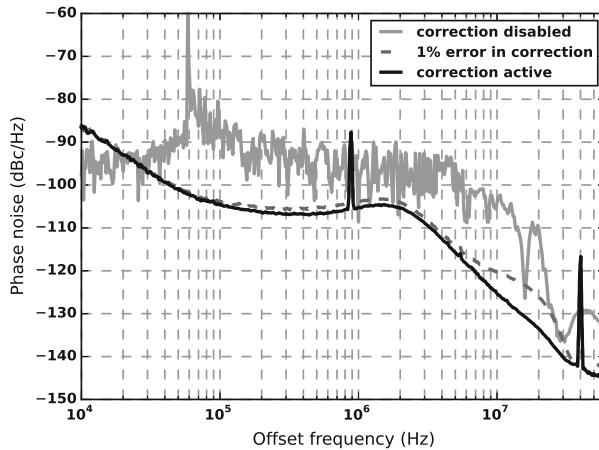


Fig. 2.25 Measured effect of DTC gain mismatch. 1% error in gain was intentionally applied

Figure 2.25 shows the effect of enabling the DTC gain correction mechanism. Without correction, phase noise is simply not acceptable. If 1% error is intentionally introduced to the optimal DTC gain, large spurs can be observed. Finally, optimal performance is obtained if the background calibration is tracking the DTC gain.

Figure 2.26 shows the effect of the MASH modulator. Higher DSM order is preferable; however, it increases the required range of the DTC.

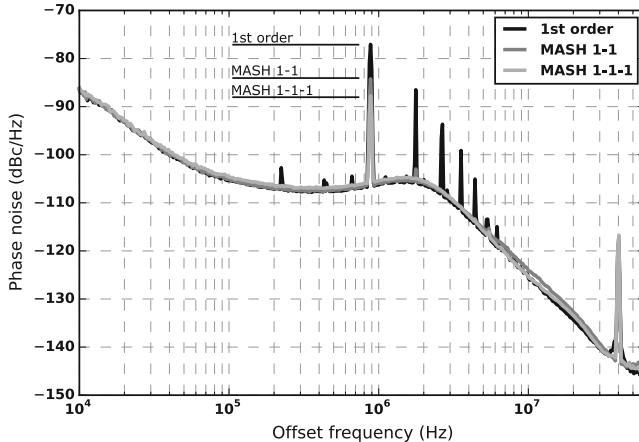


Fig. 2.26 Measured phase noise as a function of the $\Delta\Sigma$ modulator order. The higher the order, the lower the spurs

2.5.2 Remaining Fractional Spur

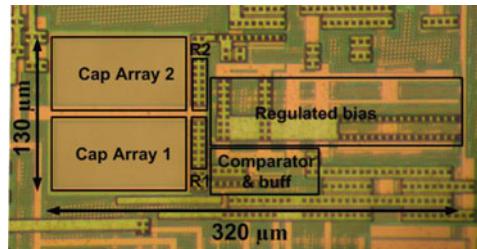
The integrated jitter is degraded when a fractional spur appears in-band. This spur is directly proportional to the multiplication factor: $f_{\text{spur offset}} = f_{\text{ref}} \cdot (N - \text{floor}(N))$. We believe that one of the reasons for a large in-band spur is the nonregulated supply of the VCO. The VCO is designed to be operated together with an LDO, which is not present on this chip. The sensitivity to the supply of the VCO is in fact higher than sensitivity to the tuning voltage. In high-power mode of the oscillator (1.4 V), the supply sensitivity is reduced and the fractional spur drops at approximately -49 dB.

The second contributor to the fractional spur is the DTC nonlinearity as depicted in Fig. 2.21 and discussed in Sect. 2.5.3. Similarly to a nonlinear charge pump in classical analog PLLs, the DTC represents bottleneck for linearity in the phase-error detection path of a subsampling PLL. Introducing the nonlinear curve from Fig. 2.21 in the modeled subsampling PLL environment results in simulated spur levels that fit to the measurements.

2.5.3 DTC-Related Measurements

The effective DTC area is 0.04 mm^2 (Fig. 2.27) and it operates at 0.9 V supply, consuming up to 0.580 mA . There is a dedicated path to measure the DTC as a standalone circuit. Measurements of the DTC transfer curve proved to be rather challenging since the DTC resolution ($\text{LSB} = 550 \text{ fs}$) is below the resolution of available oscilloscopes. In order to correlate measurements of the same static

Fig. 2.27 Micrograph of the DTC



transfer curve, a high number of averages had to be taken. A transfer curve has been extracted from 70 different measurements by Agilent oscilloscope 86100D/86112A. The averaged measured INL and DNL are represented in Fig. 2.21. Measurements of the DTC transfer curve have been made using high-speed oscilloscope and a large number of averages. The minimal propagation time of the signal through the DTC is 560 ps, and the maximal delay is 1.11 ns. The DTC delay range covers 5 VCO periods at 9.2 GHz.

DTC performance characterization in terms of phase noise, supply rejection, and linearity is further investigated throughout measurements of the PLL spectrum profile. VCO operates in a low-noise mode on a separate 1.4-V supply. This allows investigating the limitations of the system mainly set by the DTC. As indicated above, the in-band phase noise of a fractional-N subsampling PLL is dominated by the DTC noise. It is multiplied by N^2 when transferred to the PLL output. Measured PLL output phase noise profiles for different settings of the DTC are depicted in Fig. 2.28. We give direct measurements of the DTC phase noise, too. These show a 1/f component that overdominates the white noise. This is because only the falling edge was optimized for phase noise, and hence, the PLL in-band is not limited by the DTC 1/f corner. The white noise floor of the DTC is at -156 dBc/Hz for minimal, and -154 dBc/Hz for maximal delay, directly setting the PLL in-band performance. While the PLL is driven in integer- N mode, the DTC is transparent, i.e., driven by a constant code. The in-band phase noise is, of course, best for minimal delay. The in-band phase noise level at 300 kHz in integer- N operation (when N is 250 and the PLL output at 10 GHz) is -106 dBc/Hz. For a forced maximal DTC delay, the PLL in-band phase noise increases with 2 dB.

When operating in fractional- N mode, the PLL in-band phase noise does not significantly degrade and is around -105 dBc/Hz. However, some spurious content appears (mainly limited by the DTC now, since the VCO operates on the higher supply). The overall system performance in terms of spurious content and total RMS jitter (10k–30M) is summarized in Fig. 2.29. The figure shows how the system behaves for every synthesized fractional frequency between 10 and 10.0035 GHz with a step of 15.26 kHz. Above that, fractional spurs are outside of the PLL bandwidth and are negligible. The plots show the worst fractional spur level,

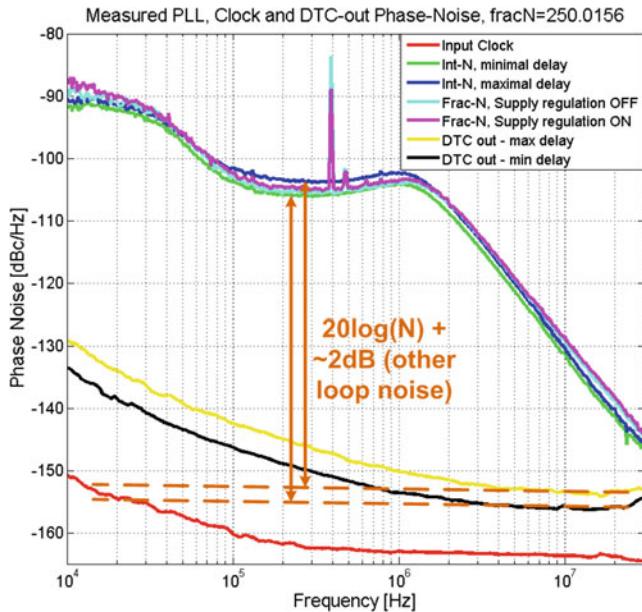


Fig. 2.28 PLL output at 10 GHz for different DTC settings

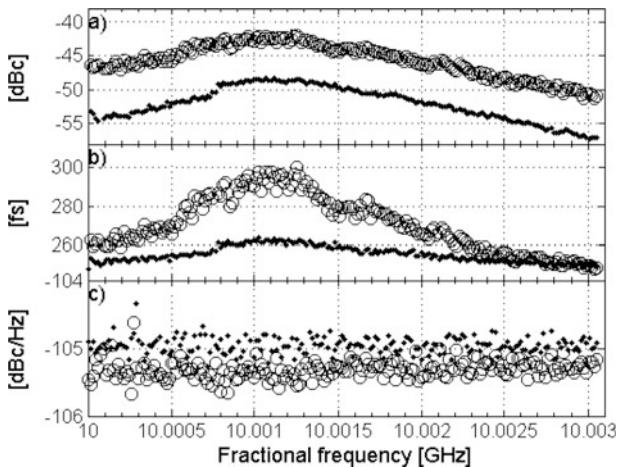


Fig. 2.29 Measured PLL performance with DTC regulation ON (black dot) and OFF (open circle). (a) Worst fractional spur. (b) RMS jitter. (c) In-band phase noise

integrated RMS jitter and in-band phase noise level, respectively, with the regulated supply enabled and disabled. With supply regulation disabled, the worst fractional spur is at -43 dBc, the worst integrated RMS jitter is at 303 fs, and the in-band phase noise is -105.5 dBc. Enabling supply regulation gives -49 -dBc worst fractional

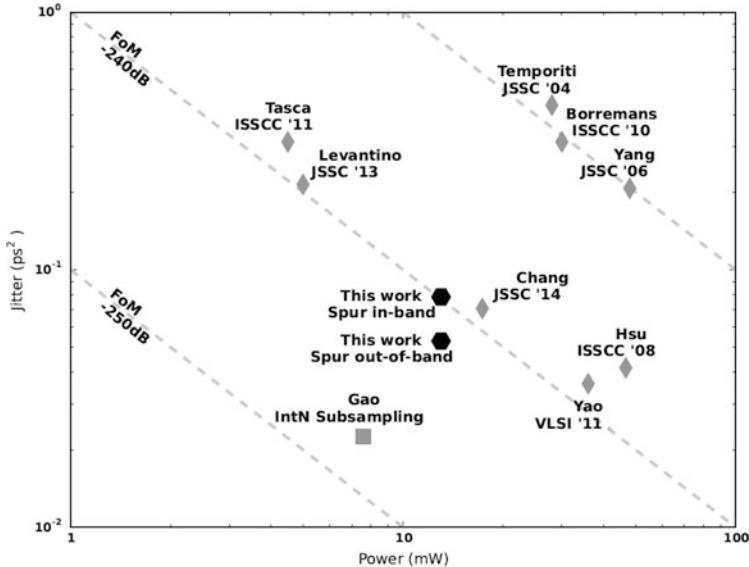


Fig. 2.30 Figure-of-merit comparison of recent fractional- N synthesizers

spur, 270-fs worst integrated jitter, and -105 -dBc/Hz in-band phase noise level. Clearly, the measurements show that the system benefits from the supply regulation, and that a supply ripple influences the overall system performance.

2.5.4 Performance Summary and Comparison to the State of the Art

Generally applied figure-of-merit [Gao09a] of PLL synthesizers is defined as

$$\text{FoM} = 10 \cdot \log \left[\left(\frac{\text{RMS jitter}}{1 \text{ s}} \right)^2 \cdot \left(\frac{\text{Power}}{1 \text{ mW}} \right) \right] \quad (2.5)$$

Table 2.1 and Fig. 2.30 show the summary of performance and FoM comparison for a few recent low-jitter fractional- N synthesizers. Figure-of-merit of the presented fractional- N subsampling PLL reaches -241.5 dB with out-of-band worst-case spur or -240 dB with the spur in-band. Excellent FoM is achieved thanks to the very low phase noise, but also thanks to the simplicity of the subsampling loop which can be designed low power. Compared to [Chang14], which is also a DTC-enhanced subsampling PLL, in-band phase noise (after scaling to 10 GHz) is close to 6 dB lower, which may be a benefit of working with a 28-nm technology. On the other hand, nanometer-scale technologies suffer from large $\frac{1}{f}$ noise, which in our

Table 2.1 Performance summary and comparison of the FNSSPLL to other low-jitter fractional- N CMOS PLLs

	This work	[Chang ^[14]]	[Yang ^[6]]	[Levantino ^[13]]	[Tasca ^[1]]	[Yao ^[1]]
Type of PLL	Analog	Analog	Analog	Analog	Digital	Digital
Technology	28 nm	180 nm	180 nm	65 nm	65 nm	55 nm
Tuning range (GHz)	9.2–12.7	2.2–2.4	2.5–3.2	3.0–4.0	2.9–4.0	5.9–8.0
Reference freq (MHz)	40	48	33	40	40	40
Bandwidth (MHz)	1.8	0.5	0.2	0.5	0.3	0.5
In-band phase noise (dBc/Hz) ^a	-104	-99.2	-89	-96.2	92.5	-103
Phase noise at 20-MHz offset (dBc/Hz) ^b	-138	-128	-139	-123.2	-128	-144
RMS jitter (fs) ^c	230–280	266–400	455	463	560	190
Integrated phase noise (dBc) ^{a,c}	-39.8 to -38.1	-38.5 to -35	-34.8	-33.8	-32	-41.5
Worst fractional spur (dBc)	-43	-53	-74	-42	-42	-70
Reference spur (dBc)	-60	-55	-78	-71	-72	-94
Power (mW)	13	17.3	48	5	4.5	36
Figure-of-merit (FoM)	-241.5 to -240	-239.1	-230	-239.7	-238.5	-239

^aScaled to 10 GHz by $20 \log\left(\frac{f_c}{10 \text{GHz}}\right)$ ^bScaled to 10 GHz and extrapolated from existing data to 20-MHz offset^cIncluding in-band or out-of-band spurs

case, dominates noise profile of the DTC. Our FoM is only slightly better than [Levantino13], though achieved for almost three times larger N and with a three times larger bandwidth. The design is on-par with the lowest-jitter digital PLL in [Yao11], though consuming only a third of its power and not using a reference doubler.

2.6 Conclusion

In this chapter, we have proposed a methodology of enhancing the low phase noise subsampling PLL to work with fractional- N multiplication factors. This methodology introduces a digital-to-time converter in the path of the reference clock, assisted by a simple digital controller. Open-loop modulation of the DTC is possible thanks to the fact that the quantization error introduced by the integer- N PLL is known a priori. We propose an effective online calibration mechanism to adjust the modulation to the PVT variations of the DTC. Moreover, we propose a number of techniques to improve spurious performance of the system limited by the resolution of the DTC.

A fractional- N subsampling PLL prototype reaches 280 fs of RMS jitter in worst-case fractional spur scenario and 204 fs in integer- N mode while consuming 13 mW. The synthesizer has a tuning range from 9.2 GHz to 12.7 GHz. Compared to the state-of-the-art synthesizers (see Table 2.1), it can be seen that this work falls within lowest phase noise analog fractional- N synthesizer to date. The in-band phase noise level of -104 dBc/Hz challenges state of art of all fractional- N synthesizers.

The following two chapters take this basic fractional- N subsampling PLL and develop it in two directions: (1) We recognize that the loop can still be optimized for performance gap minimization between integer- N and fractional- N operation. As shown in Chap. 3, this can be achieved by digital linearization of the phase comparison path. (2) In the second step (Chap. 4), we perform loop digitization. The FNSSPLL is in its basic form a digitally intensive PLL with an analog loop core (loop filter). With a digital loop, we could still achieve additional area savings; however, this should not come at a compromise in performance.

References

- [Auvergne00] D. Auvergne, J.M. Daga, M. Rezzoug, Signal transition time effect on CMOS delay evaluation. *IEEE Trans. Circuits Syst. I: Fund. Theory Appl.* **47**(9), 1362–1369 (2000)
- [Borremans10] J. Borremans, K. Vengattaramane, V. Giannini, J. Craninckx, A 86 MHz-to-12 GHz digital-intensive phase-modulated fractional- N PLL using a 15 pJ/Shot 5 ps TDC in 40 nm digital CMOS, in *2010 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. (IEEE, San Francisco, 2010), pp. 480–481

- [Chang14] W.-S. Chang, P.-C. Huang, T.-C. Lee, A fractional-N divider-less phase-locked loop with a subsampling phase detector. *IEEE J. Solid State Circuits* **49**(12), 2964–2975 (2014)
- [Chillara14] V.K. Chillara, Y.-H. Liu, B. Wang, A. Ba, M. Vidojkovic, K. Philips, H. de Groot, R.B. Staszewski, 9.8 An 860 μ W 2.1-to-2.7 GHz all-digital PLL-based frequency modulator with a DTC-assisted snapshot TDC for WPAN (Bluetooth Smart and ZigBee) applications, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. (IEEE, San Francisco, 2014), pp. 172–173
- [Gao09a] X. Gao, E. Klumperink, P. Geraedts, B. Nauta, Jitter analysis and a benchmarking figure-of-merit for phase-locked loops. *IEEE Trans. Circuits Syst. Express Briefs* **56**(2), 117–121 (2009)
- [Gao09b] X. Gao, E. Klumperink, M. Bohsali, B. Nauta, A Low noise sub-sampling PLL in which divider noise is eliminated and PD/CP noise is not multiplied by N^2 . *IEEE J. Solid State Circuits* **44**(12), 3253–3263 (2009)
- [Gao10] X. Gao, E. Klumperink, G. Soccia, M. Bohsali, B. Nauta, Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector. *IEEE J. Solid State Circuits* **45**(9), 1809–1821 (2010)
- [Hershberg14] B. Hershberg, K. Raczkowski, K. Vaesen, J. Craninckx, A 9.1–12.7 GHz VCO in 28 nm CMOS with a bottom-pinning bias technique for digital varactor stress reduction, in *ESSCIRC 2014 - 40th European Solid State Circuits Conference (ESSCIRC)*. (IEEE, Venice Lido, 2014), pp. 83–86
- [Levantino04] S. Levantino, L. Romanò, S. Pellerano, C. Samori, A. Lacaita, Phase noise in digital frequency dividers. *IEEE J. Solid State Circuits* **39**(5), 775–784 (2004)
- [Levantino13] S. Levantino, G. Marzin, C. Samori, A. Lacaita, A wideband fractional-N PLL with suppressed charge-pump noise and automatic loop filter calibration. *IEEE J. Solid State Circuits* **48**(10), 2419–2429 (2013) [Online]. <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6576220>
- [Markulic14] N. Markulic, K. Raczkowski, P. Wambacq, J. Craninckx, A 10-bit, 550-fs step Digital-to-Time Converter in 28 nm CMOS, in *ESSCIRC 2014 - 40th European Solid State Circuits Conference (ESSCIRC)*. (IEEE, Venice Lido, 2014), pp. 79–82
- [Marucci14] G. Marucci, A. Fenaroli, G. Marzin, S. Levantino, C. Samori, A.L. Lacaita, 21.1 A 1.7 GHz MDLL-based fractional-N frequency synthesizer with 1.4 ps RMS integrated jitter and 3 mW power using a 1b TDC, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. (IEEE, San Francisco, 2014), pp. 360–361
- [Marucci15] G. Marucci, Techniques for high-efficiency digital frequency synthesis. PhD Thesis, Politecnico di Milano, Italy, 2015
- [Marzin12] G. Marzin, S. Levantino, C. Samori, A.L. Lacaita, “A 20 Mb/s phase modulator based on a 3.6 GHz digital PLL with –36 dB EVM at 5 mW power. *IEEE J. Solid State Circuits* **47**(12), 2974–2988 (2012)
- [Marzin14] G. Marzin, S. Levantino, C. Samori, A.L. Lacaita, 2.9 A background calibration technique to control bandwidth in digital PLLs, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. (IEEE, San Francisco, 2014), pp. 54–55
- [Pavlovic11] N. Pavlovic, J. Bergervoet, A 5.3 GHz digital-to-time-converter-based fractional-N all-digital PLL, in *2011 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. (IEEE, San Francisco, 2011), pp. 54–56
- [Raczkowski15] K. Raczkowski, N. Markulic, B. Hershberg, J. Craninckx, A 9.2–12.7 GHz wideband fractional-N subsampling PLL in 28 nm CMOS with 280 fs RMS jitter. *IEEE J. Solid State Circuits* **50**(5), 1203–1213 (2015)

- [Riley93] T.A. Riley, M.A. Copeland, T.A. Kwasniewski, Delta-sigma modulation in fractional-N frequency synthesis. *IEEE J. Solid State Circuits* **28**(5), 553–559 (1993)
- [Ru15] J.Z. Ru, C. Palattella, P. Geraedts, E. Klumperink, B. Nauta, A high-linearity digital-to-time converter technique: constant-slope charging. *IEEE J. Solid State Circuits* **50**(6), 1412–1423 (2015)
- [Schreier05] R. Schreier, G.C. Temes, *Understanding Delta-Sigma Data Converters*, vol. 74. (IEEE Press, Piscataway, 2005)
- [Sjoland02] H. Sjoland, Improved switched tuning of differential CMOS VCOs. *IEEE Trans. Circuits Syst. II: Analog Digit. Signal Process.* **49**(5), 352–355 (2002)
- [Staszewski04] R.B. Staszewski, K. Muhammad, D. Leipold, C.-M. Hung, Y.-C. Ho, J.L. Wallberg, C. Fernando, K. Maggio, R. Staszewski, T. Jung et al, All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS. *IEEE J. Solid State Circuits* **39**(12), 2278–2291 (2004)
- [Swaminatha07] A. Swaminathan, K. Wang, I. Galton, A wide-bandwidth 2.4 GHz ISM band fractional-N PLL with adaptive phase noise cancellation. *IEEE J. Solid State Circuits* **42**(12), 2639–2650 (2007)
- [Szortyka14] V. Szortyka, Q. Shi, K. Raczkowski, B. Parvais, M. Kuijk, P. Wambacq, 21.4 A 42 mW 230 fs-jitter sub-sampling 60 GHz PLL in 40 nm CMOS, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. (IEEE, San Francisco, 2014), pp. 366–367 [Online]. <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6757472>
- [Tasca11] D. Tasca, M. Zanuso, G. Marzin, S. Levantino, C. Samori, A. Lacaia, A 2.9–4.0-GHz fractional-N digital PLL with bang-bang phase detector and 560-fs RMS integrated jitter at 4.5-mW power. *IEEE J. Solid State Circuits* **46**(12), 2745–2758 (2011)
- [Temporiti10] E. Temporiti, C. Weltin-Wu, D. Baldi, M. Cusmai, F. Svelto, A 3.5 GHz wideband ADPLL with fractional spur suppression through TDC dithering and feedforward compensation. *IEEE J. Solid State Circuits* **45**(12), 2723–2736 (2010)
- [Vaucher00] C. Vaucher, I. Ferencic, M. Locher, S. Sedvallson, U. Voegeli, and Z. Wang, A family of low-power truly modular programmable dividers in standard 0.35- μ m CMOS technology. *IEEE J. Solid State Circuits* **35**(7), 1039–1045 (2000) [Online]. <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=848214>
- [Yang06] Y.-C. Yang, S.-A. Yu, Y.-H. Liu, T. Wang, S.-S. Lu, A Quantization noise suppression technique for $\Delta\Sigma$ fractional-N frequency synthesizers. *IEEE J. Solid State Circuits* **41**(11), 2500–2511 (2006)
- [Yao11] C.-W. Yao, L. Lin, B. Nissim, H. Arora, T. Cho, A low spur fractional-N digital PLL for 802.11 a/b/g/n/ac with 0.19 ps RMS jitter, in *2011 Symposium on VLSI Circuits - Digest of Technical Papers*. (IEEE, Honolulu, 2011), pp. 110–111
- [Zhuang12] J. Zhuang, R.B. Staszewski, A low-power all-digital PLL architecture based on phase prediction, in *2012 19th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*. (IEEE, Seville, 2012), pp. 797–800

Chapter 3

A Background-Calibrated Subsampling PLL for Phase/Frequency Modulation



3.1 Introduction

The PLL presented so far (Chap. 2) describes a divider-less loop, based on a high-gain subsampling phase-error detection core that can enable low phase noise (low-jitter) operation in integer-N and in fractional-N modes. The two modes, however, still operate with a substantial performance gap between them. This chapter proposes ways of closing this gap. The goal is to maintain the same levels of spectral purity as in integer-N mode by ensuring low fractional quantization-error residue and highly linear phase-error detection through error randomization and background calibration.

Furthermore, the chapter presents ways of implementing an FNSSPLL-based phase modulator. Wide-band phase modulation is achieved despite the fact that a PLL operates with limited bandwidth. Thanks to the devised background calibration, the system provides extremely high accuracy of the modulated signal (expressed in terms of error-vector magnitude or EVM), overcoming typical limitations normally imposed by the PLL phase noise, mismatch, and nonlinearity in the digital-to-modulated phase conversion. The described environment can, as such, be used as basis for the phase-modulating path of a polar transmitter that can potentially enable higher-order modulation schemes which maximize the spectral efficiency, achieving high bit/s/Hz values.

The chapter is based on material from [Markulic16a, Markulic16b].

3.1.1 PLL-Based Phase Modulation

A typical phase modulator for a polar TX is based on a fractional-N PLL with a bandwidth carefully selected for optimal filtering of phase noise from the oscillator and reference, including in-loop and quantization noise. The basic concept of $\Delta\Sigma$

division factor dithering for fractional synthesis [Riley93] can easily be extended to achieve phase modulation. Modulation through the single injection point is limited by the PLL low-pass filtering. For wideband modulation, the loop bandwidth can obviously be increased. This however comes at the price of suboptimal phase noise filtering. The problem of limited modulation bandwidth can be to some extend prevailed by accurately matched data predistortion [Pamarti04, Perrott97] (see Fig. 3.1). The digital predistortion filter needs thus to accurately model inverse of the PLL filtering profile. This can be rather challenging in analog environments that are unfortunately rather susceptible to environmental changes, i.e., PVT variations. Moreover, for higher modulation bandwidths (i.e., data rates much larger than the loop bandwidth), strong data predistortion is necessary at high frequencies. This easily induces clipping in the phase-error comparison path (PD) that leads to nonlinear PLL behavior and consequential performance degradation.

A much more convenient wideband modulation technique is two-point injection, depicted in Fig. 3.2, adopted by, e.g., [Dürdödt01], for analog PLLs. The basic principle is to inject the modulation signal simultaneously in front of the VCO (point-two) and in the reference path, through a programmable divider (point-one). The VCO frequency instantaneously changes with respect to the modulation signal. At the same time, the divider changes its division factor to compensate for the frequency shift, matching the divided output to the input phase/frequency. In case of an accurate match between the two injection points, there is an overall all-pass of the data to the PLL output. The phase-error detector and therefore the loop sense no disturbance due to the modulation.

Two-point injection can be analyzed in frequency domain. Injection point-two has a high-pass profile to the PLL output, while point-one implements a low-pass. There is an all-pass transmission of the modulation data to the PLL output if the bandwidths of the two profiles match.

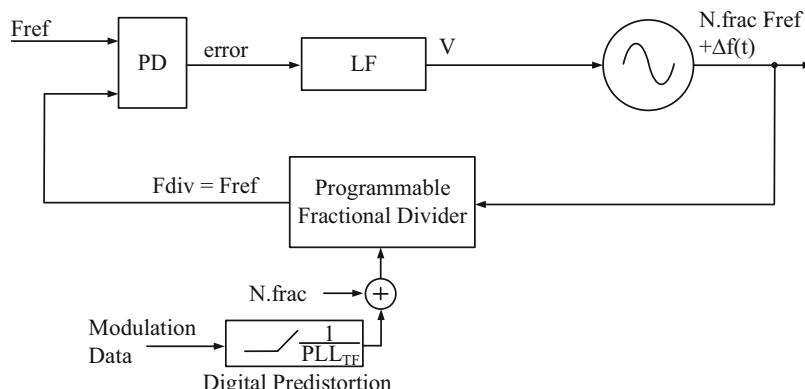


Fig. 3.1 Basic principle of single-point phase/frequency modulation with data predistortion. PLL_{TF} stands for the PLL phase signal transfer function from the input to the output. Note that modulation data propagates through the whole loop

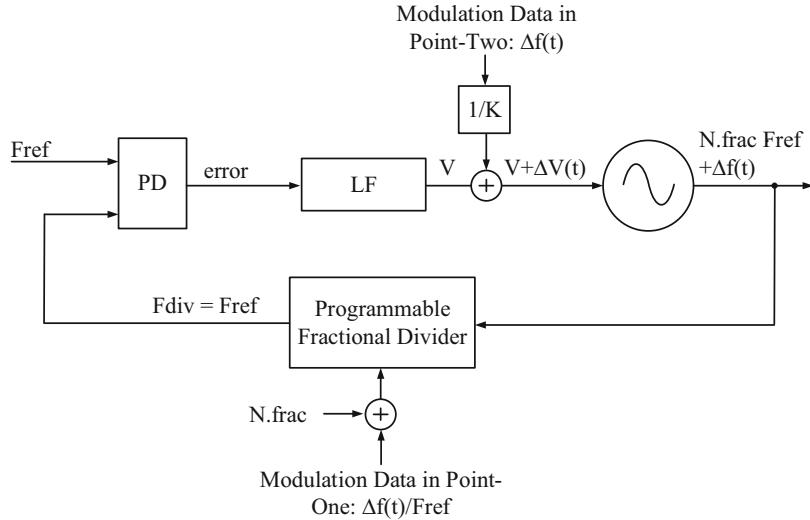


Fig. 3.2 Basic principle of two-point phase/frequency modulation. Note that the modulation data propagates from the point-two to the output, but does not propagate through the loop (cancellation in point-one)

Digital PLLs [Staszewski04] that exploit time-to-digital conversion rather than voltage-domain phase-error processing during frequency synthesis opened up new opportunities in PLL-based phase modulation [Staszewski04, Boos11, Mehta10, Marzin12]. A digital PLL offers good control over the mostly digital loop [Marzin14], and hence it becomes easier to match the modulation injection paths [Marzin12]. Unfortunately, as discussed in Sect. 2.3.2, the quantization error of the TDC used for phase-error detection in digital PLLs typically limits the spectral purity of the phase modulator.

3.1.2 A DTC-Based Fractional-N Subsampling PLL for Phase Modulation

For extremely low noise performance, a phase modulator could use a DTC-based fractional-N subsampling PLL as proposed in [Raczkowski15, Markulic14] and presented in previous chapter. A DTC indeed offers low time quantization excess (see Sect. 2.3.2), while the subsampling phase-error detection core maintains very high phase-error detection gain. This gain suppresses all the in-loop phase noise generators. Since wide PLL bandwidths can be used for VCO phase noise filtering, the loop reaches an optimum in which the output spectral purity is mainly limited by reference noise.

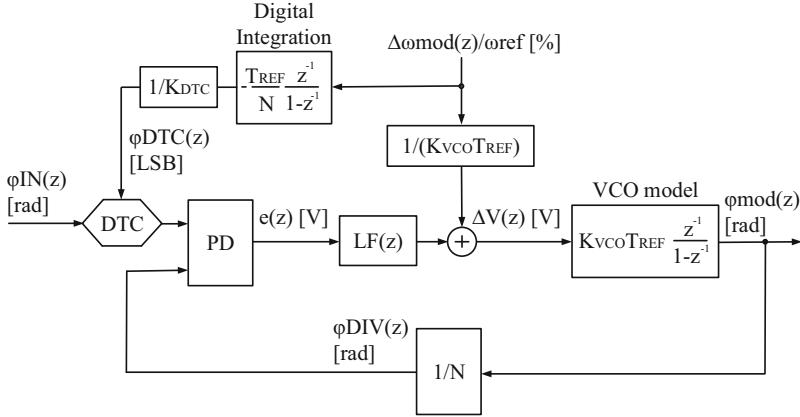


Fig. 3.3 Basic principle of two-point phase/frequency modulation in a DTC-based PLL. The modulation data propagates from the point-two to the output, but does not disturb the loop (cancellation in point-one)

To enable two-point modulation with a fractional-N subsampling PLL, based on a DTC (Fig. 3.3), the problem of injection path matching must be addressed, as well as the challenge of injection point nonlinearity neutralization, both in an analog environment [Markulic16a].

In the following sections, we present a two-point phase modulator in 28-nm CMOS which uses an analog fractional-N, DTC-based subsampling PLL (FNSSPLL). In Sect. 3.2, we start with a quick overview of the FNSSPLL and introduce digital DTC nonlinearity randomization and self-calibration enhancements. Section 3.3 presents the two-point phase modulation capabilities of the FNSSPLL, where the modulation inaccuracies are continuously calibrated in the background. Finally, Sect. 3.4 reports system measurements, followed by conclusions in Sect. 3.5.

3.2 A Self-Calibrated DTC-Based FNSSPLL

3.2.1 Basic Operation of the FNSSPLL

Figure 3.4a depicts a simplified schematic of an FNSSPLL, as described in Chap. 2. For convenience, we briefly summarize the operation of the loop here. The high-frequency VCO output sinewave is directly subsampled at the reference rate in this loop, near the zero-crossing. The G_m produces a current linearly proportional to the sampled voltage which represents the phase difference between reference and VCO phases. That current is steered to the low-pass filter during a fixed signal-

independent pulse width time window [Gao09b], and results in appropriate VCO control adjustment. A frequency acquisition loop (Chap. 2) is implemented for initial frequency lock (not shown in the figure for simplicity). The DTC with ps accuracy serves for fractional residue compensation during fractional-N synthesis. It forces *near-to-zero* voltage sampling by dynamically adjusting the delay, i.e., phase of the reference signal (Fig. 3.4b), such that the PD gain remains determined by the linear part of the sinewave (near the zero-crossing). The DTC requires low phase and quantization noise since its noise is multiplied by N^2 in the transfer to the PLL output [Markulic14]. Moreover, any kind of nonlinearity in the phase comparison path (nonlinear DTC) induces noise folding and fractional spurs [Levantino14].

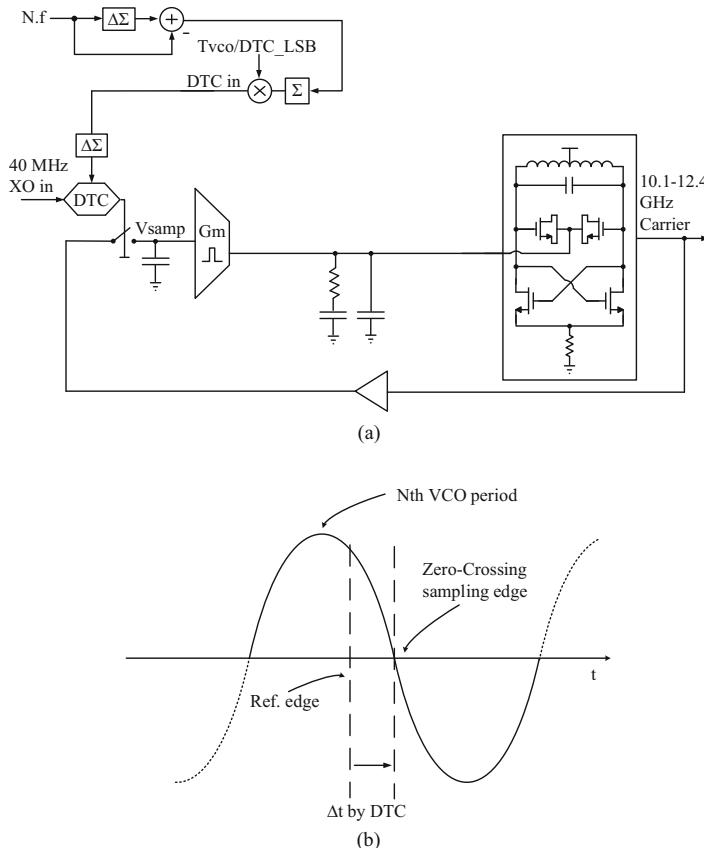


Fig. 3.4 A DTC-based fractional-N subsampling PLL. (a) Simplified schematic. (b) Time-domain operation

3.2.2 The Random-Jump for DTC Quantization Noise Randomization

The DTC input code calculation chain is depicted in Fig. 3.5a. This digital compensation ensures that subsampling events still occur near a zero-crossing, even for a noninteger PLL multiplication number $N.f$. The fractional part of the multiplication number is accumulated and creates a periodic saw-tooth signal Acc . After appropriate scaling, the required value DTC_code_frac is truncated to map on the available discrete codes of the DTC. This truncation is typically performed by a $\Delta\Sigma$ stage which shapes the quantization noise.

In this design, we propose a DTC delay calculation method which removes the colored noise distribution from the DTC quantization error. A pseudo-random integer is added to the original Acc value (Fig. 3.5a) to enforce the DTC to

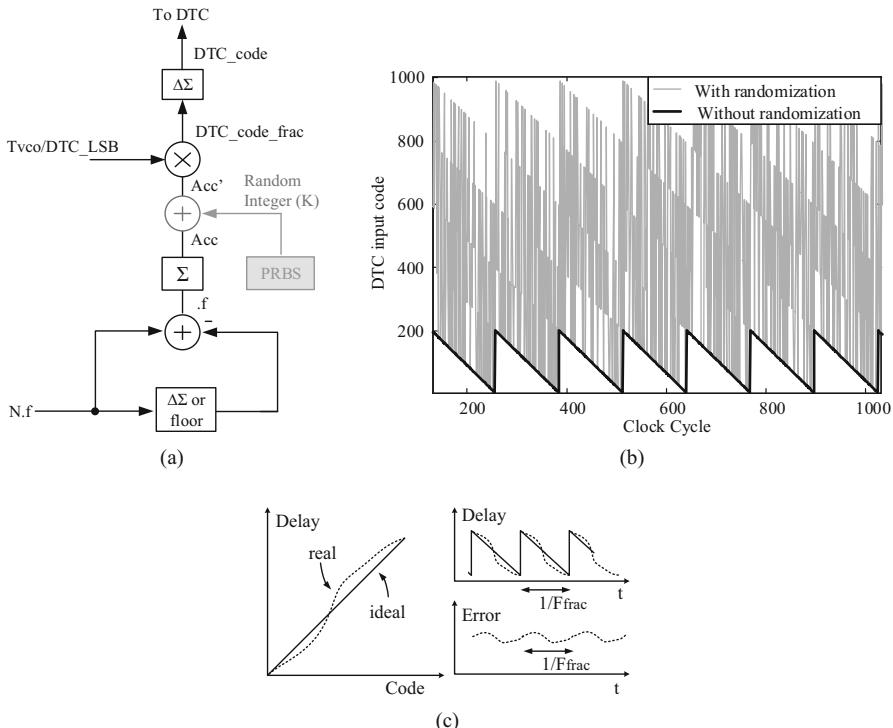


Fig. 3.5 (a) DTC input code calculation path in the fractional-N subsampling PLL with pseudo-random bit sequence (PRBS) generated random integer number (Random-Jump); (b) DTC input with and without Random-Jump calculation path in the time domain. (c) Generation of fractional spurs because of DTC INL

sample random neighboring VCO zero-crossings. The subsampling operation is not altered since phase lock does not discriminate between adjacent integer-N periods [Raczkowski15], and the phase information captured by the subsampler is nearly identical in either case. However, by dithering the exact VCO period with which to subsample, different DTC input codes that induce different quantization errors can now be used for the same fractional residue. A direct consequence of this is a “decolored” quantization noise spectrum. This is shown in Fig. 3.5b for a random integer jump in the range [0:4] generated by the pseudo-random binary sequence (PRBS). The method does not induce additional quantization noise in the system, since the quantization error remains within 0.5 ps, but the spurious content is masked. The resulting spectra are indicated in Fig. 3.6. The method necessitates a DTC with a larger range that covers multiple VCO periods, unless a multiphase, e.g., a quadrature VCO is used. Note that this randomization method has a similar purpose as a higher-order $\Delta\Sigma$ modulator at the beginning of the computational path (for details refer to Sect. 2.3.1). The fundamental difference is that the DC-biased $\Delta\Sigma$ modulator produces a periodic signal (independent of the modulators order)—and the periodicity of the signal is directly dependent on the PLL fractional multiplication number. Simplest case for this is the illustrated 1st-order modulator in Fig. 3.5b, without randomization. The periodicity of the random jump algorithm (based on a PRBS) virtually does not exist (the sequence repeats at a very slow rate which is completely independent of the fractional multiplication number). This in turn leads to better randomization. The random integer addition is also more practical since its range is easily controlled (while the range of a $\Delta\Sigma$ modulator depends on its order).

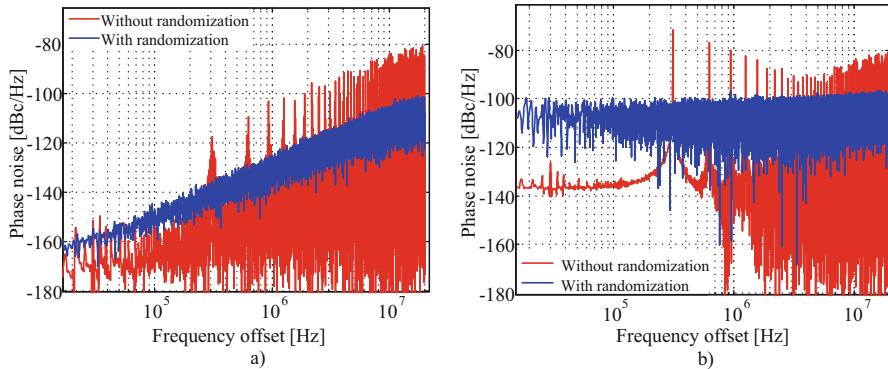


Fig. 3.6 A 10-bit 0.5 ps LSB DTC’s quantization noise spectrum around a 10 GHz fractional carrier ($40\text{ MHz} \times 253 + 2^{-7}$) before the PLL filtering (in-band) without and with randomization **(a)** linear DTC **(b)** nonlinear DTC (± 1 LSB INL error). A randomized nonlinear DTC increases in-band noise floor masking the shaped ($\Delta\Sigma$) quantization noise. Note that a smooth, first-order nonlinearity is used in this example. In the presence of higher-order nonlinearities, spurs are only suppressed and not completely masked by noise

3.2.3 *The Random-Jump for DTC Nonlinearity Randomization*

Any nonlinearity in the phase-error comparison path of a PLL will also induce fractional spurs and noise folding. In the context of analog PLLs, this problem typically originates from a nonlinear charge pump [Meninger06, Lacaita07]. Similarly, in classical digital PLLs where an analog PFD is replaced by a TDC, the TDC causes similar issues [Borremans10, Temporiti10, Staszewski05].

In DTC-based PLLs, this problem arises due to the nonlinearity of the DTC, since the delay through the DTC affects the phase-error detection. Intuitively, if the DTC has a pronounced integral nonlinearity (INL), by application of a periodic signal at its input (Fig. 3.5c) during fractional synthesis, the DTC will create periodical errors. The periodicity is related to the fractional multiplication number and results in spurs at the fractional residue frequency and its harmonics [Levantino14]. Again, the Random-Jump mechanism described earlier helps with spurious tone reduction by removing periodicity from the DTC INL-induced errors. For example, with [0:4] jump range, there are five different DTC codes and five different INL errors at five different zero-crossings which can be used for the same residual phase compensation. The cost of randomization is now, unfortunately, a higher noise floor, i.e., the spur energy is spread, not removed. Nevertheless, the reduction of spurs can be of greater importance than increased noise if the in-band noise is not a limitation, as is the case for many low-power medium performance PLLs.

Figure 3.6 shows the DTC error-induced phase noise (in-band, without PLL filtering) for a fractional 10-GHz VCO output in a simulation with a nonlinear DTC (2 LSB INL), both with and without randomization. Clearly, to correct the spurs without raising the phase noise floor, the spurs themselves must be reduced, meaning that better DTC linearity is required.

3.2.4 *Self-Calibration of the DTC Nonlinearity*

Instead of directing design effort into analog DTC linearization [Ru15, Markulic14] for low spur content at the PLL output, in this design we propose digital self-calibration. If the DTC INL is known (measured a priori), predistortion can be used to remove the predictable errors (see Fig. 3.7a). Since we use a 10-bit LSB DTC from Sect. 2.4.2, a 10-bit Look-Up Table (LUT) could be employed in front of the DTC. The LUT stores the INL value at every DTC code. At a given cycle, the appropriate INL error is simply subtracted from the original DTC input code. With this correction to the nonlinear DTC, the sampler correctly samples near the VCO zero-crossing restoring spur-free PLL operation.

For accurate predistortion, the DTC INL needs to be measured. This is not a trivial task, especially for small DTC time steps (on the order of 100 fs) [Palattella15]. Moreover, the DTC transfer curve is very sensitive to process and

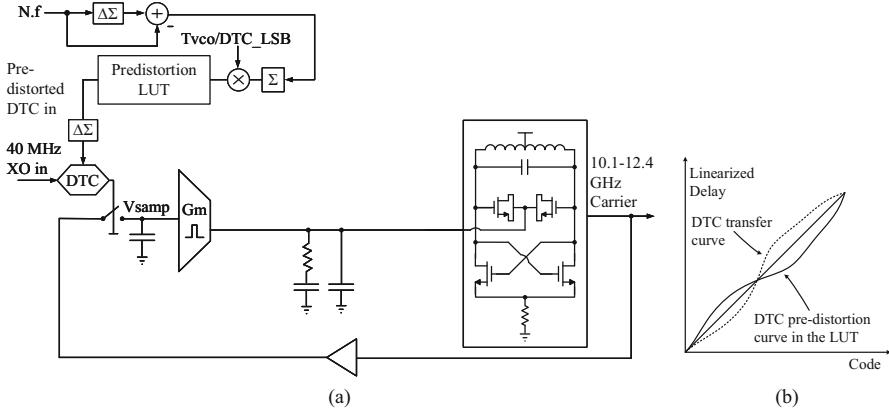


Fig. 3.7 (a) Concept of FNSSPLL DTC predistortion based on a lookup table (LUT); (b) predistortion principle: LUT stores a curve which mimics inverse of the DTC nonlinearity

environmental variations, meaning that the correct predistortion values will drift over time. It is therefore mandatory to have a digital self-calibration which runs in the background while the PLL operates normally.

The central phenomena used for self-calibration is the correlation between the DTC input code and the instantaneously detected phase error. Namely, if the DTC is nonlinear, the sampler will trigger at an offset from the VCO zero-crossing (Fig. 3.8). The nonzero sampled voltage that results is proportional to the timing error at that particular DTC code (plus noise from other sources). The transconductor then outputs a nonzero current into the loop filter. This error current contains information about the DTC (INL) timing error per particular code. The main idea of the proposed calibration scheme is to track the average value of this current¹ for each DTC input code, and to use this information to construct an accurate predistortion function that linearizes the DTC.

Measuring the magnitude of the current produced by the transconductor is cumbersome. Instead, we extract only its sign. This is enough to determine if the sampling instance came before or after the VCO zero-crossing at the respective DTC code. These two situations are depicted in Fig. 3.8. The sign information is used to gradually update the INL LUT (Fig. 3.9a) using a sign least mean square (LMS) algorithm. The current error sign is in every cycle initially extracted by a comparator. The digital (± 1) value is then scaled by $2^{-K} \ll 1$ and finally integrated to the appropriate LUT address defined by the DTC input code. While the PLL operates, the LUT coefficients gradually change towards their optimal values at a convergence

¹ Alternatively, it is possible to track the sampled voltage. However, in the presence of transconductor offset, the sampled voltage is no longer a valid representation of the phase error. The Type-2 Subsampling loop settles in a zero-phase-offset condition which translates to a zero-current state and not in the zero-sampled-voltage state [Gao10].

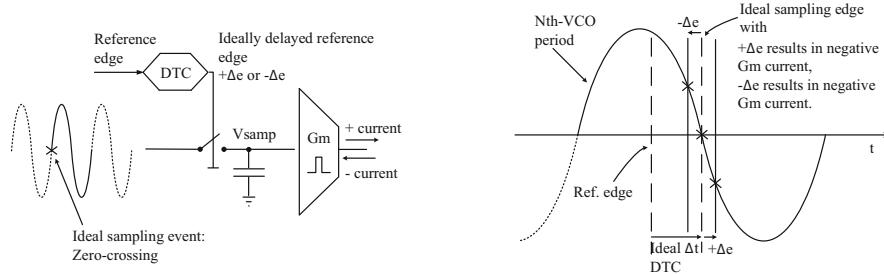


Fig. 3.8 The sign of the Gm output current is correlated to the error of the DTC code

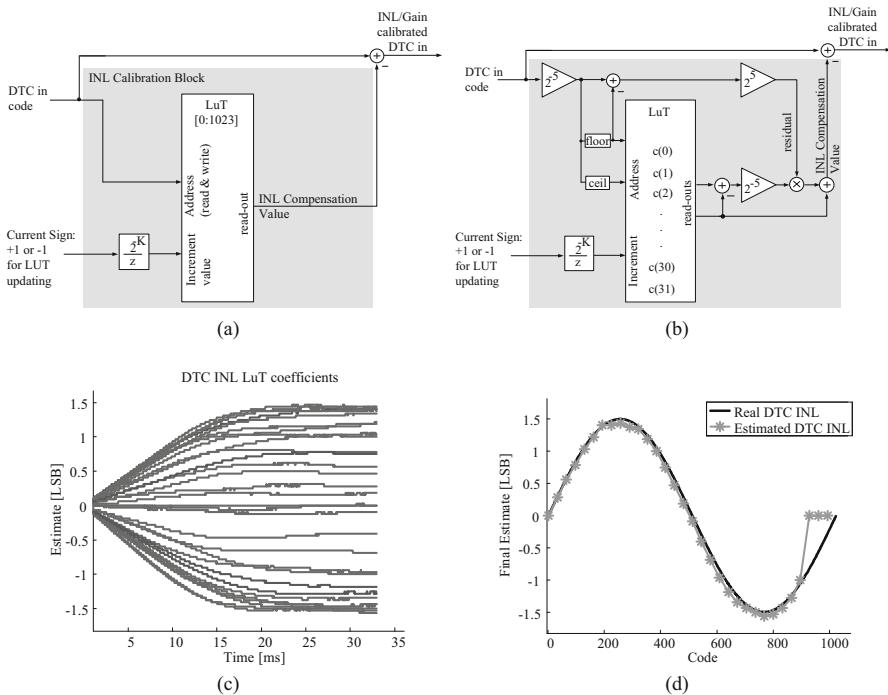


Fig. 3.9 (a) Digital self-calibration of the DTC with a 1024-entry LUT; (b) digital self-calibration of the DTC with a 32-entry LUT and piece-wise linear approximation; (c) time-domain simulation of the FNSSPLL with 32-entry LUT-based DTC calibration with a nonlinear DTC, and (d) final INL estimates

speed dependent on K (Fig. 3.9b). They settle to values that accurately represent the DTC INL, and only move if the INL changes (due to PVT), because the error sign becomes a zero-mean stream per particular code once converged.

To reduce complexity and area, we implement a 32-entry piece-wise-linearized equivalent of the 1024 entry LUT, as shown in Fig. 3.9b. The 32 correction coefficients are spread across 1024 DTC INL values, with linear approximation in-between. The error cancellation is negligibly compromised if the INL curve changes monotonically between two neighboring correction coefficients. In every cycle, the coefficients are updated based on the extracted error sign. The correction value is formed by linearly approximating between the floored and ceiled LUT address, i.e., by approximating between the two instantaneous read-outs from the LUT. To visualize this approximation, a straight gray line is indicated between the “*” symbol (representing 32 LUT correction values) in Fig. 3.9d. This approximation obviously only tracks the “global” INL characteristic of the DTC transfer curve, and cannot compensate for localized DNL errors for every code. But since the DTC capacitive load array has been sized with intrinsic matching for 10-bit accuracy, such remaining errors are negligible. We reuse the DTC architecture from the previous design (Fig. 2.11). The overall DTC nonlinearity is dominated by analog effects such as RC-input-slope-dependent (from the delay stage) comparator delay, which typically results in DTC INL curve similar to the exemplary Fig. 3.9b that can indeed be calibrated in the proposed way.

Another concern is supply ripple in the DTC, or more specifically colored supply noise which could be coupled into the delay path that uses digital gates as described in Sect. 2.4.2. This DTC uses a replica of the sensitive delay stage, in parallel to the main delay path, to ensure that there is no code-dependent power consumption (colored supply noise). In other words, the amount of charge pulled from the supply remains consistent at every reference cycle. Moreover, the comparator in the DTC can optionally work on a regulated supply to avoid code-dependent time switching instant and the related supply bounce. By taking these precautions, the DTC nonlinearity remains unaltered over time. Any other (noncolored) supply noise in the DTC can still exist and can create phase noise, and hence it is necessary to use enough DTC supply decoupling. Finally, the output DTC buffers generate steep slopes to minimize phase noise at the sampler or any dependency of the DTC on VCO frequency/swing.

A simulation of the FNSSPLL during the background calibration process is shown in Fig. 3.9c for a DTC with 3 LSB INL error. The tap 2^K gain is set to 2^{-13} for a calibration time of approximately 15 ms, with correction coefficients initialized to zero. Simulations show that there is insignificant impact on the circuit performance with INL correction resolution below approximately 0.1% of an LSB. Precision in the actual implementation is over-designed and in the order of 0.003% of an LSB.

Although not the focus of this work, obviously initializing the LUT to previously acquired values will reduce the convergence time, since the calibration must then only settle whatever small PVT-related changes have occurred in the interim, i.e., the experienced Δ INL. Also, gear-shifting techniques by initially operating with a larger tap gain will speed up the settling. Since the DTC nonlinearity is independent of the VCO operating frequency/swing, recalibration for different output frequencies is not needed, only background tracking for slow supply and

temperature variations. Finally, in the absence of spurs (due to calibration) the dithering effect of the Random-Jump technique no longer increases the noise floor appreciably.

The linear gain error of the DTC is another effect that must be calibrated, which can be done using separate background calibration algorithm, as described in Sect. 2.3.3. Here, a separate algorithm is not necessary, since it is automatically accounted for in the LUT coefficient update procedure. Just as DTC INL error affects the sampling of the VCO zero-crossing, so does DTC gain error, and as a result, any gain error is removed by the predistortion function provided by the LUT. The gain correction algorithm from Sect. 2.3.3 is still implemented in this design for comparison purposes. It can be run for gain error cancellation but never in parallel with the presented gain/INL calibration algorithm.

3.2.5 Extraction of the Current's Sign and Comparator Offset Compensation

Figure 3.10 depicts the main transconductor and the subcircuit used to extract the sign of its current. The subsampled VCO voltage biases the input pair which then steers a proportional current into the loop filter during pulser opening windows. After that pulse, the same current is steered into a sign extracting node within the given reference period. Since the capacitance at the sign extracting node $V_{extract}$ is kept low (e.g., below 100 fF), a high-voltage swing can be expected there. This is beneficial for the dynamic comparator [Malki14] which senses the node potential in comparison with V_{follow} . The output of the comparator is interpreted as +1 or -1. Importantly, the node $V_{extract}$ is reset to V_{lpf} potential, through a unity gain buffer.

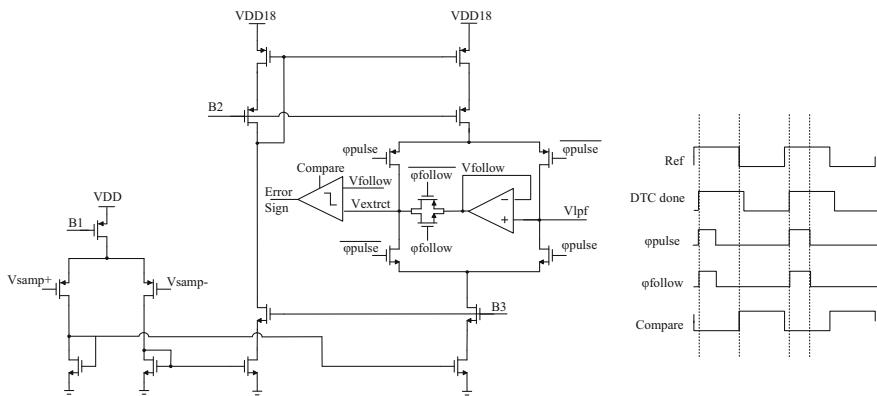


Fig. 3.10 Gm transconductor and current extraction mechanism schematic and operation time diagrams

Any input-referred offset in the comparator is destructive for DTC INL calibration. Its presence will cause the error sign signal to have a nonzero mean, even when the PLL settles into its zero phase-offset condition (type-II loop). Consequently, the DTC INL calibration block receives a stream of plus-ones and minus-ones with a shift in their mean value that the LUT interprets as error and corrects. This results in a slow drift of the LUT correction coefficients at a speed proportional to the offset level. We solve this with digital background compensation (Fig. 3.11a) that cancels the comparator-induced offset before sending the error sign to the INL calibration block, thereby restoring the average error sign value to zero in the settled state. This is achieved by periodically overruling the Error Sign bit with a fixed high or low value to bring the average output back to zero. The overrule period should be proportional to the magnitude of the offset.

To determine this period, the raw Error Sign output is fed into an integrator, which gradually begins to accumulate in the direction related to the offset sign. Its output drives another integrator (after appropriate scaling by $\mu \ll 1$) which will eventually overflow at a rate depending on the magnitude of the comparator offset. The overflow flag is used to address an MUX, steering the appropriate artificial error sign value to the INL calibration algorithm. The original Error Sign signal stream is “cleaned” of offset by these periodic injections of override values, which stops the coefficient drift. The first integrator settles to a value which is proportional to the offset.

This algorithm runs in the background in parallel to the INL calibration. If override injections are triggered infrequently, they appear as low-level noise to the INL calibration and are averaged out. However, an excessive amount of offset results in the integrator overflowing frequently, which dilutes the useful information about INL error. For normal settling of the algorithm, the offset has to be kept below 1σ of the full input swing (assuming Gaussian distribution at the comparator input with standard deviation σ). This is easily achieved by good analog design and layout.

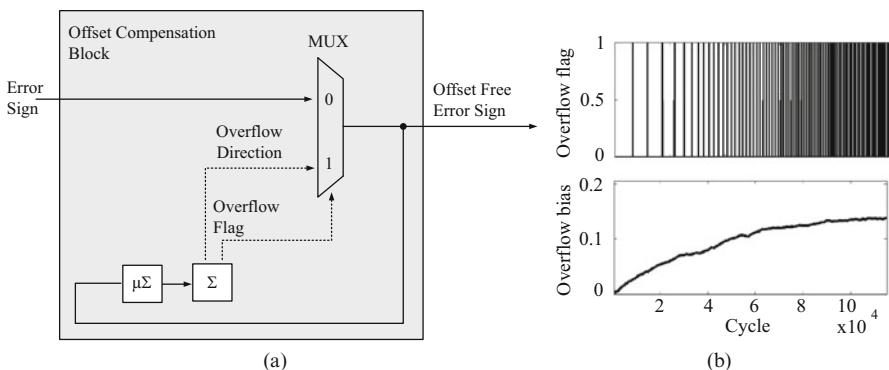


Fig. 3.11 (a) Digital background offset calibration implementation. (b) Background offset calibration simulation with 0.15σ comparator input swing offset. When the loop settles, the overflow flag is activated approximately every 10th cycle

Another important detail is that the system indeed requires two integrators connected in series. An alternative implementation with only a single integrator would not be able to track offset in the background. A single integrator would overflow with frequency proportional to the offset but also dependent on the integration gain. To overflow at the exact frequency of interest, this gain would need accurate manual tuning, which is inappropriate for this application.

3.3 Two-Point Phase Modulator Based on the FNSSPLL

The synthesizer is also enabled with functionality for high-speed (10 Mb/s), high-performance (<−40 dB at 10.24 GHz) GMSK phase modulation. The objective is to exploit the subsampling architecture's excellent phase noise performance for high fundamental EVM performance during modulation.

Similar to other fractional-N PLLs, phase modulation can be achieved by direct modulation of the DTC. The DTC can force controlled “phase errors” (modulation signal) into the loop by purposely moving away from the VCO zero-crossings. The modulation signal is then transferred to the PLL output, although due to the presence of the loop filter the achievable data rate is limited by the loop bandwidth. Higher data rates could be enabled by predistortion of the modulation signal [Perrott97] in front of the DTC. However, the predistortion transfer function is not easily determined since it depends on inaccurate analog components. Furthermore, it can initiate clipping in the PD: the linear range of the subsampling PD is valid only for small phase errors.

To achieve data rates above the loop bandwidth, we use two-point modulation, following similar principles discussed in [Staszewski04, Mehta10, Marzin12]. To realize the second injection point (point-two), we add a separate 8-bit, 50-kHz/LSB frequency modulating bank in the VCO (Fig. 3.12) that enables 10-Mb/s GMSK modulation. Since the modulating bank is clocked at F_{REF} (40 MHz), the highest reasonable modulation bandwidth is 20 MHz, according to the Nyquist–Shannon

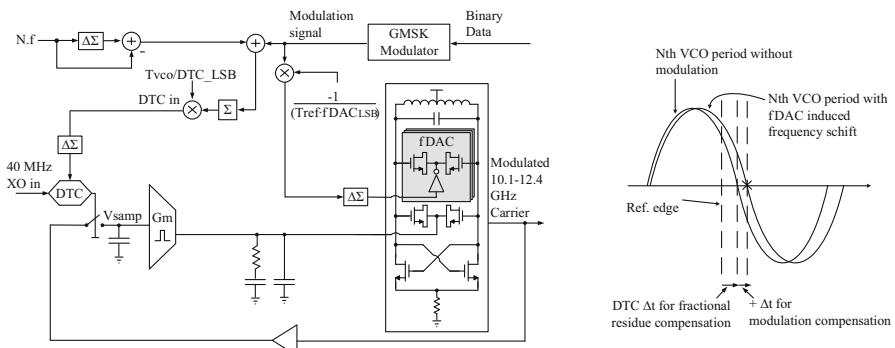


Fig. 3.12 Simplified two-point modulation schematic based on the FNSSPLL and time-domain operation

sampling theorem. We restrict the speed of GMSK to 10 Mb/s to avoid degradation by the spectral replicas. During modulation, the input data is sent through a GMSK modulator which generates a frequency modulating signal for the VCO's modulating fDAC, inducing appropriate instantaneous frequency shifts. In the time domain, this means that the VCO instantaneously changes its period and the amount of accumulated phase over one reference period changes. Therefore, the DTC receives the same modulation signal (expressed in phase) on top of the regular fractional residue compensation signal. It must delay the sampling event by an amount that will keep the VCO zero-crossing sampling unperturbed. Importantly, there are no phase errors in the PD induced by the modulation. The modulation signal transfers unfiltered to the output, i.e., the injection between the two points is cancelled within the PD (sampler) and the loop does not “sense” the modulation data.

3.3.1 *Modulating fDAC INL Calibration*

The EVM of this modulation scheme can be strongly degraded in the presence of nonlinearities and gain errors in the digital-to-modulated phase conversion that occurs in both the DTC and the modulating fDAC. The DTC INL calibration as explained in Sect. 3.2.4 is readily reused here. Indeed, during two-point FNSSPLL modulation the zero-crossing subsampling condition still holds, and the PD works within its linear range. The DTC input code and phase error remain correlated in the presence of INL in the same way as in the synthesizer mode. The modulator's injection point-one (DTC) can, therefore, already be randomized and background calibrated. We show next that a similar algorithm can be implemented for calibration of the modulating fDAC at injection point-two.

Nonlinearity of the modulating fDAC arises due to capacitor mismatch and nonlinear capacitance-to-frequency conversion in the LC tank. We use an INL background calibration technique to linearize its behavior (Fig. 3.13). In the presence of nonlinearity (or gain error), the instantaneous frequency shift induced by the fDAC will be wrong, i.e., the modulated PLL output period will be larger or smaller than expected. This leads to the DTC sampling away from the ideal zero-crossing (late or early) and, consequently, to the transconductor injecting an error current into the loop filter. This current and its extracted sign is strongly correlated with the fDAC input code, and can be exploited for fDAC self-calibration (Fig. 3.13). To linearize the fDAC, we use a predistorting 16-value LUT that linearly approximates between 256 unique fDAC inputs. To correctly correlate the PD error with the fDAC input code, it is necessary to track the derivative of the PD error since the VCO integrates the phase and the modulation signal appears in the frequency domain. The correction coefficients are again gradually increased/decreased (depending on the PD sign) at appropriate addresses, until they converge to positions that cancel out the fDAC nonlinearity, i.e., where the PD output becomes a zero-mean stream per particular fDAC input code. Since DTC and fDAC INL calibration loops cancel uncorrelated static errors independently of each other, they can run simultaneously in the background.

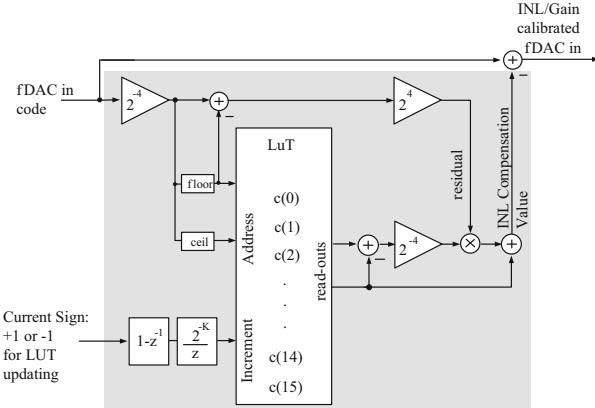


Fig. 3.13 fDAC calibration implementation details

As in the case of the DTC, correction for gain error of the fDAC is also inherently covered by this background calibration technique, since gain error also disturbs the zero-crossing condition of sampler. The LUT correction coefficients are simply adjusted in the background for appropriate gain-error cancellation, on top of the INL correction. Alternatively, a dedicated background calibration algorithm (as proposed in [Marzin12] and Sect.) can be enabled that serves only for initial gain calibration. The latter algorithm is used only for comparison purposes and it is not run in parallel to the presented one.

3.3.2 Delay-Spread Cancellation

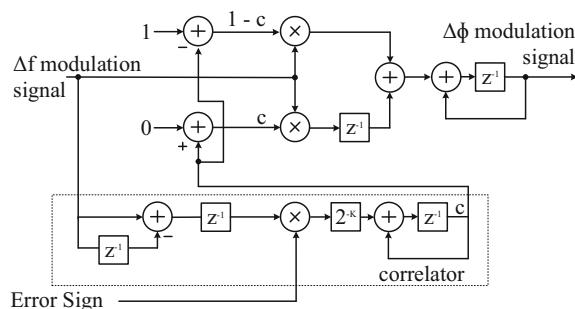
Besides the gain match and linearity requirements, accurate phase modulation necessitates minimal “delay spread” between injection points. This effect originates from mismatch in the time instant at which the modulation signal is injected into the two points. For example, we expect the VCO to instantaneously change its operating frequency (with respect to input code) exactly at the rinsing edge of the reference signal, which is the same edge at which the DTC reacts. This change is obviously noninstantaneous, i.e., the VCO operates for a small portion of time in transition between the original and the new, desired frequency. Simulations show that the consequential *time mismatch* (the VCO accumulates the desired amount of phase a little *later* than originally expected) is small (approximately 100 ps), constant, and not data dependent. Since this spread is small in comparison to the full period of the reference clock cycle (throughout which the VCO accumulates the major part of the phase), the 10-MHz GMSK data remains substantially unfiltered by the loop. In this case thus, even without any compensation, the DTC still samples near the zero-crossing. The degradation imposed by the delay spread is proportional

to the modulation speed. With the same amount of delay spread, i.e., time injection mismatch between the two points, the accumulated error grows proportionally with the modulation speed. Simulations show that at 10-MHz modulation BW (used in this design), delay spread imposes no issues (degradation masked below other circuit noise) as long as it is in the order of 100 ps, for EVM close to -40 dB around a 10-GHz carrier.

A more severe example of delay spread happens in an optional operating mode of the PLL, where the VCO changes its oscillation frequency on the falling edge of the clock (instead of on the rising edge). The VCO thus accumulates the amount of phase as desired by the modulation process—in between the falling edges of the clock, while the DTC still operates from the rising edge. Without the appropriate compensation in the DTC path, this leads to undesired data filtering, i.e., to a situation where the sampler detects modulated phase as an error.

A delay-spread cancellation algorithm (see Fig. 3.14) similar to the one proposed in [Marzin12] has been implemented and is functional on chip. The goal is to inform the DTC about the exact moment at which the VCO changes its operating frequency, with respect to the reference edge at which the DTC reacts. The DTC then produces a modified phase-error compensation signal that takes delay spread into account. The zero-crossing subsampling is still maintained, and the modulation data propagates to the output without filtering. The operation principle of the delay-spread cancellation algorithm is as follows: in the presence of no delay spread, the coefficient c is set to 0, and the two-tap filter in Fig. 3.14 has no influence on the phase data propagation. In the presence of a delay spread, the two-tap filter predistorts the DTC input code to compensate for the effect itself in the digital domain. To do that, the coefficient c needs to be set to $(\tau_{\text{VCO}} - \tau_{\text{DTC}})/T_{\text{REF}}$, where τ_{DTC} and τ_{VCO} stand for the delay from the reference edge at which the DTC and VCO react, respectively, and T_{REF} is the reference period. In this way, the code sent to the DTC compensates for the exact phase integration that happened before its triggering, with respect to both frequencies at which the VCO operated in the previous cycle. For the algorithm to operate in the background, the coefficient c is set (while the modulator operates normally) by correlating the detected error sign with the appropriate modulation data. The product of this two signals is integrated over time (with a very small gain $K < 1$) in order to slowly adjust c . Once when

Fig. 3.14 Delay-spread background cancellation algorithm



c settles to the appropriate value that compensates the delay spread, the input to the correlator becomes a zero mean product (and c only dithers around the optimal value). The algorithm can still be left to operate in the background tracking potential environmental changes. As discussed above, the simulations show that there is no impact on the circuit performance with resolution of c below approximately 0.4%. Precision in the actual implementation is over-designed and in order of 0.01%.

3.4 Experimental Results

A complete system overview is depicted in Fig. 3.15. The prototype IC was fabricated in TSMC 28-nm bulk digital CMOS technology, with an active area of 0.25 mm^2 (Fig. 3.16). It operates on 0.9 V and 1.8 V supplies (IO interface and the Gm stage). Optionally, the VCO can be put in a high-power/performance mode at a 1.4-V supply.

The frequency synthesizer consumes 5.6 mW in total, of which 1.8 mW is for the loop components, 2.7 mW for the VCO, and 1.1 mW for digital circuitry. The VCO tuning range is 10.1–12.4 GHz. The PLL output spectrum profile is first compared in modes with and without DTC Random-Jump (Fig. 3.17). An approximate 2 MHz BW was chosen as a compromise between VCO in-band noise suppression and reference path noise floor. As expected, turning the randomization algorithm on

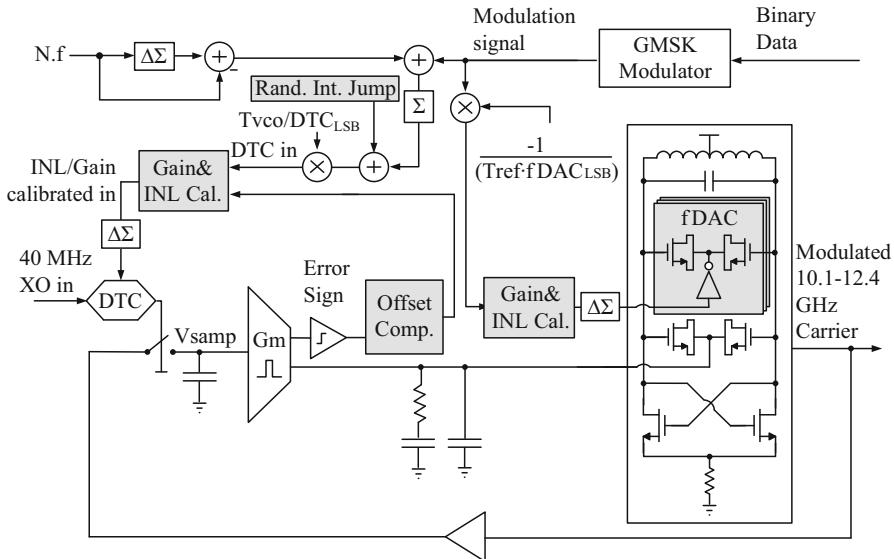


Fig. 3.15 DTC-based FNSSPLL capable of self-calibrated fractional synthesis and two-point modulation

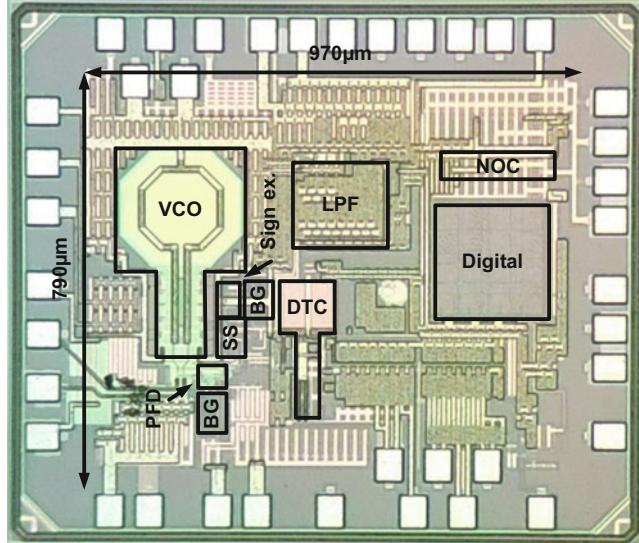


Fig. 3.16 Die microphotograph

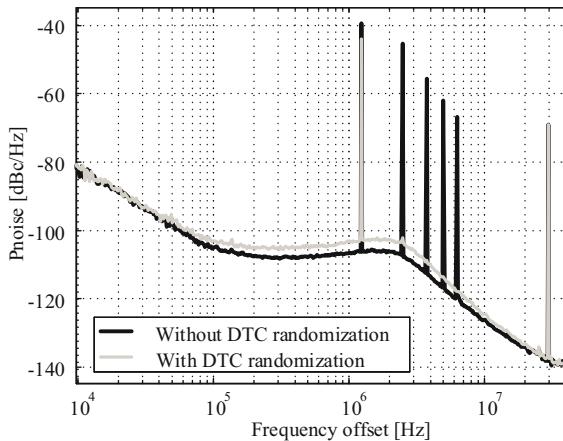


Fig. 3.17 A comparison between PLL output phase noise profile without and with DTC random-jump. Spurs are indicated in dBc (the DTC INL calibration is not enabled). The higher-order spurs disappear with randomization; however, the dominant spur remains only substantially suppressed and is not completely masked by noise. This is due to a larger, higher-order DTC nonlinearity (note that the DTC uses no supply regulation in these measurements)

dithers the fractional spurs, spreading their power across the spectrum. Higher harmonics disappear and the fundamental spur decreases. The cost is, of course, a higher in-band noise floor.

The background calibration of the DTC nonlinearity is enabled next, in parallel with the Random-Jump dithering and comparator offset compensation, which

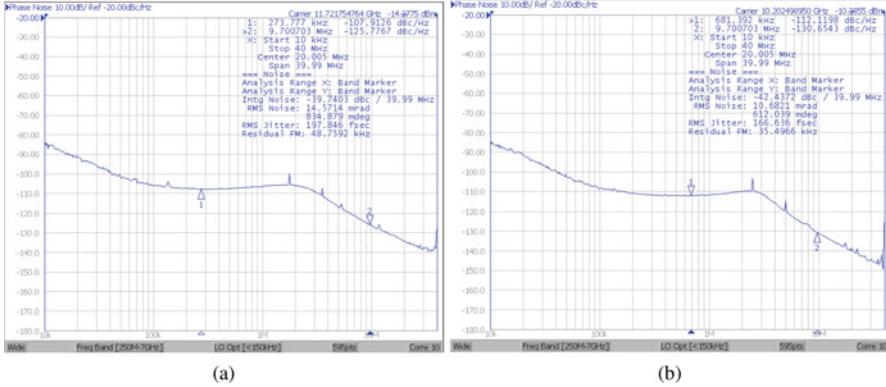
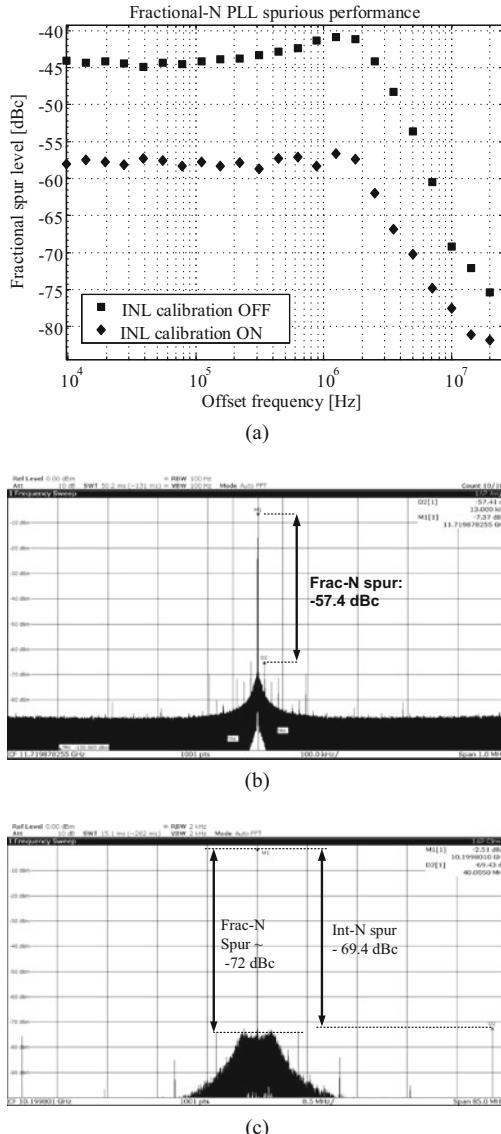


Fig. 3.18 Measured output phase noise profile: (a) low-power VCO, and (b) high-power VCO. The integration range for RMS jitter calculation is 10 kHz–40 MHz and includes all spurs (including worst fractional and integer)

significantly improves the output spectral purity. The measured in-band phase noise around a close-to-integer fractional 11.72-GHz carrier is -107.9 dBc/Hz (Fig. 3.18a). The measured RMS jitter is 198 fs with an integration range from 10 kHz to 40 MHz and all spurs (worst fractional and integer-N) included. In integer-N mode, the RMS jitter is 176 fs (only 22 fs below) which proves that the system's nonidealities were well identified and successfully mitigated in this design. The measured spurious performance is shown in Fig. 3.19. The worst fractional spur before calibration appears at -41 dBc but drops by 15.6 dB after calibration to -56.6 dBc. The reference spur is at -69 dBc. The PLL achieves a competitive -246.6 – -247.6 dB FOM (worst-case/best-case fractional-N mode). Integer-N FOM, with the digital circuitry disabled, is -248.5 dB. Moreover, the PLL can run in high-performance mode where the VCO operates at up to 22.4 mW from a 1.4-V supply (Fig. 3.18b). The RMS jitter (all spurs included) is then between 167 and 147 fs for worst-case fractional and integer-N mode. The small 20 fs difference, or 1 dB in terms of total integrated phase noise, shows again a well-calibrated environment. The in-band phase noise drops to -113.1 dBc/Hz. Due to excessive VCO power consumption, the FOM is now reduced to -242.6 – -241.5 dB in the worst/best case. A summary of the PLL performance and comparison to other similar work [Chen15, Huang14, Gao09a] is depicted in Table 3.1.

During modulation, we operate the VCO in high-power mode. The modulating fDAC was mistakenly designed to produce 5 times larger LSB than targeted (241 kHz/LSB instead 50 kHz/LSB) at low power, i.e., at low output VCO swings. By running the VCO in high-power mode (6.4–22.4 mW), the LSB reduces to 113–83 kHz/LSB. This is due to an increased swing at the VCO output which minimizes the effective ΔC [fF] of a digital cell [Hung06] in the modulating fDAC. With a higher power consumption cost, we ensure that the fDAC quantization noise does not limit the EVM performance. Figure 3.20 shows the measured 10-Mb/s

Fig. 3.19 (a) Fractional spur with and without calibration at different fractional offsets; (b) calibrated spectrum analyzer output at a deep in-band fractional channel after calibration; (c) spectrum analyzer plot at out-of-band fractional channel after calibration



GMSK spectrum and constellation with all the background calibration algorithms enabled, around a fractional carrier close to 10.24 GHz. The modulated output has an undistorted spectral profile with bandwidth larger than 10 MHz even though the PLL loop BW is approximately 2.5 MHz. The best EVM RMS with DTC/DAC calibration is at -40.5 dB, which is only 1.9 dB above the fundamental limit of the total integrated phase noise of a high-power mode synthesizer. Notably, EVM scales with $20 \log_{10}(N)$ which means that the modulator would provide -49.6 dB

Table 3.1 Performance summary of the FNSSPLL2 and comparison to other fractional-N CMOS PLLs

	This work	[Chen15]	[Huang14]	[Gao16]
Architecture	Subsampling analog	Subsampling digital	Subsampling analog	Sampling digital
Reference (MHz)	40	49,15	48	40
Output (GHz)	10.1–12.4	2.6–3.9	2.2–2.4	2.7–4.3
Tuning range (%)	20	40	9	46
Bandwidth (kHz)	1800	700	500	500 ^c
In-band PN (dBc/Hz) ^a	−108	−98	−98	−99 ^c
Ref. Spur (dBc)	−69	−60	−55	−78
Worst frac. spur (dBc)	−56.6	−62	−53	−54
RMS jitter (fs) best/worst	176/197 (10k–40M)	226/240 (1k–100M)	266/400 (10k–30M)	159/n.a.
Power (mW)	5.6	11.5	17.3	8.2
Frac-N FOM (dB) ^b worst/best	−246.6/−247.6	−241.8/−242.3	−235.6/−239.1	n.a./−246.8 (10k–40M)
Area (mm ²)	0.25	0.23	0.75	0.3
Process (nm)	28	65	180	28

^a Scaled to 11.72 GHz

^b $FOM_{PLL} = 10 \log_{10} \left[\left(\frac{\sigma_{I,PLL}}{1s} \right)^2 \left(\frac{P_{PLL}}{1mW} \right) \right]$

^c Estimated from source analyzer measurement

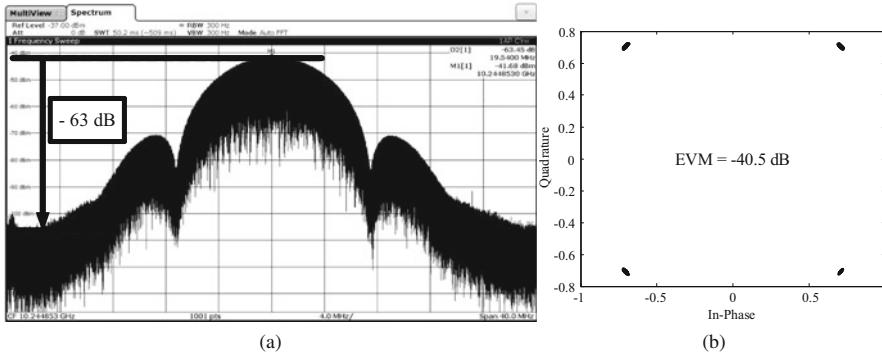


Fig. 3.20 GMSK spectrum and EVM with self-calibration enabled (10 Mb/s, close to a 10.24-GHz fractional carrier)

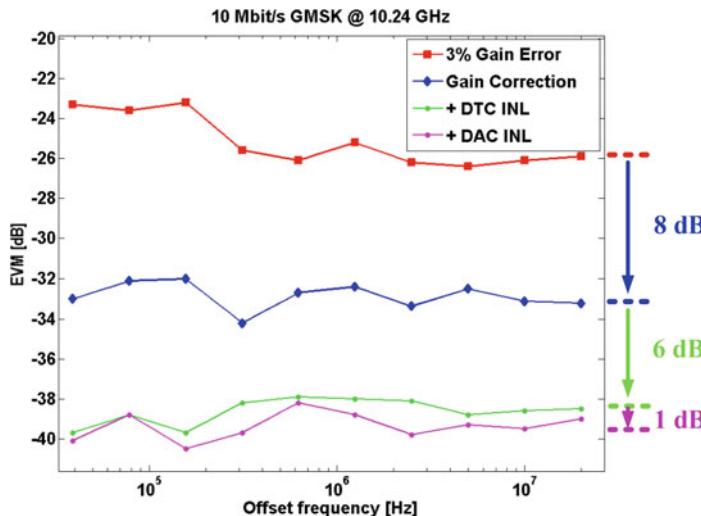


Fig. 3.21 Measured EVM at different fractional offsets with and without calibration

EVM at 3.6 GHz. To prove the importance of calibration, we show EVM at different fractional offsets (Fig. 3.21) with and without calibration. Using only gain correction results on average in 8-dB EVM improvement. The DTC and fDAC INL calibration results in additional 7-dB EVM improvement (6 dB after DTC calibration and then additional 1 dB after fDAC calibration). Enabling the delay-spread algorithm does not change the results, as predicted. The DTC INL calibration proves to have a larger impact on the overall system performance. The overall performance overview and comparison to state-of-the-art is given in Table 3.2, with a note that this phase modulator's EVM is state of the art and exceeds similar work.

Table 3.2 Performance summary of the phase/frequency modulator and comparison to the state of the art

	Low power	High power	[Marzini12]	[Li15]	[Shanan09]	[Xu14]
Architecture	Direct-FM	Direct-FM	Direct-FM	Direct-FM	Direct-FM	Direct-FM
Reference (MHz)	40	40	n.a.	52	26	
Output (GHz)	10.1–12.4	2.9–4.0	1.5–1.8	2.4	1.7–2.1	
Power dissipation (mW)	8.1	25.4	5	10	16	6.9
Modulation type	GMSK	GMSK	GMSK	GMSK	GMSK	GMSK
Data rate (Mb/s)	10	10	10	10	2	1.08
Energy/bit (nJ/bit)	0.8	2.5	0.5	1	8	6.4
Integrated PN (dBc) ^a	−41.7	−42.4	−29.9	n.a.	n.a.	−20.3
EVM RMS (dB) ^a	−37.4	−40.5	−26.9	n.a.	−5.4	−13.9
Out-of-band emission (dB _R)	−55	−63	−56	n.a.	n.a.	
Area (mm ²)	0.25		0.5	0.5	1.1	0.49
Process (nm)	28	65	65	180	180	65

^a Scaled to 10.24 GHz

3.5 Conclusion

We have presented a DTC-based FNSSPLL that achieves state-of-the-art performance during synthesis (-246.6 – -247.6 -dB FOM worst-case/best-case fractional-N mode) and wideband, 10-Mbit/s GMSK modulation (-40.5 -dB EVM around a 10-GHz carrier) in 28-nm CMOS. This performance is derived from the analog subsampling phase-error detection core and efficient digital background calibration. This PLL can be used for low-power, low-noise LO synthesis. In the modulator mode, the system prototype offers an attractive solution for accurate, wideband phase modulation that can be used as part of a polar transmitter.

The presented low-jitter PLL operates without any significant performance gap between integer and fractional modes and enables wideband modulation. With respect to this, we recognize that the loop can still be improved in a different direction: *digitization*, which should ideally come without compromise in performance. Replacing the analog core (analog loop filter) of the FNSSPLL with a digital one can still add to area savings. Additionally, the phase modulator could be employed in a highly accurate polar transmitter architecture. These challenges are explored in the following chapter.

References

- [Boos11] Z. Boos, A. Menkhoff, F. Kuttner, M. Schimper, J. Moreira, H. Geltinger, T. Gossmann, P. Pfann, A. Belitzer, T. Bauernfeind, A fully digital multimode polar transmitter employing 17b RF DAC in 3G mode, in *2011 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, Piscataway, 2011), pp. 376–378
- [Borremans10] J. Borremans, K. Vengattaramane, V. Giannini, J. Craninckx, A 86MHz-to-12GHz digital-intensive phase-modulated fractional-N PLL using a 15pJ/Shot 5ps TDC in 40nm digital CMOS, in *2010 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, Piscataway, 2010), pp. 480–481
- [Chen15] Z.-Z. Chen, Y.-H. Wang, J. Shin, Y. Zhao, S.A. Mirhaj, Y.-C. Kuan, H.-N. Chen, C.-P. Jou, M.-H. Tsai, F.-L. Hsueh et al., 14.9 Sub-sampling all-digital fractional-N frequency synthesizer with -111 dBc/Hz in-band phase noise and an FOM of -242 dB, in *2015 IEEE International Solid-State Circuits Conference-(ISSCC)* (IEEE, Piscataway, 2015), pp. 1–3
- [Dürdodt01] C. Dürdodt, M. Friedrich, C. Grewing, M. Hammes, A. Hanke, S. Heinen, J. Oehm, D. Pham-Stäbner, D. Seippel, D. Theil et al., A low-IF Rx two-point Sigma-Delta modulation Tx CMOS single-chip bluetooth solution. *IEEE Trans. Microwave Theory Tech.* **49**(9), 1531–1537 (2001)
- [Gao09b] X. Gao, E. Klumperink, M. Bohsali, B. Nauta, A low noise sub-sampling PLL in which divider noise is eliminated and PD/CP noise is not multiplied by N^2 . *IEEE J. Solid State Circuits* **44**(12), 3253–3263 (2009)
- [Gao09a] X. Gao, E. Klumperink, P. Geraedts, B. Nauta, Jitter analysis and a benchmarking figure-of-merit for phase-locked loops. *IEEE Trans. Circuits Syst. Express Briefs* **56**(2), 117–121 (2009)

- [Gao10] X. Gao, E. Klumperink, G. Soccia, M. Bohsali, B. Nauta, Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector. *IEEE J. Solid State Circuits* **45**(9), 1809–1821 (2010)
- [Gao16] X. Gao, O. Burg, H. Wang, W. Wu, C.-T. Tu, K. Manetakis, F. Zhang, L. Tee, M. Yayla, S. Xiang et al., 9.6 A 2.7-to-4.3 GHz, 0.16 psrms-jitter, -246.8 dB-FOM, digital fractional-N sampling PLL in 28nm CMOS, in *2016 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2016), pp. 174–175
- [Huang14] P.-C. Huang, W.-S. Chang, T.-C. Lee, 21.2 A 2.3 GHz fractional-N dividerless phase-locked loop with -112dBc/Hz in-band phase noise, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, Piscataway, 2014), pp. 362–363
- [Hung06] C. Hung, R.B. Staszewski, N. Barton, M. Lee, D. Leipold, A digitally controlled oscillator system for SAW-less transmitters in cellular handsets. *IEEE J. Solid State Circuits* **41**(5), 1160 (2006)
- [Lacaita07] A.L. Lacaita, S. Levantino, C. Samori, *Integrated Frequency Synthesizers for Wireless Systems*. (Cambridge University Press, New York, 2007)
- [Levantino14] S. Levantino, G. Marzin, C. Samori, An adaptive pre-distortion technique to mitigate the DTC nonlinearity in digital PLLs. *IEEE J. Solid State Circuits* **49**(8), 1762–1772 (2014)
- [Li15] X. Li, L. Sitao, X. Liu, N. Xu, W. Rhee, W. Jia, Z. Wang, A 10 Mb/s hybrid two-point modulator with front-end phase selection and dual-path DCO modulation, in *2015 IEEE International Wireless Symposium (IWS)* (IEEE, Piscataway, 2015), 1–4
- [Malki14] B. Malki, T. Yamamoto, B. Verbruggen, P. Wambacq, J. Craninckx, A 70 dB DR 10 b 0-to-80 MS/s current-integrating SAR ADC with adaptive dynamic range. *IEEE J. Solid State Circuits* **49**(5), 1173–1183 (2014)
- [Markulic14] N. Markulic, K. Raczkowski, P. Wambacq, J. Craninckx, A 10-bit, 550-fs step Digital-to-Time Converter in 28nm CMOS, in *ESSCIRC 2014 - 40th European Solid State Circuits Conference (ESSCIRC)*, Sept 2014, pp. 79–82 (2014)
- [Markulic16a] N. Markulic, K. Raczkowski, E. Martens, P.E.P. Filho, B. Hershberg, P. Wambacq, J. Craninckx, 9.7 A self-calibrated 10Mb/s phase modulator with -37.4dB EVM based on a 10.1-to-12.4GHz, -246.6dB-FOM, fractional-N subsampling PLL, in *2016 IEEE International Solid-State Circuits Conference (ISSCC)*, Jan 2016, pp. 176–177 (2016)
- [Markulic16b] N. Markulic, K. Raczkowski, E. Martens, P.E.P. Filho, B. Hershberg, P. Wambacq, J. Craninckx, A DTC-based subsampling PLL capable of self-calibrated fractional synthesis and two-point modulation. *IEEE J. Solid State Circuits* **51**(12), 3078–3092 (2016)
- [Marzin12] G. Marzin, S. Levantino, C. Samori, A.L. Lacaita, A 20 Mb/s phase modulator based on a 3.6 GHz digital PLL with -36 dB EVM at 5 mW power. *IEEE J. Solid State Circuits* **47**(12), 2974–2988 (2012)
- [Marzin14] G. Marzin, S. Levantino, C. Samori, A.L. Lacaita, 2.9 A background calibration technique to control bandwidth in digital PLLs, in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, Piscataway, 2014), pp. 54–55
- [Mehta10] J. Mehta, R.B. Staszewski, O. Eliezer, S. Rezeq, K. Waheed, M. Entezari, G. Feygin, S. Vemulapalli, V. Zoicas, C.-M. Hung et al., A 0.8 mm 2 all-digital SAW-less polar transmitter in 65nm EDGE SoC, in *2010 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, Piscataway, 2010), pp. 58–59
- [Meninger06] S.E. Meninger, M.H. Perrott, A 1-MHZ bandwidth 3.6-GHz 0.18-um CMOS fractional-N synthesizer utilizing a hybrid PFD/DAC structure for reduced broadband phase noise. *IEEE J. Solid State Circuits* **41**(4), 966–980 (2006)

- [Palattella15] C. Palattella, E.A. Klumperink, J. Ru, B. Nauta, A sensitive method to measure the integral nonlinearity of a digital-to-time converter based on phase modulation. *IEEE Trans. Circuits Syst. Express Briefs* **62**(8), 741–745 (2015)
- [Pamarti04] S. Pamarti, L. Jansson, I. Galton, A wideband 2.4-GHz delta-sigma fractional-NPLL with 1-Mb/s in-loop modulation. *IEEE J. Solid State Circuits* **39**(1), 49–62 (2004)
- [Perrott97] M.H. Perrott, T.L. Tewksbury III, C.G. Sodini, A 27-mW CMOS fractional-N synthesizer using digital compensation for 2.5-Mb/s GFSK modulation. *IEEE J. Solid State Circuits* **32**(12), 2048–2060 (1997)
- [Raczkowski15] K. Raczkowski, N. Markulic, B. Hershberg, J. Craninckx, A 9.2–12.7 GHz wideband fractional-N subsampling PLL in 28 nm CMOS with 280 fs RMS jitter. *IEEE J. Solid State Circuits* **50**(5), 1203–1213 (2015)
- [Riley93] T.A. Riley, M.A. Copeland, T.A. Kwasniewski, Delta-sigma modulation in fractional-N frequency synthesis. *IEEE J. Solid State Circuits* **28**(5), 553–559 (1993)
- [Ru15] J.Z. Ru, C. Palattella, P. Geraedts, E. Klumperink, B. Nauta, A high-linearity digital-to-time converter technique: constant-slope charging. *IEEE J. Solid State Circuits* **50**(6), pp. 1412–1423 (2015)
- [Shanan09] H. Shanan, G. Retz, K. Mulvaney, P. Quinlan, A 2.4 GHz 2Mb/s versatile PLL-based transmitter using digital pre-emphasis and auto calibration in 0.18 μ m CMOS for WPAN, in *IEEE International Solid-State Circuits Conference-Digest of Technical Papers, 2009. ISSCC 2009* (IEEE, Piscataway, 2009), pp. 420–421
- [Staszewski04] R.B. Staszewski, K. Muhammad, D. Leipold, C.-M. Hung, Y.-C. Ho, J.L. Wallberg, C. Fernando, K. Maggio, R. Staszewski, T. Jung et al., All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS. *IEEE J. Solid State Circuits* **39**(12), 2278–2291 (2004)
- [Staszewski05] R.B. Staszewski, J.L. Wallberg, S. Rezeq, C.-M. Hung, O.E. Eliezer, S.K. Vemulapalli, C. Fernando, K. Maggio, R. Staszewski, N. Barton et al., All-digital PLL and transmitter for mobile phones. *IEEE J. Solid State Circuits* **40**(12), 2469–2482 (2005)
- [Temporiti10] E. Temporiti, C. Weltin-Wu, D. Baldi, M. Cusmai, F. Svelto, A 3.5 GHz wideband ADPLL with fractional spur suppression through TDC dithering and feedforward compensation. *IEEE J. Solid State Circuits* **45**(12), 2723–2736 (2010)
- [Xu14] N. Xu, W. Rhee, Z. Wang, A hybrid loop two-point modulator without DCO nonlinearity calibration by utilizing 1 bit high-pass modulation. *IEEE J. Solid State Circuits* **49**(10), 2172–2186 (2014)

Chapter 4

A Background-Calibrated Digital Subsampling Polar Transmitter



4.1 Introduction

Today's wireless community is facing the problem of spectral scarcity. With many communication standards sharing the same environment, it becomes highly desirable to make use of spectrally efficient links. The upcoming 802.11ax WI-FI standard, for example (backward compatible with 802.11g/n/ac), envisions complex 1024 QAM constellation schemes with modulation bandwidths greater or equal to 80 MHz, with OFDM and close sub-carrier spacing around 2.4 and 5 GHz [Working Gr17]. The linearity and noise specifications of the supporting radio front-ends become in this context extremely challenging.

On the transmit side, polar architectures have proved to be attractive, digital-friendly solutions [Ba16, Fulde17, Madoglio17, Zheng15] that can make use of power efficient amplitude modulators [Chowdhury11, Yoo13] ensuring overall system power efficiency. At the same time, for high spectral purity, the digital-to-transmitted signal conversion needs to be stringently linear, substantially uncorrupted by the quantization or environmental noise. Unfortunately, the efficient PA architectures come typically at the price of distorted Amplitude Modulation (AM) [Chowdhury11, Chowdhury12]. The problem of nonlinearity is similarly present in wideband PLL-based Phase Modulation (PM) as described in Sect. 3.3. To overcome these difficulties and to ensure linearity, transmitters make use of offline predistortion [Madoglio17, Zheng15, Chowdhury11, Yoo11] correcting in the digital domain for the presence of AM-to-AM, PM-to-PM nonlinearities, or cross-correlation mechanisms such as AM-to-PM distortion.

For accurate predistortion, it is not enough to rely on estimates from simulations since they are prone to modeling inaccuracies. Moreover, a silicon implemented chip is sensitive to PVT variations, hence an optimal predistortion setting is likely to fluctuate over time. Measurement-based predistortion (that covers all possible PVT variations) can significantly increase production/testing expense in large vol-

ume fabrication. Background tracking and distortion cancellation (self-calibration) become therefore an attractive system feature, especially in environments that make use of complex constellations, where no linearity related compromise is allowed.

In this chapter, we present a polar transmitter which consists of a low-noise digital subsampling PLL that allows two-point PM and a DPA for AM. The architecture of the proposed system, unlike in a typical polar transmitter, places the DPA within the phase-locked loop. This means that the PLL's phase-error detection mechanism, i.e., subsampling happens directly at the DPA output, instead of at the VCO output. Thanks to that, the phase-locked loop becomes sensitive not only to phase errors but also to *modulation amplitude*. This feature enables accurate system linearization and (AM-to-PM) distortion filtering. Analog nonidealities such as PM-to-PM, AM-to-AM can be detected and cancelled digitally, in the background, while the transmitter operates normally. To highlight this system's attribute and to distinguish it from a typical polar transmitter (where the PLL and amplitude modulator are separated), we refer to it as: the background-calibrated digital *subsampling polar transmitter*, SSPTX.

In Sect. 4.2, we introduce the system's architecture and discuss the SSPTX prototype specifications. The study of background linearization and distortion suppression techniques enabled by the proposed architecture are discussed in Sect. 4.3 and 4.4, respectively. Analog building blocks are presented in Sect. 4.5. Measured results are discussed in Sect. 4.6. The paper is summarized in the concluding Sect. 4.7.

4.2 System Overview

A simplified block diagram of the proposed digital SSPTX followed by its small signal phase-domain model is depicted in Fig. 4.1.

4.2.1 A Digital Fractional-N Subsampling PLL

The system operates as a PLL when there is no modulation, in which the phase-error detection is performed by the subsampling mechanism (see Fig. 4.2a), described throughout the Chaps. 2 and 3 of this book, well known in analog subsampling PLLs [Gao09b, Markulic16a, Markulic16b, Raczkowski15, Markulic14]. In the PLL mode, where no AM or PM is present, we assume $AM(z)=1$ and $\Delta\omega_{mod} = 0$, i.e., the code sent to the DPA is constant, and there are no modulation data injected in the phase/frequency modulation path through the DAC nor the DTC. The high frequency VCO output is buffered through the DPA that ideally has no influence on the phase signal propagation to the output. The DPA drives a $50\ \Omega$ load (antenna) and its output is a sinewave, since the matching network filters the unwanted harmonics. This sinusoidal DPA output is subsampled at the reference rate and the

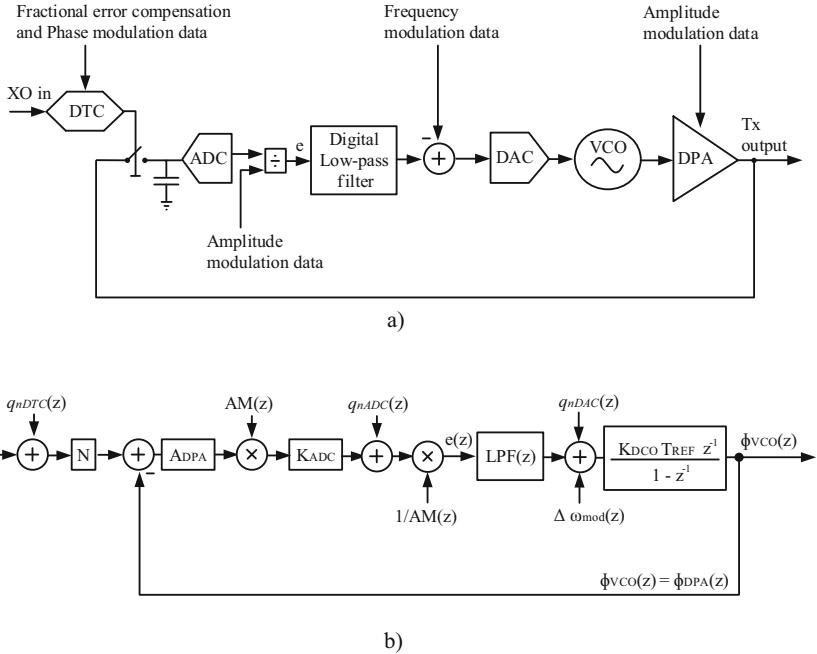


Fig. 4.1 (a) Simplified block diagram of the proposed digital SSPTX. (b) Small signal phase-domain model: $q_{nDTC}(z)$, $q_{nADC}(z)$, and $q_{nDAC}(z)$ represent quantization noise signals from the DTC, ADC, and the DAC, respectively. ϕ_{DTC} contains, both, fractional residue compensation information and phase modulation data. $\Delta\omega_{mod}$ is the scaled frequency modulation data that are perfectly cancelled in the phase-error comparison path by the DTC. N is the PLL multiplication number

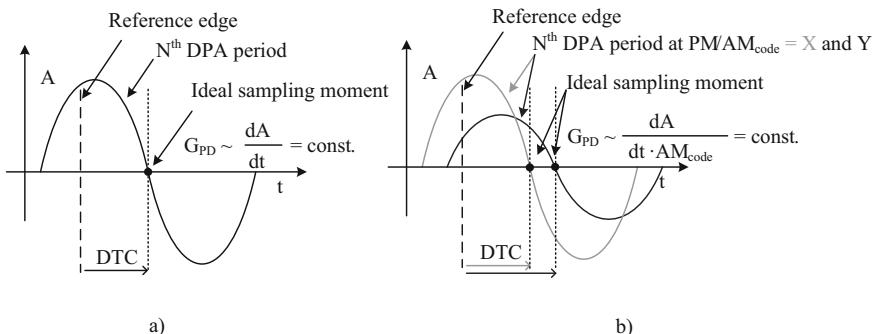


Fig. 4.2 Basic SSPTX phase-error detection principle in PLL and TX modes. (a) DTC delay for fractional error compensation ensures near-to-zero crossing sampling. G_{PD} is linearly proportional to the DPA output slew rate. (b) DTC delay for fractional error and PM compensation ensures near-to-zero crossing sampling. Two point AM injection ensures constant G_{PD}

exact sampling edge is controlled by a DTC. The DTC is transparent during integer-N PLL operation while, in fractional-N operation mode, it serves for fractional residue compensation. In other words, the DTC systematically delays the sampling instant by an a priori known value to force near-to-zero voltage sampling during fractional-N lock (Chap. 2). If there is no mismatch between the input and the output phase, the output sinewave zero-crossings are sampled, i.e., the ADC outputs a zero. If there is some mismatch between the output and the input phase, a non-zero voltage is subsampled, and the ADC produces a code linearly proportional (small signal approximation) to the detected phase error. This code is digitally low-pass filtered and sent into a DAC to adjust the VCO's output phase, forcing a zero-phase-offset condition between the input and the output of the PLL. The DAC controlled VCO is modeled as a DCO in Fig. 4.1 for simplicity.

The phase-error detection gain is, specifically for a subsampling PLL, directly proportional to the PLL output amplitude [Gao09b] and in this context it is defined by:

$$G_{\text{PD}} = A_{\text{DPA}} \cdot K_{\text{ADC}}, \quad (4.1)$$

where A_{DPA} stands for DPA output amplitude and K_{ADC} is the ADC conversion gain (we assume for now a loss-less sampler and no amplification in front of the ADC). Thanks to the fact that G_{PD} can be set to *a large value*, analog subsampling PLLs operate in a condition where the in-band noise generators are highly suppressed [Gao09b, Gao15]. The PLL output becomes then mainly corrupted by the VCO noise which can be efficiently filtered in a wide-PLL-bandwidth configuration since the reference noise is typically low. This likely enables high phase noise performance and PLL FOM [Gao09a].

A *digital* subsampling loop differs performance-wise from its original analog predecessor in added quantization noise at several nodes in the system. To achieve the same performance level in a digitized loop the quantization error residue is carefully handled, as addressed next.

The DTC quantization noise (see Sect. 3.2 from Chap. 3) is randomized by the random jump algorithm and shaped by a delta-sigma modulator, hence it has very little influence in-band, while out-of-band it is compensated by the low-pass PLL filtering. Moreover, as shown later in Sect. 4.2.3, the DTC quantization noise can be subtracted in the digital domain from the sampled phase-error information, avoiding its propagation to the output all together.

The ADC quantization error transfer function to the PLL output is given by:

$$\frac{\Phi_{\text{VCO}}(z)}{q n_{\text{ADC}}(z)} = \frac{1}{AM \cdot A_{\text{DPA}} \cdot K_{\text{ADC}}} \cdot \frac{G_{\text{OL}}(z)}{1 + G_{\text{OL}}(z)}, \quad (4.2)$$

where $G_{\text{OL}}(z)$ is the open loop gain defined as:

$$G_{\text{OL}}(z) = A_{\text{DPA}} \cdot K_{\text{ADC}} \cdot LPF(z) \cdot \frac{K_{\text{DCO}} \cdot T_{\text{REF}} \cdot z^{-1}}{1 + z^{-1}}, \quad (4.3)$$

where $LPF(z)$ is the loop filter transfer function, AM is the amplitude modulation data (set to 1 in PLL mode), and T_{REF} is the period of the input clock. The ADC quantization noise is low-pass filtered and compressed by G_{PD} . A high G_{PD} value therefore ensures no degradation in-band. The ADC quantization not being the performance bottleneck can be motivated by pointing to PLLs with high performance even in bang-bang phase-error detection mode [Tasca11] where only a single bit phase-error data converter (instead of a multi-bit ADC) is used in a DTC-based PLL. In other words, the ADC only quantizes the DTC error residue, which is already *small enough*, hence the quantization levels are relaxed. Finally, a high G_{PD} from (Eq. 4.1) imposes no issues with loop stability since the digital low-pass filter can be easily adjusted (whereas an analog loop might necessitate a large capacitive filter).

The last quantizer in the digitized subsampling loop that potentially degrades performance with respect to a purely analog subsampling PLL implementation is the DAC driving a varactor in the VCO, effectively implementing a DCO (see model in Fig. 4.1). Note that the DCO output phase quantization noise is high-passed in transfer to the output. In this design, the DAC is implemented with a $\Delta\Sigma$ modulator (clocking at the reference rate) in front of it, to randomize and shape the quantized residue. Thanks to the fact that we make use of a DAC driving a varactor (to achieve digital frequency control), it is possible to implement a second analog pole in the system. This is done by a simple RC filter at the DAC output that suppresses the high-passed DAC quantization noise (see Fig. 4.4 and the basic PLL loop). The filter thus ensures an additional 20 dB/decade noise roll-off after the second pole. Finally, the DAC quantization noise can be kept *low enough* in this environment (see Sect. 4.2.3) since a varactor can normally be sized to low sensitivity.

4.2.2 Phase/Frequency and Amplitude Modulation

The system makes use of the two-point injection scheme so that the phase/frequency modulation bandwidth is not limited by the PLL cut-off [Markulic16a, Marzin14]. This allows data independent loop filtering profile setting that ensures optimal phase noise management with respect to the low-passed (DTC, reference) and high-passed (VCO) noise contributors. Namely, the frequency modulation data $\Delta\omega_{\text{mod}}$ sent to the TX output through the DAC–VCO–DPA chain is cancelled by the DTC (Φ_{DTC}) in the phase-error detector (sampler) so that the loop remains undisturbed, while the data travels to the output unaltered. This is achieved by sending the same modulation data (with opposite sign), expressed as accumulated phase, to the DTC—which then re-adjusts the delay for zero-crossing sampling. Due to this, the small signal approximation for PD gain still holds, as given by Eq. (4.1). In frequency domain, the DAC injection path implements a high-pass, while the DTC injection path implements a low pass to the output—which when added together ensure an all-pass phase/frequency data transfer [Marzin12].

The DPA is placed within the PLL, after the VCO. An ideal DPA does not induce any nonidealities in the phase signal propagation path. This is relevant for the subsampling process which remains unaltered with respect to an implementation where the VCO output is subsampled (Fig. 4.2a), as in a subsampling PLL-based phase/frequency modulator (see Chap. 2). On the other hand, the DPA modulates amplitude of the output sinewave, as desired by the polar modulation process. The amplitude modulation signal from Fig. 4.1 can therefore take any instantaneous value from the rational number range $<0,1]$ representing amplitude levels at the TX output between zero power (0) and full power (1), with appropriate number of quantized levels in between. The DPA amplitude change over time could induce phase-error detection gain (G_{PD}) fluctuations (see Fig. 4.2 which depicts the effect for some two AM codes X and Y) and a consequential PLL bandwidth modulation. To avoid this, the amplitude modulation data are also used in scaling of the ADC output, through a simple digital division. One might describe this process as *two-point amplitude modulation* since (similarly to two-point phase/frequency modulation) the same data are injected in two places within the loop with a cancelling effect. The goal is to transfer the AM data to the output without any data bandwidth limitation or any disturbance of the loop.

The AM data changes over time in a pseudo-random fashion. Even though the loop dynamics remain constant (two-point AM injection), the ADC quantization noise and the sampler noise filtering are influenced by AM. In the noise transfer to the output, as described by Eq. (4.2), the AM value can be modeled with the amplitude modulation signal's *mean*. This value depends on the desired constellation and related Peak-to-Average Power Ratio (PAPR). For example, in 1024 QAM with equal occurrence probability of every symbol, the amplitude modulation data distribute at the range $<0,1]$ as shown in Fig. 4.3. The mean of distributed values (0.44 in this case) is then used in Eq. (4.2) to model influence of the appropriate noise generators and to re-adjust the loop parameters if necessary.

To ensure modulation quality guaranteed by the PLL phase noise floor (within the desired modulation BW), the quantization error induced by the modulating DAC and

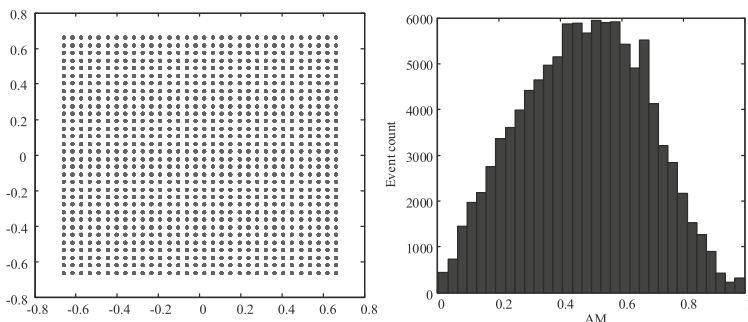


Fig. 4.3 Generic 1024 QAM constellation (left) and distribution of used amplitudes with uniform occurrence probability of any symbol (right)

the amplitude modulator (DPA) should be *small enough*. The total integrated phase noise (limited by environmental noise) translates then directly into the measured EVM.

4.2.3 Prototype Targets and Building Block's Specifications

A detailed overview of the system's block diagram is depicted in Fig. 4.4. This prototype's target is a visible 1024 QAM constellation at 2.5 MHz bandwidth (with possible extension to 10 MHz) around a 5.5 GHz carrier. Note that the far-out noise is not optimized in this design, we concentrate exclusively on the in-band accuracy expressed in terms of EVM.¹

To reach the targets we reuse the 10-bit DTC with 0.5 ps/LSB from Sect. 2.4.2, ensuring that the integrated DTC quantization noise remains at least 20 dB below the total integrated phase noise of the PLL.

We opt for a 6bit ADC with $K_{\text{ADC}} = 1.5 \text{ LSB/mV}$ to cover the complete range of possible phase errors without clipping at peak output amplitudes of $A_{\text{DPA}} = 1.25 \text{ V}$ (12 dBm). A signal amplifier with a gain of A compensates for sampler loss.

The VCO operates at central frequency of 11 GHz (with targeted 20% tuning range) and is divided by 2–5.5 GHz before it drives the DPA. This feature enables harmonic rejection mixing in the DPA as discussed later in Sect. 4.5.2. During phase modulation, the output phase may necessitate shifts as large as $\pm\pi$ radians in a single clock cycle, to reach any desired point in the constellation. The largest corresponding instantaneous carrier frequency shift can then be extracted from: $\Delta\omega_{\text{mod}} \cdot T_{\text{REF}} = \pm\pi$, i.e., $\Delta f_{\text{mod}} = \pm 1/(2T_{\text{REF}})$ (note that the DAC is clocked at the reference rate in this system). The DAC produces aliases around harmonics of F_{REF} , hence the largest modulation bandwidth allowed by the Nyquist criterion is $F_{\text{REF}}/2$. Here we restrict the PM bandwidth to $F_{\text{REF}}/4$. The DAC-VCO combinations provide a 50 kHz/LSB frequency step (around 11 GHz carrier), with 11-bits, covering in total a range of $\pm 25.6 \text{ MHz}$ after VCO's output division by 2. The Least Significant Bit (LSB) weight ensures that the integrated quantization noise remains substantially below the total Integrated Phase Noise (IPN) of the PLL. The amplitude modulator (DPA) operates with 10 bit resolution.

Note that a subsampling loop cannot discriminate between frequencies that are $N \cdot F_{\text{REF}}$ apart from each other. To ensure accurate frequency acquisition, this PLL makes use of a Phase/Frequency Detection (PFD) mechanism which runs in parallel to the subsampling loop, with a predefined dead-zone [Gao09b]. In the presence of frequency difference between the desired and current PLL output, the PFD loop produces appropriate UP/DOWN pulses (after enough phase error is accumulated) that are scaled and then integrated in the loop filter (see Fig. 4.4). After the correct

¹An obvious way to improve the far-out noise is to use a higher clock speed for the phase modulating DAC, more than the 40 MHz crystal oscillator reference used now.

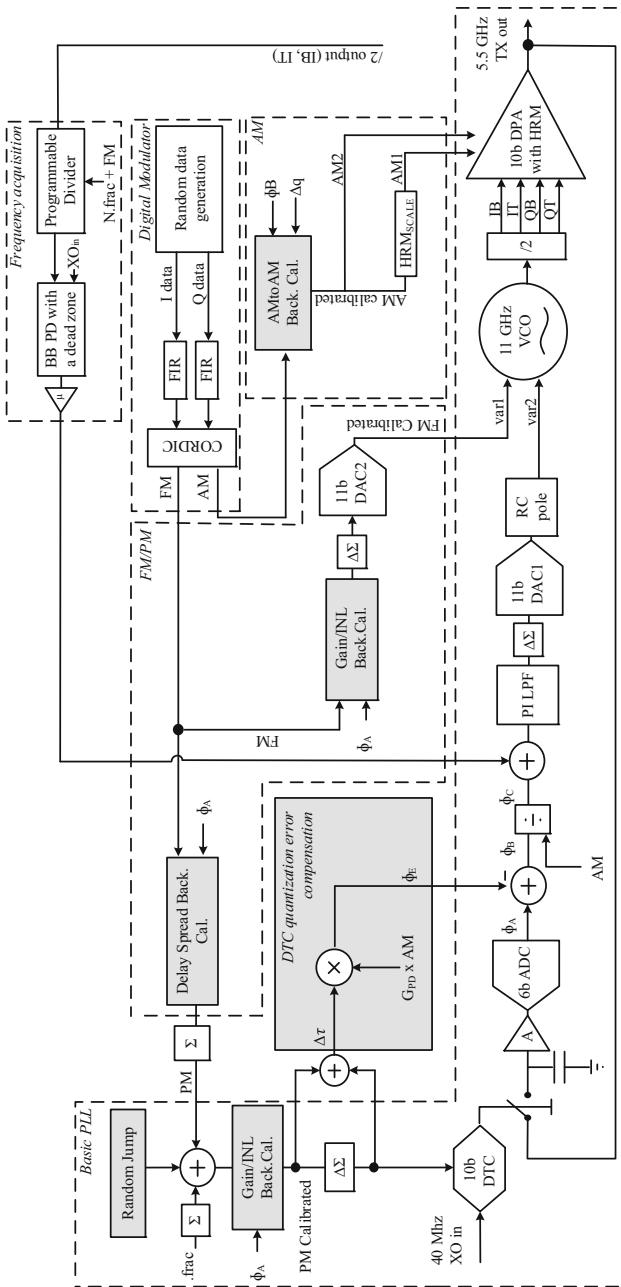


Fig. 4.4 Detailed overview of the SSPTX. Gray blocks serve for background calibration

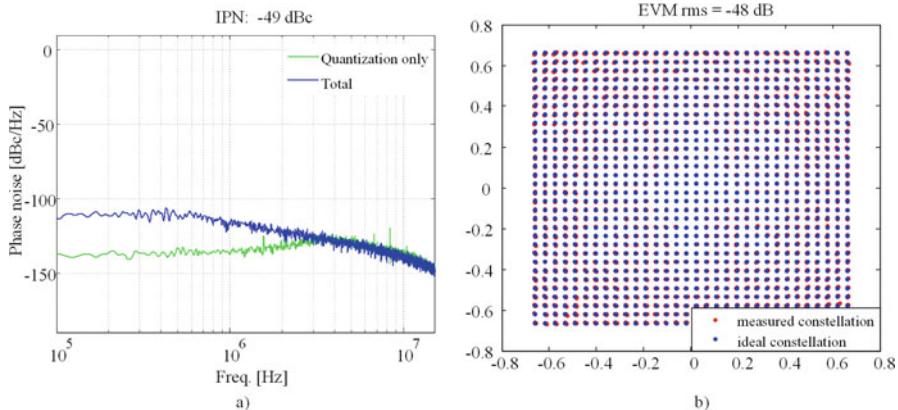


Fig. 4.5 (a) Simulated PLL performance—quantization noise only (green) and all noise sources included (blue) at a fractional 5 GHz output ($F_{ref} = 40$ MHz). (b) Simulated output constellation of a linear system as depicted in Fig. 4.4 at 2.5 MHz modulation bandwidth and 1024 QAM with all modeled noise sources enabled

frequency had been acquired (or is close enough to the solution), the subsampling loop takes over and the PFD loop is disabled for power savings.

Figure 4.5a shows simulated spectral profile of the proposed PLL during synthesis of a fractional channel around 5 GHz. Results indicated in green show a simulation where no thermal nor 1/f noise is present, i.e., the output is corrupted solely by quantization noise. The overall integrated quantization phase noise is <-60 dBc, which is well below the expected thermal (and flicker) contributions. Indicated in blue is the PLL output when all the modeled noise sources are included. The crystal oscillator phase noise floor is modeled with -165 dBc/Hz at 40 MHz while the simulated DTC noise floor is at -160 dBc/Hz. A sampling capacitor of 20 fF is assumed for modeling of the sampler's kT/C noise. The VCO noise models are fit to match measured VCO phase noise performance from [Hershberg14] (that is reused in this design). The target is to achieve <-47 dBc IPN around 5 GHz that ensures a visible 1024 QAM constellation. The simulated constellation of a fully linear transmitter is depicted in Fig. 4.5b), during 2.5 MHz bandwidth modulation with 1024 QAM, around a 5 GHz carrier. The 25 Mbit/s data transmission is mainly corrupted by phase noise, hence the measured EVM of -48 dB resembles to the level of the PLL's total integrated phase noise (the 1dB degradation comes from added AM/FM quantization errors).

4.3 Digital Linearization Techniques

The transmitter had to this point been considered ideal. Any nonidealities in the digital-to-transmitted output conversion will result in a disturbance that is sensed by the loop, either as a phase error, a G_{PD} gain (bandwidth) modification, or a

combination of these two. The disturbances at the scaled phase-error detector output can be correlated with the injected signal to identify and digitally null the distortion mechanisms.

Distortion is typically a consequence of analog nonlinearities (such as DAC/DTC/DPA INL), or cross-correlation/mismatch between two different injection points in the loop (such as AM-to-PM distortion or two-point injection mismatch). The following sections identify typical culprits for TX's output signal corruption and explain ways to mitigate them through self-linearization, during normal TX operation.

4.3.1 PM-to-PM Background Calibration

The SSPTX loop gets disturbed by colored noise in the presence of a gain-inaccurate/nonlinear DTC and/or modulating DAC, or similarly, in the presence of timing mismatch between the data injection in the two PM points. These effects were analyzed in Sect. 3.2 in a phase modulating analog subsampling PLL, where background calibration is proposed to mitigate the undesired behavior. Indicated mechanisms fall under the category of PM-to-PM distortion in context of the SSPTX (since they result in nonlinear translation of the ideal digital PM signal to the real and corrupted PM output). Techniques applied in Sect. 3.2 are similarly reused here in a digital loop.

As long as the phase modulation is linear, there is no correlation between the input DTC/DAC codes and the detected phase error (Φ_A in Fig. 4.4). This changes in the presence of gain-imbalanced or nonlinear DTC/DAC and the correlation can then be tracked and minimized through digital background calibration. The process is similar to the one described in Sect. 3.2, independently of AM in the SSPTX, since the Φ_A sign used for error tracking remains unbiased by potential PD gain (G_{PD}) fluctuations.

The principle of predistortion remains similar to the one described in Sect. 3.3.1, and is summarized here: if the code dependent error is predictable at any code, original input data may be modified to avoid making the error in the first place. The modified digital input then induces the desired analog effect [Markulic16b, Levantino14]. The correction coefficients used per input code may be stored in a LUT. To avoid making an excessively large LUT with 2^N coefficients (where N is the data converter's number of bits) a limited set of correction coefficients is instead spread across the full input dynamic range, with suitable approximation in between [Markulic16b].

In this design we recognize that the predistortion can be optimized with respect to the particular DAC architecture. For example, to reduce potentially large INL/DNL excess, designers often opt for segmented arrays that offer a good compromise between area/routing efficient binary and matching/linearity superior thermometric implementations [VDB13]. In segmented DACs, the transfer curve often exhibits specific INL shape that has “saw-tooth” alike behavior, since the systematic

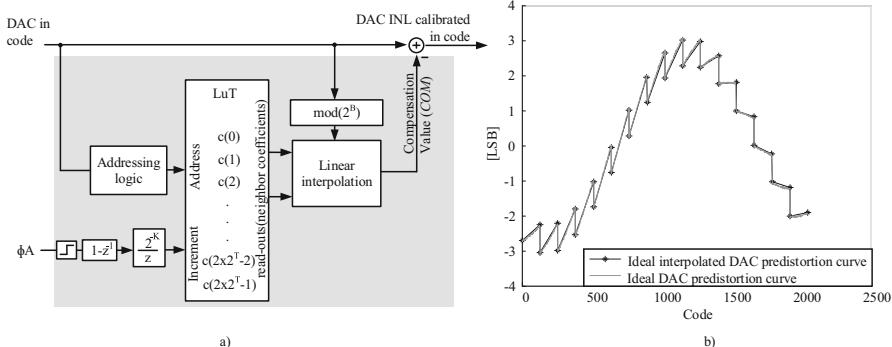


Fig. 4.6 DAC INL background calibration algorithm (T and B stand for the DAC's number of thermometrically and binary coded bits, respectively): (a) implementation; (b) ideal (gray) and ideal interpolated (black) INL predistortion curve for $B = 4$, and $T = 7$

nonlinearity (of an, e.g., varactor in this design) combines with the largest DNL errors at the transition between the segmented parts (of an, e.g., resistive DAC array in this design). This behavior is exemplified in Fig. 4.6a for a 11 bit modulating DAC with 4 bit thermometric (T) and 7 bit binary (B) segmentation, which is used here (see Sect. 4.5.1.4).

In contrast to the implementation from Sect. 3.3.1, the LUT correction coefficients are here grouped in pairs of two, to create neighboring correction segments that match the used DAC segmentation (see Fig. 4.6b). With this correction coefficient grouping (instead of a simple piece-wise linear mapping across the full DAC range as in Sect. 3.3.1), the predistortion principle becomes more suitable for converter arrays with segmented sections. Assume a segmented array with T most significant bits (MSBs) implemented thermometrically and B least significant bits (LSBs) implemented in a binary switching scheme. The correction coefficients $c[0 : (2 \cdot 2^T - 1)]$ are then used at the input codes $n \cdot 2^B$ and $[(n + 1) \cdot 2^B - 1]$, where n is an integer from $[0, 2^T - 1]$. For the input codes that fall in-between the ones represented by the coefficients c , a linear interpolation between the nearest pair is used. The compensation values (COM) subtracted from the original input code, in , are then calculated as:

$$COM = \left\{ c\left(\left\lceil \frac{in}{2^B} \right\rceil\right) - c\left(\left\lfloor \frac{in}{2^B} \right\rfloor\right) \right\} \cdot \frac{in \bmod 2^B}{2^B} + c\left(\left\lfloor \frac{in}{2^B} \right\rfloor\right), \quad (4.4)$$

for every clock cycle. In practice this means that 2 neighboring coefficients predistort for binary segments with piece-wise linear spread. The mathematical manipulations behind the computation of the COM value are easily implemented in the digital domain (multiplication in powers of two only need a shift operator). The final correction coefficient estimates are depicted by a red * symbol in Fig. 4.6 b for the presented 4-bit thermometric + 7-bit binary segmented DAC array. The straight

red line in-between the coefficients represents linearly approximated values used at the appropriate input codes (x-axes).

The LUT correction coefficients are updated in the background. At every clock cycle, the measured phase error Φ_A sign is differentiated (data injected in the frequency domain, but measured in the phase domain [Markulic16b]), scaled (by 2^{-K}) and then accumulated to the appropriate LUT address, determined by the used input code. As the algorithm runs in the background, the coefficients slowly change and then settle to a value which accurately cancels the nonlinearity. The measured phase error Φ_A then becomes a zero-mean stream per certain input code, and the correction curve is stabilized as depicted in Fig. 4.6b, in red. The calibration speed is determined by K , i.e., by the gain with which the sign of the detected phase error is scaled before accumulation at the desired address. Note that increasing the calibration speed normally leads to more noise in the correction process. Visualization of the background correction coefficient estimation process can be found in Sect. 4.3.3.

The background INL correction is similarly implemented for the DTC (see Sect. 3.2.4), except that the LUT addressing is driven by DTC input codes (independently of DAC input). Note also that there is no differentiation of the detected error and that there is no need for segmented correction as for the DAC since the DTC was optimized for 10-bit matched capacitor array (systematic DTC nonlinearity is not dominated by mismatch). The final PM-to-PM distortion mechanism, the two-point error injection mismatch (or delay spread) is background cancelled as in [Marzin12]. Note that the PM-to-PM correction algorithms can independently run in the background since they cancel orthogonal problems (as discussed in Sect. 3.2).

4.3.2 *AM-to-AM Distortion Background Calibration*

4.3.2.1 *Background G_{PD} Estimation*

In this section we explain how the DTC produced quantization errors can be used to estimate the exact analog value of G_{PD} (that is dependent on the output RF amplitude, as explained in Sect. 4.2). The DTC quantization errors can be subtracted digitally from the ADC output so that they do not propagate through the loop, since they are predictable and since they contain no useful phase-error information (note the DTC quantization error compensation block in Fig. 4.4). This process can be executed with high accuracy with respect to the DTC (quantization error estimate) since its gain is known and background calibrated (see Sect. 4.3.1). The absolute accuracy of phase-error detection gain (G_{PD}) is, however, difficult to predict at design time. Note that for now we assume linear *AM* scaling from <0:1], or 0 to 100% of output amplitude.

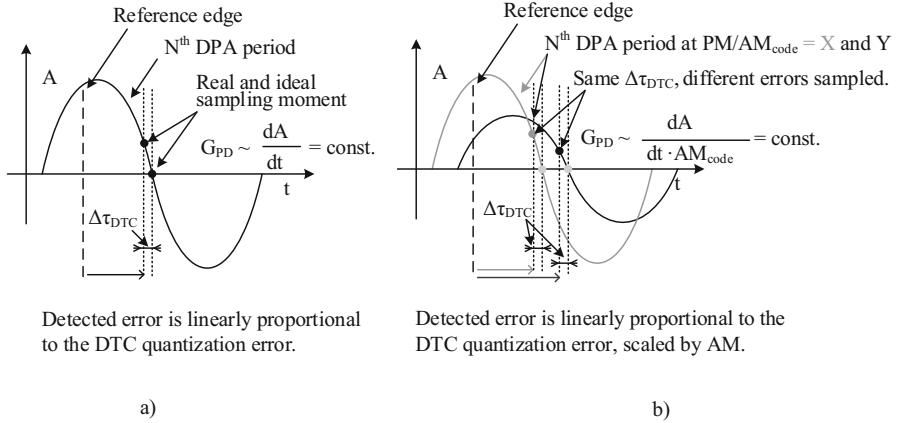


Fig. 4.7 Ideal and real sampling event in the presence of DTC quantization error in (a) PLL mode and (b) SSPTX mode at two different AM codes (A and B). G_{PD} is proportional to the slope around the zero-crossing (i.e., sinewave amplitude). Note that the same time-quantization excess leads to different sampled amplitude during AM

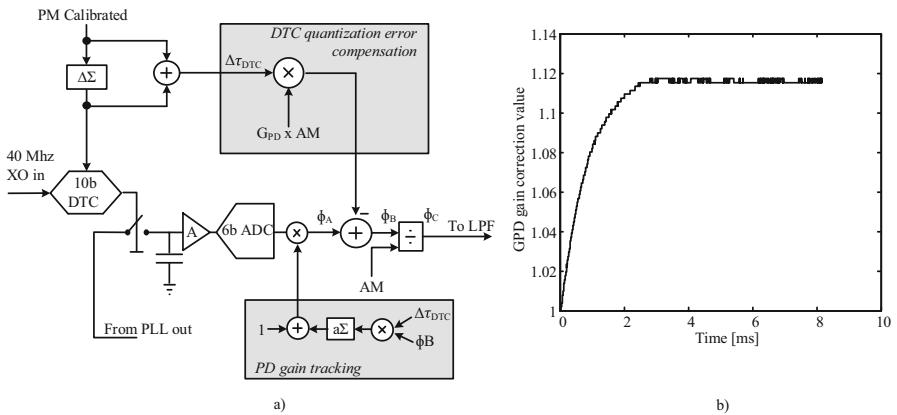


Fig. 4.8 Phase-error detection gain background calibration: (a) block diagram implementation indicated in gray within the SSPTX; (b) background calibration with 10% gain error

During amplitude modulation, as depicted in Fig. 4.7, the same amount of time error, or DTC quantization excess in this context, leads to a different sampled voltage error, since the DPA output operates at different amplitudes and the PD detection slope changes around the ideal zero-crossing. For accurate compensation in the digital domain, it is hence necessary to scale the original time deviation (DTC quantization error) by the appropriate instantaneous AM code and the original G_{PD} value (see Fig. 4.8a).

The detected phase error, Φ_B , under the influence of DTC quantization error $\Delta\tau$ can be expressed as:

$$\Phi_B = \Delta\tau 2\pi f N \frac{1}{AM(1 + G_{OL})} (G_{PDreal} \cdot AM_{real} - G_{PD} \cdot AM), \quad (4.5)$$

where $G_{PDreal} \cdot AM_{real}$ represent the analog phase-error detection gain and AM amplitude scaling, respectively, that may differ from the expected digital values $G_{PD} \cdot AM$. f is the operating frequency of the DPA and G_{OL} is the PLL open loop gain defined with Eq. (4.3). Note that expression 4.5 represents a small signal approximation valid when the PLL locks, and when the DTC quantization errors are *small enough*. To minimize Φ_B , $G_{PDreal} \cdot AM_{real}$ needs to correspond to $G_{PD} \cdot AM$. To address this issue, we propose self-estimation of the G_{PD} value in the digital domain, as depicted in Fig. 4.8. The goal is to minimize correlation between the quantization signal $\Delta\tau$ and the phase-error detector output Φ_B . If there is no gain error, the G_{PD} value is simply scaled by the integer 1 and the integrator in the correction loop receives a zero-mean stream, since Φ_B is then not at all influenced by $\Delta\tau$. The scaling value changes if a gain error is present in the system. The integrator slowly drifts as in Fig. 4.8b (speed determined by the coefficient $a \ll 1$) in the positive or negative direction, depending on the gain error sign, and then finally saturates at the desired position. As the gain error disappears, the input to the integrator becomes a zero-mean stream again. In other words, the PD gain is digitally scaled to fit the G_{PDreal} value. The digital estimate can later be used for optimal scaling of the loop filtering properties. The algorithm can run in the background continuously tracking potential environmental changes, or it can be disabled once when the correct estimate is acquired.

Note that different training sequences could be used instead of the DTC quantization errors to achieve the same result. For example, larger timing errors could be on purpose induced in the system to produce *more information*, i.e., larger detected error amplitudes. This principle is useful at lower AM codes where the information is limited by the given ADC resolution. Care must be taken to still approximately operate in the linear range of the PD. Simulations show that by application of the first order $\Delta\Sigma$ modulator in front of the DTC, as depicted in Fig. 4.4, enough information is produced for the G_{PD} calibration algorithm to consistently converge at all the AM codes used in 1024 QAM. A second order $\Delta\Sigma$ modulator that has a larger output swing can be optionally used to speed up the calibration and produce more information at low AM codes.

4.3.2.2 Background DPA INL Cancellation

The described G_{PD} background calibration successfully resolves potential gain mismatch, but the same principle can also be extended for compensation of nonlinear AM_{real} scaling, or in other words, for background predistortion of a nonlinear DPA. Figure 4.9a exemplifies a transfer function of a nonlinear DPA

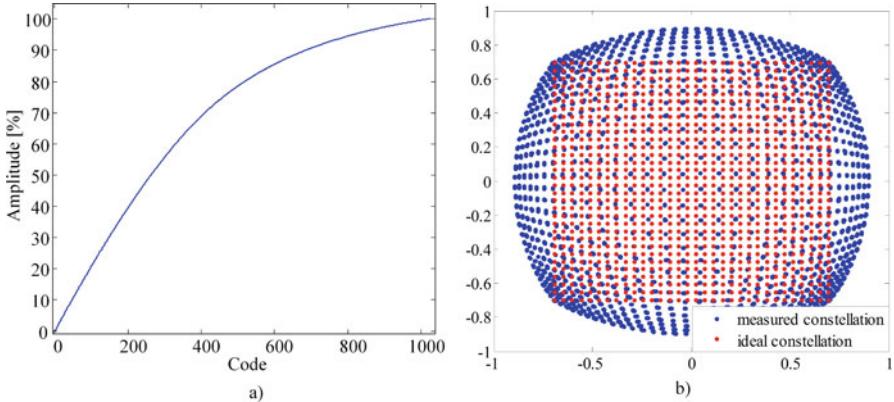


Fig. 4.9 (a) Simulated transfer curve of a 10 bit class D⁻¹ DPA; (b) Constellation diagram of a transmitter using the nonlinear DPA

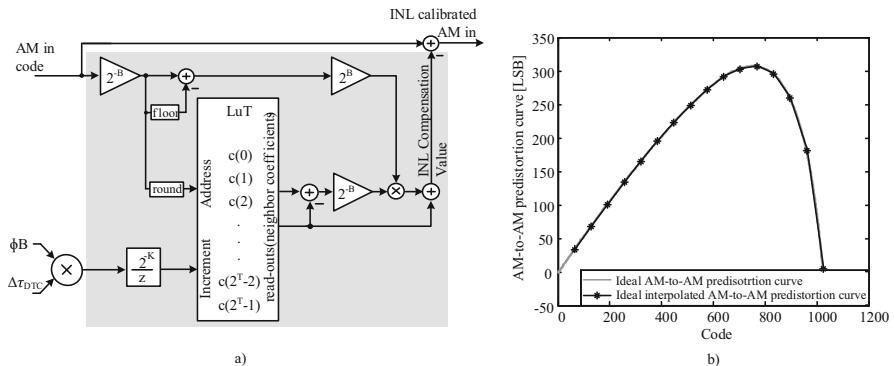


Fig. 4.10 (a) DPA INL background calibration algorithm implementation (T = thermometrically coded bits, B = binary coded bits); (b) Ideal and interpolated AM-to-AM predistortion curve

(simulated 10-bit class D⁻¹ DPA model similar to [Chowdhury12]). Constellation diagram of a transmitter employing the DPA for amplitude modulation is clearly distorted, as shown in Fig. 4.9b. The constellation points do not appear in circles linearly distanced from the constellation center, which consequentially imposes heavy degradation of the measured EVM. Instead of pushing effort into analog linearization that typically comes with an efficiency, complexity, and area expense, we make use of digital background calibration, as explained next.

Predistortion of the 10-bit DPA proposed in this design is implemented through a 16 position LUT, as depicted in Fig. 4.10 (note that the DPA uses a 4-bit thermometric and 6-bit binary segmentation, hence $T = 4$ and $B = 6$). Instead of sending the original input code to the DPA, predistorted inputs are used, which force accurate digital to amplitude conversion. In other words, if the DPA amplitude error is known at every code, it is always possible to choose a digital input that

results with the desired output amplitude. The 16 correction coefficients are simply linearly spread across the full input range and they represent values that need to be subtracted from the original input to achieve the desired DPA response.

The central phenomena used for populating the LUT (i.e., for implementation of the background calibration process) are the correlation between Φ_B and $\Delta\tau$ signals per particular AM input code (Fig. 4.10). The Φ_B signal represents the detected phase error, ideally perfectly free of DTC quantization error ($\Delta\tau$) influence, as explained in the previous section. In time domain, this means that the Φ_B signal, as shown by (4.5), remains a zero-mean pseudo-random stream of data. If, however, AM scaling is nonlinear, this stops being true across certain amplitudes. In other words, the Φ_B and $\Delta\tau$ signals correlate at AM inputs that induce nonlinear amplitude scaling. To track this, the product of Φ_B and $\Delta\tau$ is at every cycle accumulated to the appropriate AM-determined address in the LUT, as shown in Fig. 4.10. The correction coefficients therefore change over time (speed determined by the tap gain factor 2^{-K}). They propagate towards the point where the nonlinear AM scaling becomes accurately compensated. The values within the LUT then saturate, and the Φ_B and $\Delta\tau$ signal product becomes a zero-mean stream, per particular input DPA code. The final AM-to-AM predistortion estimation curve ideally settles to the values indicated in Fig. 4.10 by the red * symbol. Piecewise linear approximation is used at the codes in-between. Visualization of the background calibration process can be found in Sect. 4.3.3. The algorithm can be enabled in the background, as the transmitter operates normally, since it does not interfere with its operation. This correction loop compensates gain mismatch as well, hence it does not operate in parallel to the PD gain error estimation. In other words, the DPA INL calibration loop gets enabled after the G_{PD} estimation settles (and freezes the gain estimation).

4.3.3 Phase-Domain Matlab Simulations of Background Calibration

Figure 4.11 shows full system simulation with all the background calibration algorithms enabled simultaneously. The DTC, DAC, and DPA nonlinearities are modeled as described in the section above. The algorithms run independently and do not influence each other since the nonlinearities that are being cancelled produce disturbances that can be correlated to separate/independent digital inputs. In other words, there is a single optimal position to which all the loops converge to maximize the desired error suppression. The correction coefficients are continuously updated in the background and, as visible from the figure, at about 150 ms of operation time they settle to the desired positions, predicting accurately the necessary predistortion profile. The simulation is on purpose prolonged for the algorithms to run in the background, tracking potential environmental change. Since there is no modeled

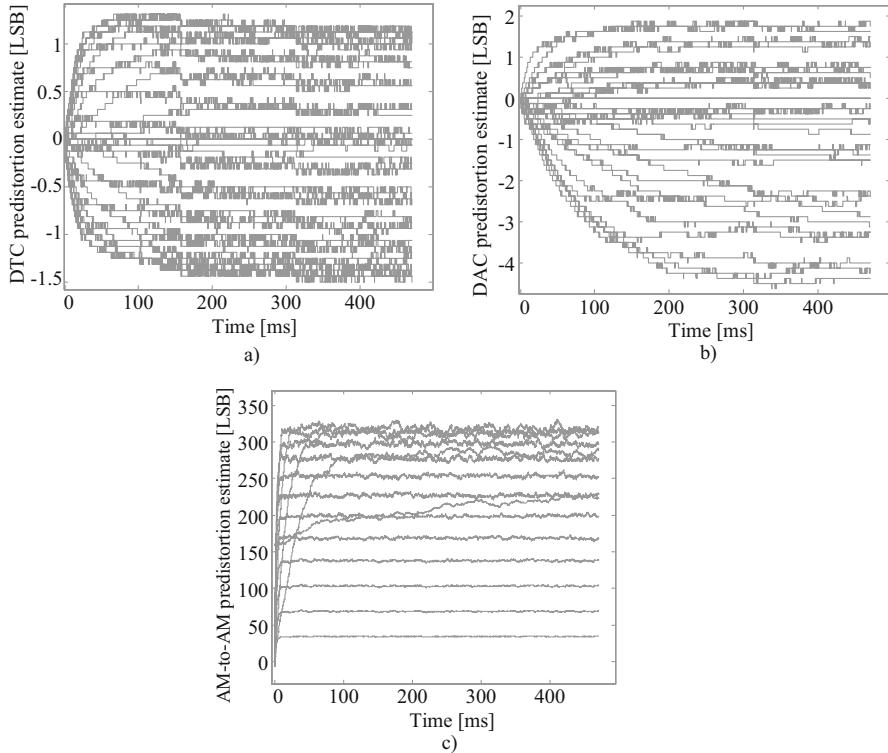


Fig. 4.11 Background SSPTX calibration. (a) DTC INL background predistortion estimation. (b) DAC INL background predistortion estimation. (c) DPA INL (AM to AM) background predistortion estimation

distortion fluctuation over time, the correction coefficients simply dither around the optimal position. If the estimation gain is low, the added noise in the TX path is negligible.

Figure 4.12a shows output constellation before the background calibration is enabled. The output is clearly heavily distorted both in amplitude and in phase domain. After the calibration is enabled, the constellation is linearized as the correction algorithms formulate optimal predistortion curves (Fig. 4.12b).

4.4 Built-in AM-to-PM Distortion Filtering

Besides that it enables background AM-to-AM distortion calibration, there exists a second benefit of placing the amplitude modulating block within the PLL. An important disturbance mechanism often encountered in transmitters is AM-to-PM distortion [Yoo11, Chowdhury11]. In the moderate/high output power range,

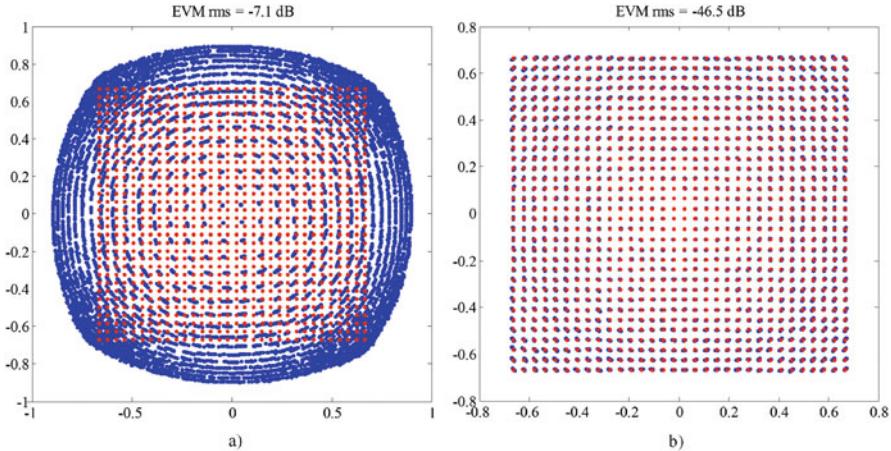


Fig. 4.12 SSPTX **(a)** before and **(b)** after the background calibration is enabled

this effect may easily lead to spectral leakage and EVM degradation. Amplitude dependent power consumption, amplitude dependent DPA load characteristics, or amplitude dependent magnetic coupling between the VCO and DPA are just some of the typical culprits for the unwanted outcome.

The problem of AM-to-PM distortion is typically neutralized through predistortion [Yoo11, Chowdhury11]. Measurement-based predistortion can be time consuming and costly, especially when covering PVT variations. Moreover, predistortion becomes increasingly complex in the presence of memory effects, i.e., when AM-to-PM leakage is not purely static but depends on an arbitrary number of previous AM samples (at some arbitrary weight).

In the SSPTX, the DPA-produced phase deviations are high-pass filtered in transfer to the output, in the same fashion as for the VCO originated phase noise. This system's property comes directly as a consequence of placing the DPA in series to the VCO, where the transfer function of the phase signal to the output behaves as depicted and numerically expressed in Fig. 4.13a. The suppression of the unwanted AM induced phase deviation signal Φ_{DPA} is frequency dependent. Low frequency AM-to-PM distortion is strongly suppressed. As the frequency of Φ_{DPA} increases, the suppression rolls off with 20 dB per decade, until it disappears at frequencies above the PLL cut-off.

Frequency content of Φ_{DPA} depends on the AM excitation signal and on the distortion mechanism properties (see Fig. 4.13b). To verify the performance of distortion suppression in a SSPTX, we introduce a static AM-to-PM distortion curve (simulated class D⁻¹ DPA) in the Matlab simulation environment as given in Fig. 4.14a. Note that this is the simplest form of distortion that excludes any memory effects in the system. The PLL loop bandwidth is set to approximately 5 MHz, while the PLL modulation BW remains at 2.5 MHz. As predicted, SSPTX efficiently suppresses undesired behavior, achieving approximately 3 dB better result in terms

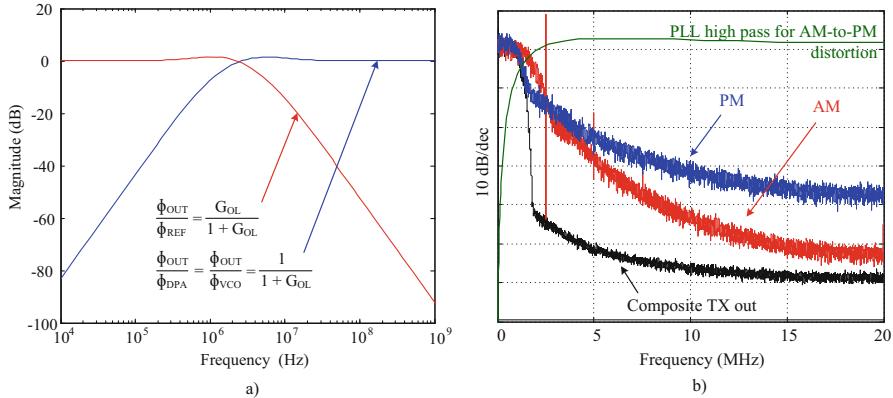


Fig. 4.13 DPA induced phase deviations are high-pass filtered in transfer to the output of the SSPTX, similarly as the VCO phase noise. (a) PLL filtering profiles. PLL BW set to 2.5 MHz in this example. (b) TX signal composition and high-pass filtering profile of AM-to-PM distortion in linear scale

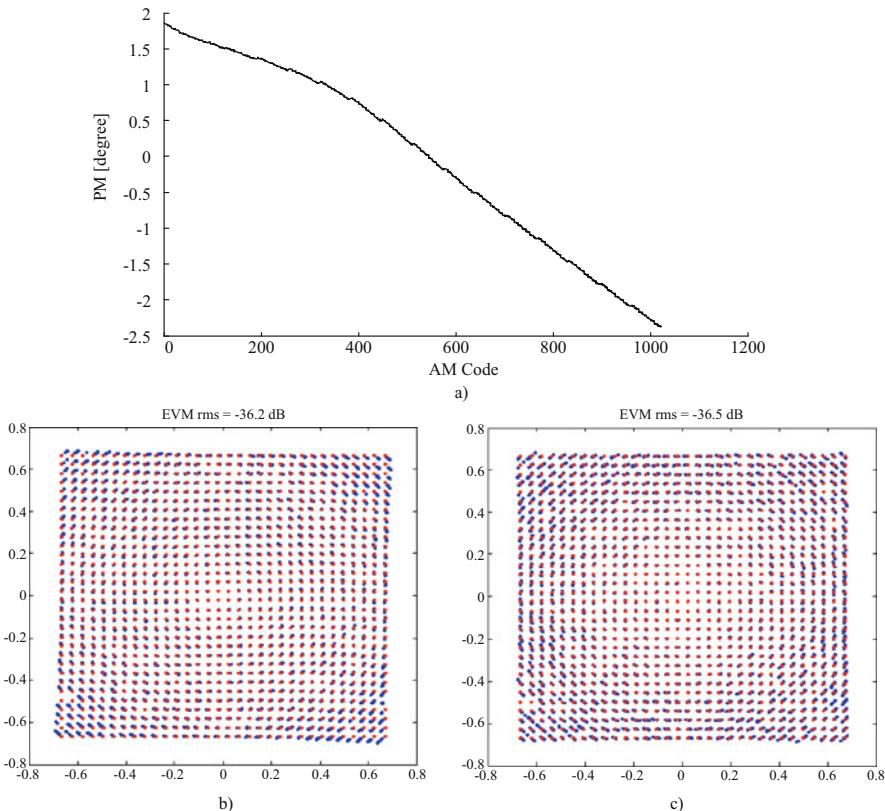


Fig. 4.14 AM-to-PM suppression in a SSPTX in comparison to a typical polar TX with AM modulator out of the PLL. (a) Modeled static AM-to-PM distortion. (b) EVM in a Polar TX with AM to PM out of the PLL. (c) EVM in a SSPTX

of EVM than a typical polar TX with the amplitude modulating block out of the loop—the AM induced phase rotation from Fig. 4.14b is significantly reduced in the SSPTX (Fig. 4.14c).

4.5 Analog Building Blocks

4.5.1 Subsampling Path: From Sampler to Code

A simplified overview of the SSPTX subsampling path is depicted in Fig. 4.15a. Phase-error detection is performed directly at the DPA output, where the signal can be represented by a (constant amplitude and frequency) sinewave, between the samples. The timing diagram of the phase-error detection and data handling process is summarized in Fig. 4.15b.

The DTC [Markulic14] delays the sampling event with respect to the crystal oscillator input rising edge. The minimal delay imposed by the DTC is approximately 500 ps, while the highest value in its output dynamic range reaches approximately 1 ns. After a sample is acquired, the *sampler done* signal goes high and triggers the dynamic amplification of the sampled signal. The amplifier uses 2 ns to provide a valid input for the ADC. The timing window for the analog-to-digital conversion is 5 ns wide. The ADC resolves 6 bits in this time frame. After the *ADC done* pulse, the phase-error detection process is reset and the digital subsystem starts phase-error information processing. There is 4.5 ns dedicated for the digital filter to

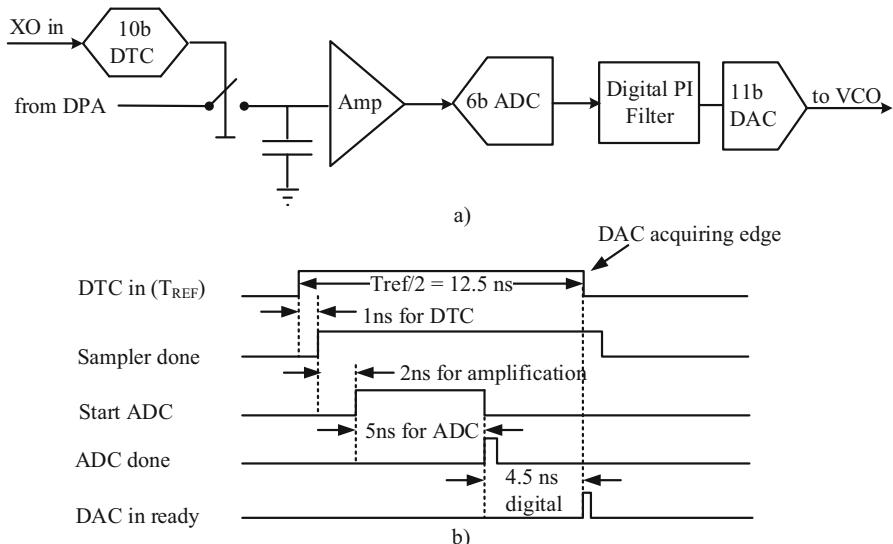


Fig. 4.15 SSPTX (a) subsampling path block diagram and (b) timing diagram

provide new digital output in front of the falling edge of the clock, at which the transmitter acquires new values for the DACs.

4.5.1.1 Sampler

The sampler operates at the high-power DPA output, hence it is important to take care of its reliability, linearity (signal independence), and potential sensitivity to supply disturbances, not to introduce any distortion in the transmitted signal. The peak DPA output power target is 12 dBm which translates to max 2.5 V peak-to-peak over a $50\ \Omega$ load.

The sampler schematic is presented in Fig. 4.16. The sampler is directly connected to the DPA output, differentially at both sides with a twin dummy structure to equalize loading of the DPA and make it independent of the sampling phase [Gao10]. The structure operates in two phases. In the track phase M2 switch is closed, and M1 switch is opened. Capacitors C1 (DC blocking) and C2 and the resistors embody an attenuating network that divides the input signal swing by an approximate factor of 12 as it propagates to the top plate of C2. The attenuation is strong enough to ensure linear tracking, independent of AM signal, even for the largest DPA output power. Moreover, the same mechanism ensures only a small portion of the DPA output current flows through the structure (avoiding potential large signal problems). Note that the sampler should not introduce any

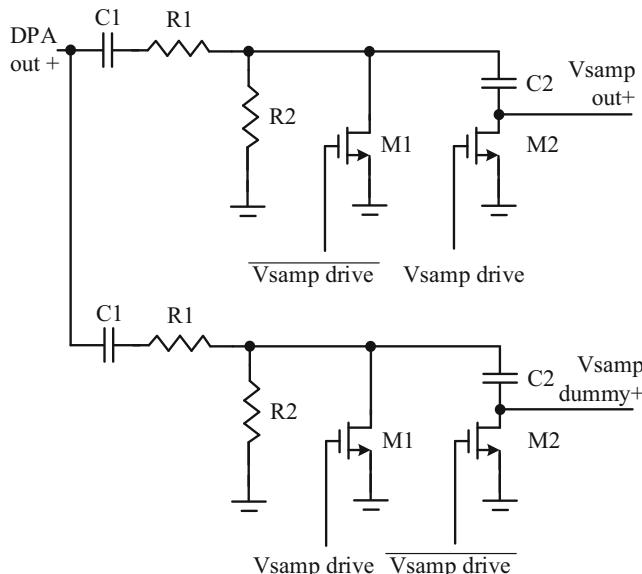


Fig. 4.16 Sampler schematic. Note that the figure shows only single path sampling (with a dummy for load equalization) while in reality sampling is differential at both DPA sides

AM-to-AM, or AM-to-PM distortion in the track phase, i.e., signal tracking should be consistent at the given modulation BW. The parasitic amplitude/phase signal deviation introduced by the structure at the bandwidth of interest is insignificant in post PEX verification. In hold phase, the M2 switch from Fig. 4.16 opens, storing charge in the C2 capacitor. To avoid large amplitude swing across C2 that can potentially induce charge leakage through M2, we make use of switch M1 that closes in this phase. This architecture enables pseudo-differential sampling, in which charge leakage cannot be dismissed as a common mode signal. Top plate of C2 remains grounded and experiences very little swing. The sampled voltage is then readily used from the bottom plate of C2 for the input stage of the dynamic amplifier, in front of the ADC.

Although not under investigation of this work, direct sampling at the DPA output might result with down-conversion of the environmental noise picked up by the antenna (for example, large neighboring blockers). This problem can be addressed with careful matching network design in the future.

4.5.1.2 Signal Amplifier

The amplifier schematic is depicted in Fig. 4.17. In the sampler tracking phase (see Sect. 4.5.1.1) input is set to common mode ground, biasing a pMOS differential pair. This is the reset phase for the amplifier, when the Φ_{reset} is high (switches closed) and Φ_{amplify} signal is low (switches open). The amplifier output is set by V_{GS} of M2 which is designed to be around $VDD/2$, as desired by the comparator of the Successive Approximation Register (SAR) ADC in the following stage.

In the hold phase of the sampler, Φ_{reset} switches are opened and the amplifier moves into amplification mode. The differentially sampled voltage ($V_{\text{samp A}} - V_{\text{samp B}}$) is held at the amplifier's input and the Φ_{amplify} switches are closed. The

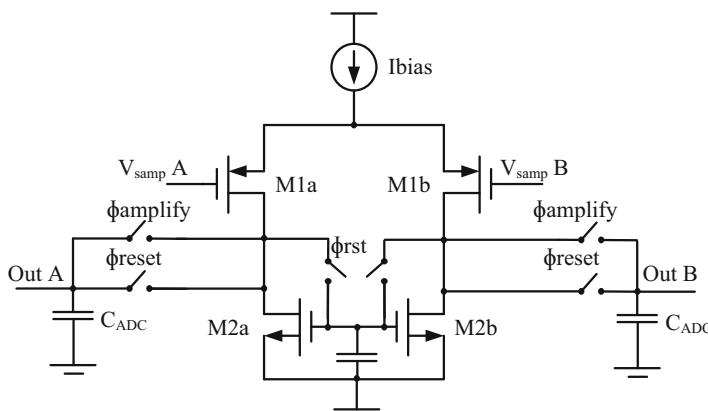


Fig. 4.17 Signal amplifier schematic

differential current now flows through the output capacitors and the amplifier gain is defined with:

$$A = \frac{g_{m,M1}}{C_{ADC}} \cdot \tau_{AMP}, \quad (4.6)$$

where $g_{m,M1}$ is the input differential's pair transconductance, C_{ADC} is the ADC input capacitance, and τ_{AMP} is the amplification time. The τ_{AMP} amplification window can be digitally tuned, to manipulate with the desired amplification factor. In the presence of mismatch, the amplifier can exhibit certain amount of offset that can be cancelled offline. This is achieved in the calibration mode where the input pair is shorted to ground and the transistors M2 (a and b) are digitally tuned to cancel the unwanted behavior. Monte Carlo analysis is used to confirm that the tuning range available can cover 3σ of statistically possible gain/offset variations.

Another important aspect is the amplifier linearity, as it ensures no degradation in phase-error-to-digital code conversion. Fortunately, as the PLL locks, the amplifier operates with low-amplitude inputs (zero-crossing sampling), where distortion is typically low. System level simulations show that the standard deviation of detected phase errors is $\sigma_t = 0.4$ ps, which translates to $\sigma_t = 1.4$ mV for the sampled input values (at 5.5 GHz with DPA output amplitude of 1.25 V and sampler attenuation of 12). Amplifier output voltage stays between $[-43$ mV:43 mV] during the nominal amplification by 8. This output maps to the $[-48$ mV:48 mV] ADC (6 bit 1.5 mV LSB) input range (see Sect. 4.5.1.3) with few millivolts of overhead to avoid clipping.² Simulations show that such an amplifier introduces below 0.5 LSB of INL in the phase-error detection path (assuming a 1.5 mV LSB ADC and aforementioned limited output swing), and does not impose any significant influence on the TX output.

4.5.1.3 ADC

The design of the ADC is based on the final stage of the 10 bit, 300 MS/s SAR ADC reported in [Malki16] to implement a 6 bit 1.5 mV/LSB ADC used here. The relaxed time constraint of 5 ns to resolve 6 bits enables use of only a single low-noise comparator from [Malki16]. The SAR ADC uses a DAC with differential feedback to ensure constant common mode input voltage during the conversion. The total input capacitance (including layout parasitics) seen by the amplifier is approximately 830 fF, and is used as such in Eq. (4.6) to calculate the amplifier gain. The DAC array is constructed from a unit capacitance of 0.5 fF that ensures 99% yield for 9.3 ENOB matching, which is more than enough to guarantee that ADC does not induce any distortion in the phase-error detection path of the SSPTX.

²Theoretically this clipping can still appear, however, it is statistically a very rare event that does not influence the average EVM.

4.5.1.4 DAC

The VCO contains two equivalent DAC driven analog varactors that are independently controlled by the loop filter and the digital modulator, respectively (see Fig. 4.4). We opt for DAC-driven-varactor oscillation control, rather than for a digital varactor bank since the former offers accurate step control (low LSB values), and simpler layout—the resistive DAC can be put out of the VCO core (that contains the varactor), while a digitally controlled bank necessitates normally integration within it. A negative consequence of this choice is a rather nonlinear digital-to-capacitance, i.e., digital-to-frequency conversion, typical for wide dynamic range varactors. This normally imposes no issue for the loop-filter-controlled varactor that operates with very low input swings (and is used exclusively for phase noise filtering). The modulating DAC will, however, distort the transmitted phase. The problem is in context of a SSPTX solved with digital background predistortion (see Sect. 4.3).

The DAC-based varactor control is depicted in Fig. 4.18. The 11-bit DAC architecture is segmented in 7 bit binary and 4 bit thermometric sections. The binary coded section is implemented in a classical R-2R ladder architecture, while the thermometric section uses simple resistive division (and a unit cell R). The switches connect the resistors (2R for binary and R for thermometric section) to either VDD or VSS, with respect to the input digital code. This creates an effective voltage division between VDD to VSS, and an analog output voltage that is linearity proportional to the input code. Unit resistor size R is $40\text{ k}\Omega$, and the cell's width/length are chosen for a statistical 11-bit matching.

Note that change of the DAC output voltage is not instantaneous. On the contrary, a R-2R DAC often produces (code dependent) output voltage glitches during switching, since the newly applied state (code) needs to propagate through the resistive network that is potentially rich in parasitics. The largest simulated settling

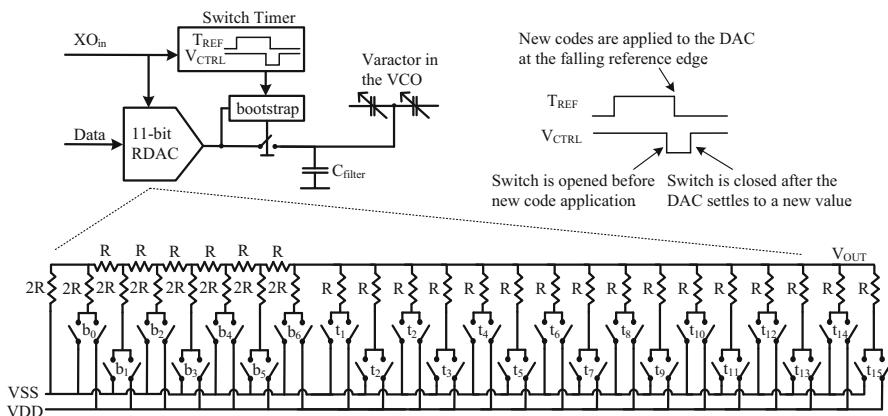


Fig. 4.18 DAC schematic and operation principle

time of the DAC is approximately 2 ns. To overcome this issue and not to allow any glitch energy dependent phase accumulation, we separate the DAC from the varactor at the switching instant. As depicted in Fig. 4.18, the switch opens and disconnects the DAC from the varactor before the new code is applied. During the switch opening time window, the DAC output settles. The switch is then again closed, and the new voltage is applied to the varactor input. To avoid code dependent settling time influenced by amplitude dependent switch resistance, we make use of bootstrapped switch control. Thanks to bootstrapping, the switch resistance is an order of magnitude smaller than the constant DAC output impedance ($R/15$) and independent of the input signal. The size of the C_{filter} capacitance from Fig. 4.18 is selected independently for the two banks. The loop filter driven varactor uses this capacitor in combination with the input impedance for implementation of the second PLL pole, and its size is approximately 8 pF. The phase modulating path, however, uses only a 1 pF capacitor, not to impose filtering on the transmitted data. Note that the non-instantaneous frequency shift of the VCO is compensated digitally with background delay spread cancellation as described in Sect. 3.3.2.

Note that the delta-sigma modulator in front of the DAC is not oversampled in this design, hence the analog high frequency pole is necessary for compression of the shaped quantization noise.

4.5.2 Inverse Class-D DPA with Harmonic Rejection Mixing (HRM)

The presented SSPTX incorporates an inverse class-D DPA (D^{-1}) [Chowdhury12]. This switching PA structure in its original form (depicted in Fig. 4.19) makes use of a single LC tank that is integrated in the output matching network (balun). For a maximized PA efficiency, this network ensures zero voltage switching where the LO

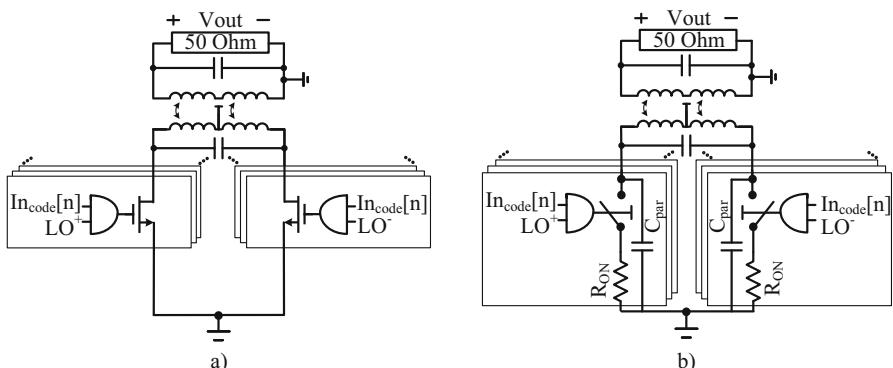


Fig. 4.19 Inverse class-D (D^{-1}) DPA. (a) Simplified schematic. (b) Model

driven switch is carefully closed only as the voltage across it equals zero. In other words, the square wave shaped current through the switch is set *out of phase* with the sinusoidal output voltage waveform. A theoretical 100% drain efficient structure assumes infinitesimally small switch impedance and, importantly, no output current (odd or even) harmonic leakage into the load [Chowdhury12]. The target of this design is to operate above 10 dBm peak power with a 0.9 V supply and a 50 Ohm load. For the proof-of-concept of the SSPTX, however, the DPA is not optimized for efficiency.

To enable amplitude modulation, the switch is changed into a digitally controlled parallel array. Note that linear code increase results with nonlinear resistance scaling (since switches are connected in parallel) and thus with nonlinear amplitude control. The unwanted behavior is resolved through background predistortion, as described in Sect. 4.3.2. Moreover, the structure's sensitivity to AM-to-PM cross-distortion (the nonideal switches operate under significantly different conditions between low and high amplitudes) gets attenuated by inherent SSPTX filtering as described in Sect. 4.4.

The AM-to-AM background calibration algorithm in this system relies on linear subsampling-based error extraction. Subsampling is unavoidably susceptible to down-folding of DPA's harmonical content. The matching network naturally filters the unwanted LO replicas, however, this process is quality factor dependent and thus imperfect.

The harmonics are likely to appear at the odd multiples of the fundamental tone, since the DPA is pseudo-differentially driven by a square wave. The subsampled fundamental sinewave shape is imperfect. Simulation-based study shows that as long as the harmonics scale linearly with the fundamental, that is, if the amplitude ratio and relative phase between the harmonics remain constants over AM codes, there is no influence on the SSPTX (besides a small G_{PD} change, which is nevertheless background tracked). A realistic, nonlinear structure, however, likely exhibits different AM-to-AM (and/or AM-to-PM) transfer mechanisms for the fundamental and the upper harmonics, respectively. The consequence can potentially be *false* AM-to-AM detection. Intuitively, if the harmonics are strong enough, the predistortion algorithm tracks linearity information about them (even though this information is not useful)—false AM-to-AM detection happens if this information differs from the information related to the fundamental tone. Fortunately, simulations show that the SSPTX output remains unbounded by these effects (no significant EVM degradation) as long as the ratio between the fundamental LO component and its harmonics is *large enough* (>30 dB).

The square wave driven switching pair of the typical class D⁻¹ DPA induces differential square wave current flow, as illustrated in Fig. 4.20a. The voltage waveform at the output contains odd harmonics of the fundamental tone. The DPA can optionally use harmonic-rejection-mixing (HRM) to avoid this [Weldon01, Lin14, Mensink05]. The basic idea is to modify the current waveform so that it does not contain the 3rd and the 5th harmonic. This can be achieved by mixing the 50% and 25% duty cycle rectangular drive (see Fig. 4.20b). The 25% duty cycle current opening window needs to fall exactly in the middle of the 50% duty cycle

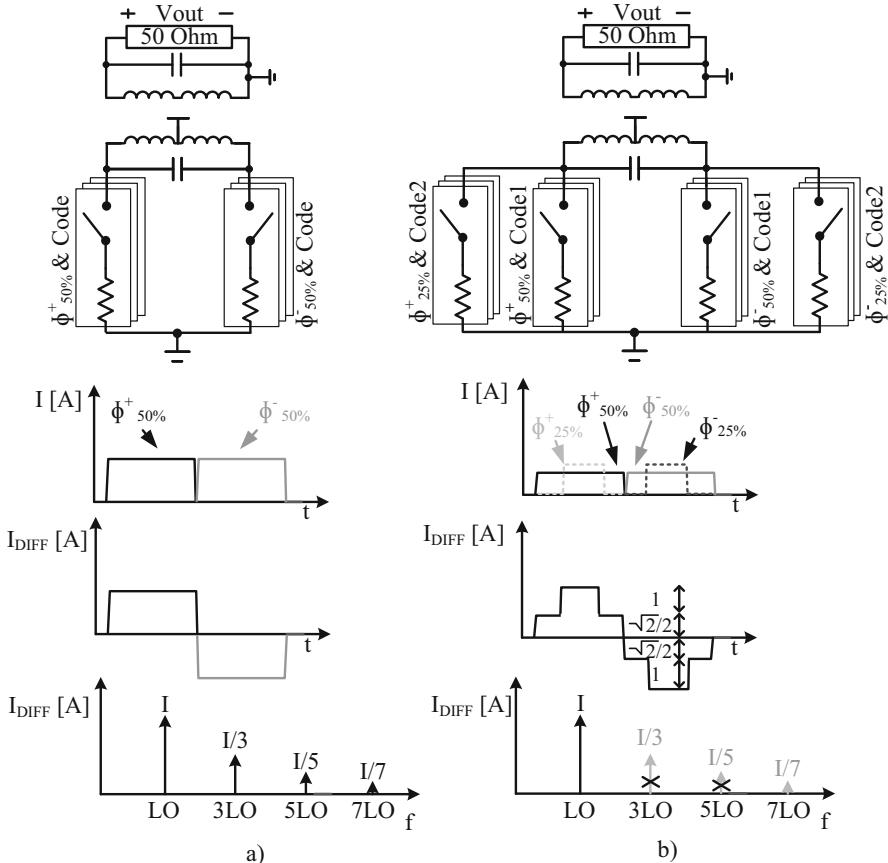


Fig. 4.20 Typical versus harmonic rejection mixed class D^{-1} DPA. **(a)** Operation principle of a typical inverse class-D DPA. **(b)** Operation principle of an inverse class-D DPA with harmonic rejection mixing

opening window and the ratio of currents over a half-phase has to be $(\sqrt{2}+1)/\sqrt{2}$ [Weldon01]. Note that any phase or amplitude imbalance results with imperfect cancellation. To achieve at least 30 dB of suppression, the current amplitude scaling error should be below 10% and the phase error below 1.5° [Weldon01].

We make use of a secondary switch bank that is put in parallel to the original array to enable HRM (Fig. 4.20b). The two 10 bit banks are physically identical, they make use of a 290 nm wide unit cell switch (28 nm in length) in a 6 LSB binary and 4 MSB thermometric segmentation. The switches are driven by AND gates that mix code with the LO signal, as desired by the polar amplitude modulation process. Note that the AM data is latched by the LO low-phase, so that no AM is permitted during the switch ON phase. One pseudo-differential bank receives the 50% duty cycle LO drive, while the other operates based on a 25% duty cycle LO signal input

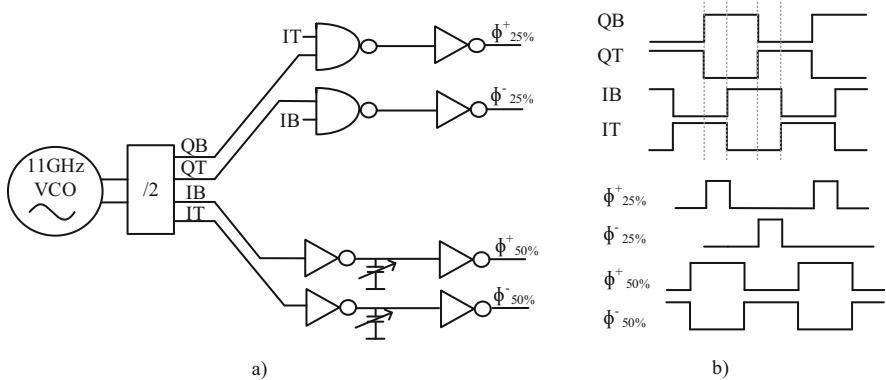


Fig. 4.21 HRM signals generation. (a) Harmonic rejection mixing implementation. (b) Harmonic rejection mixing waveforms

(delay adjusted). Instead of analog switch size scaling for the desired amplitude ratio over a period (nonlinear behavior), we use digital control. The ratio of codes is easily configured in the digital domain, where a LUT determines scaled version of the codes sent to the 50% duty-cycled bank (with respect to the codes sent to the 25% duty-cycled bank) ensuring desired switch impedance ratio over an LO period and strong HRM.

Although DPA's drain efficiency is not optimized in this design, the following observation can be made: (1) harmonics represent unnecessarily wasted power (no useful information), hence minimizing them can potentially maximize the efficiency; (2) HRM class D⁻¹ DPA can also operate in zero voltage switching condition.

The buffering from the conceptual schematic in Fig. 4.21 is implemented through custom made digital logic. The input and the output of the divider are DC shielded, with a fully reconfigurable DC biasing network to enable accurate duty cycle control of the signals driving the DPA. The inverters in the path are scaled in a tapered configuration to ensure a reliable square shape driving the final DPA buffer. A digitally controlled delay line is used on the 50% duty cycle signal to make it fall exactly in the middle of the 25% duty cycle signal's ON phase. To ensure correct phase matching we implement a 7 bit, 0.4 ps/LSB delay line. The addition of HRM in the structure obviously has a power repercussion, since the LO buffering became more complex.

The balun is integrated on chip. Full thick metal stack is used to implement the coils, ensuring low resistivity and a quality factor of 12 across the frequencies of interest. The size of the inductors in the primary and secondary are 380 pH and 410 pH, respectively. To achieve the maximal power transfer across the full range of output frequencies from 5 GHz to 6 GHz, the balun needs resonant frequency tuning. This is achieved by placing a tunable digital varactor bank in parallel to the primary inductor. The bank is 4 bit wide (plus a half-cell) with 90 fF/LSB. Note that

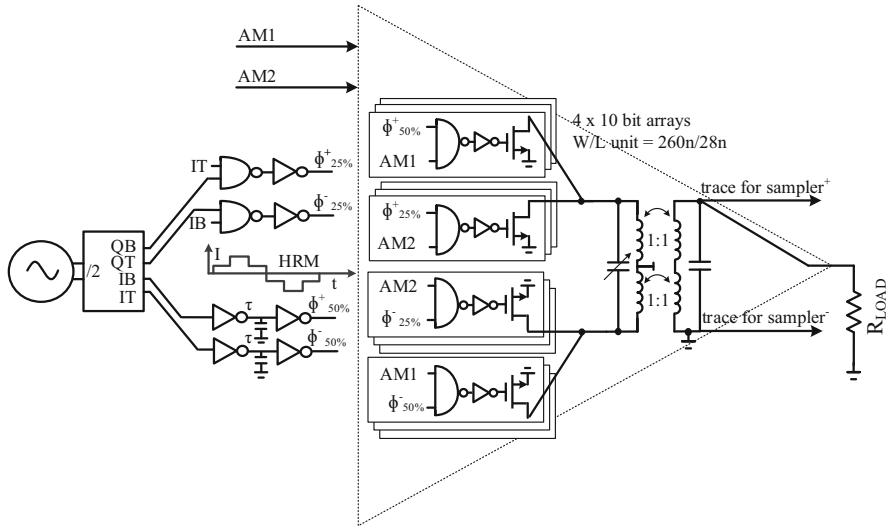


Fig. 4.22 HRM class D⁻¹ overview

a 200 fF MOM capacitor is fixed in the primary and a capacitor of 1.2 pF is fixed in the secondary.

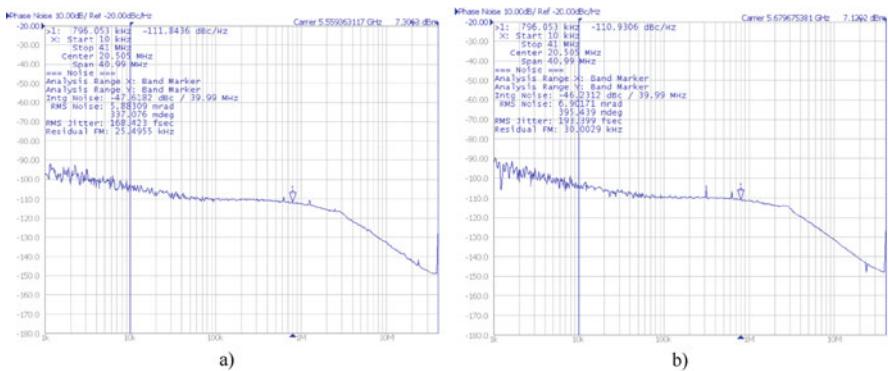
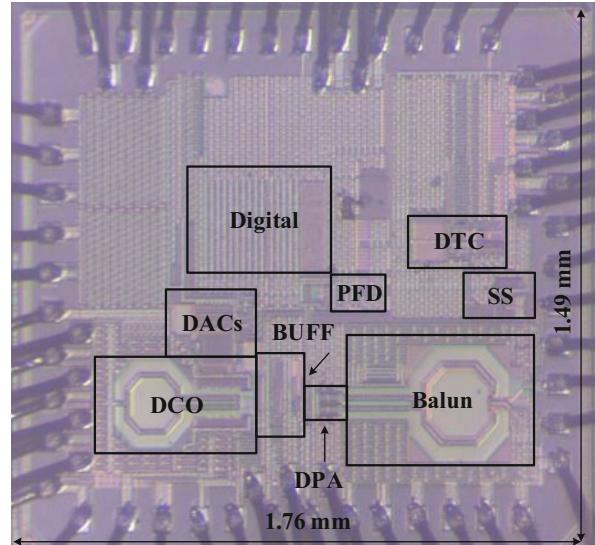
The complete DPA architecture is illustrated in Fig. 4.22. Note that the digital processor determines AM1 and AM2 input codes independently. The accurate ratio of codes that ensures desired current scaling for strong HRM can be experimentally determined and fixed. Similarly, the delay τ which ensures desired phase matching for HRM is fixed experimentally.

4.6 Measured Results

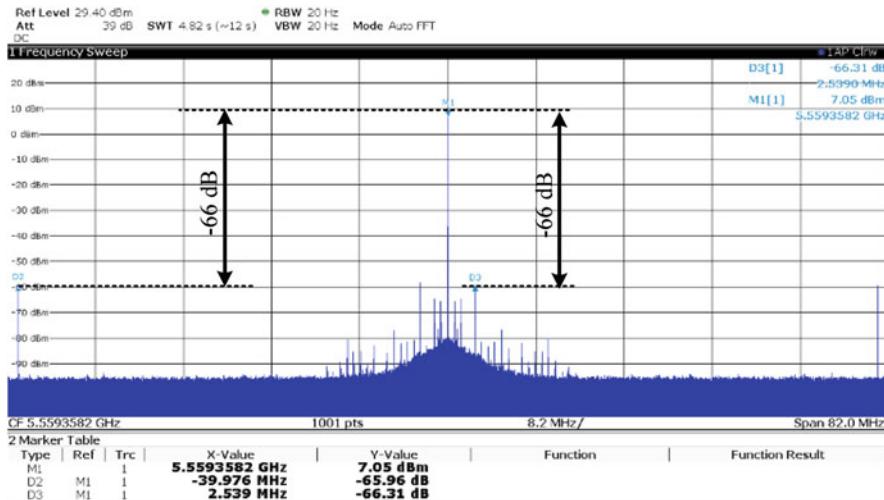
The SSPTX was fabricated in 28 nm bulk digital CMOS with the chip size of 2.62 mm² and an active area of approximately 50%. One of the measured die's micrographs is shown in Fig. 4.23. The chip operates from a 0.9 V supply, while the IO interface uses a separate 1.8 V supply.

4.6.1 Digital Subsampling PLL Measurements

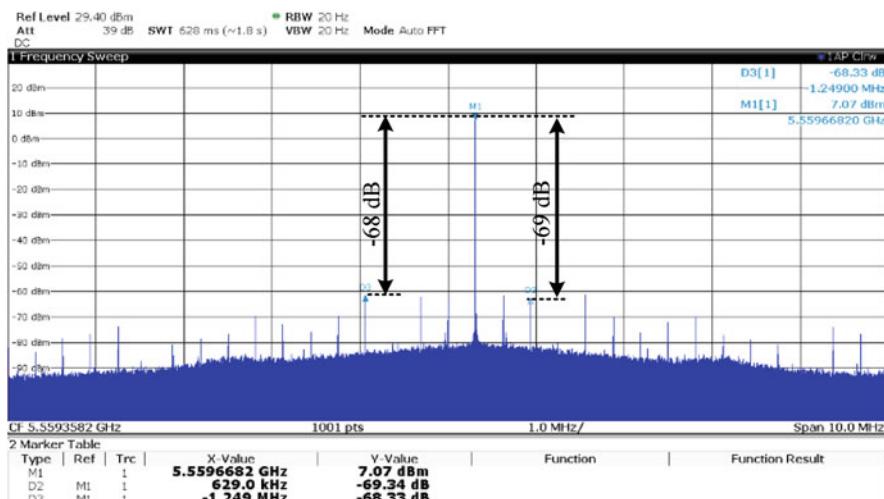
We analyze performance of the background-calibrated PLL first, when no modulation is present in the system, with 2 different VCO setting modes: (1) high-power and (2) low-power mode. The DPA is set to a constant mean output power of 7 dBm. A fractional-N output around 5.6 GHz (with in-band spurs) is synthesized

Fig. 4.23 Die micrograph**Fig. 4.24** Measured phase noise of the PLL. (a) High-power VCO mode. (b) Low-power VCO mode

from a 40 MHz crystal reference. Figure 4.24 shows measured phase noise. In high-power VCO mode, the PLL operates with 168.4 fs RMS jitter (-47.6 dBc IPN). Note that the integration range (10 kHz–40 MHz) includes all spurs. The in-band noise floor is in this mode -111.8 dBc/Hz. In low-power VCO mode, the RMS jitter and in-band noise floor are 193.4 fs (-46.2 dBc IPN) and -110.9 dBc/Hz, respectively. Thanks to the background-cancelled PD nonlinearities (DTC), the PLL operates with virtually no difference between integer-N and fractional-N modes (Δ IPN < 0.25 dB). Spurious content is similar in both modes and reported through Fig. 4.25. The integer-N spur level is -66 dBc. Its origin is in-band, since it can be reduced by choosing a lower PLL bandwidth (with suboptimal noise filtering as a consequence). The worst case measured fractional-N spur is at -66 dBc, which is the



a)



a)

Fig. 4.25 Measured spectrum of the PLL. (a) Spectrum with worst case fractional-N spur and integer-N spur. (b) Spectrum during synthesis of a channel with deep in-band spur

residue spur limited by accuracy of devised DTC predistortion and potentially other parasitic effects such as supply noise/coupling. Note that the performance of the PLL as shown in Fig. 4.24 matches the performance of the analog subsampling PLL as presented in Chap. 3 (see Fig. 3.18). The measured jitter (frequency independent) is similar in the compared systems, both, in low and in high-power modes. This stands as a proof that the loop digitization was successful and that there is no added phase noise degradation between the analog and the digital loop, as desired.

The overall DC current consumption of the PLL loop from a 0.9 V supply is as follows: DTC, signal amplifier, and ADC consume 1.3 mA, DAC consumes 0.3 mA, digital consumes 2.2 mA, and the VCO consumes 5.6 mA/10.3 mA in low-power/high-power modes. If we exclude from the PLL power budget the LO buffering and DPA that are used for delivering power to the 50Ω output load, the subsampling PLL loop operates with total 12.7 mW/8.5 mW in high-power/low-power modes. This results with a competitive fractional-N PLL FOM³ of -245 dB (with 0.5 dB less in high-power mode).

4.6.2 Digital Subsampling Polar TX Measurements

The SSPTX performance is initially analyzed at 2.5 MHz modulation bandwidth. Figure 4.26 shows TX output spectrum at -2.7 dBm average power. The measured EVM is -41.3 dB, which enables a clear 1024 QAM constellation. The measured Adjacent Channel Leakage Ratio (ACLR)1/2 is $-36/-38$ dB. This level of accuracy was predicted by simulations and proves that the distortion mechanisms had been identified and compensated well—note that all the calibration loops are operating simultaneously in the background. Figure 4.27 shows the TX output spectrum at 1.2 dBm average output power (1024 QAM). The EVM degrades to -40.1 dB EVM which is likely due to an increased amount of AM-to-PM distortion present in the system (discussed later). The measured ACLR1/2 is $-35/-38$ dB. The average power can be increased further with similar consequence, for example, at 3.7 dBm, the EVM drops to -37.1 dB which is still enough for a clear 256 QAM constellation. To prove the importance of the devised algorithms, we report measured EVM in steps, as an increasing number of nonidealities is cancelled in the background (see Table 4.1). The table clearly shows that enabling background calibration results with at least 14 dB improvement in measured EVM. Figure 4.28 shows measured predistortion curves as settled in the background.

Measurements disclose that with the increase of the DPA’s average output power delivered to the load, the quality of the transmitted signal drops. As stated earlier, this is likely due to increased cross-distortion between the AM and PM path. First, the DPA has a simulated fixed AM-to-PM distortion profile as shown in Fig. 4.14, a constant output phase dependence on the used AM amplitude. Second, AM-to-PM distortion can be dynamic as well, for example, the DPA can *pull* the VCO, i.e., influence the speed at which the direct frequency modulation happens at the VCO output—in dependence with the AM code change. As discussed in Sect. 4.4, the SSPTX filters AM-to-PM distortion in-band. Figure 4.29 shows EVM dependence on the SSPTX bandwidth setting, at 3.7 dBm average TX output power, where AM-to-PM is more pronounced. A clear improvement is seen for the

³ $FOM_{PLL} = 10 \log_{10} \left[\left(\frac{P_{DC}}{1mW} \right) \left(\frac{RMS_{jitter}}{1s} \right)^2 \right]$ as defined in [Gao09a].

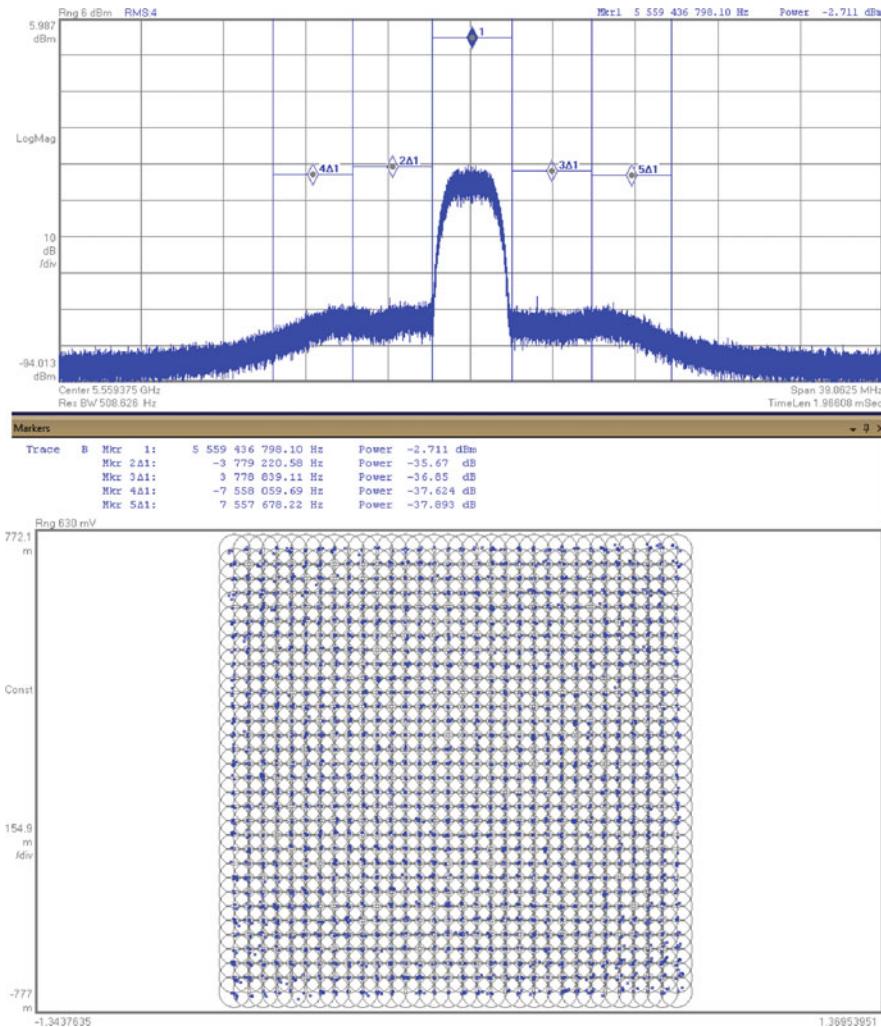


Fig. 4.26 SSPTX output at -2.7 dBm average power (1024 QAM). EVM = -41.3 dB

increased PLL bandwidth, even at frequency cut-offs that are beyond the optimal phase noise filtering point (about 1.2 MHz as seen in Fig. 4.24). In other words, SSPTX enables, in contrast to a classical polar TX, an additional degree of design freedom: optimal loop filtering setting for highest transmission quality that includes AM-to-PM distortion suppression.

The importance of *background* calibration becomes pronounced in time-variant environments. Figure 4.30 shows measured EVM with $\pm 5\%$ variation in the supply, around the nominally calibrated 0.9 V point. Without environmental change tracking, the EVM increases as the original distortion estimates fluctuate with the

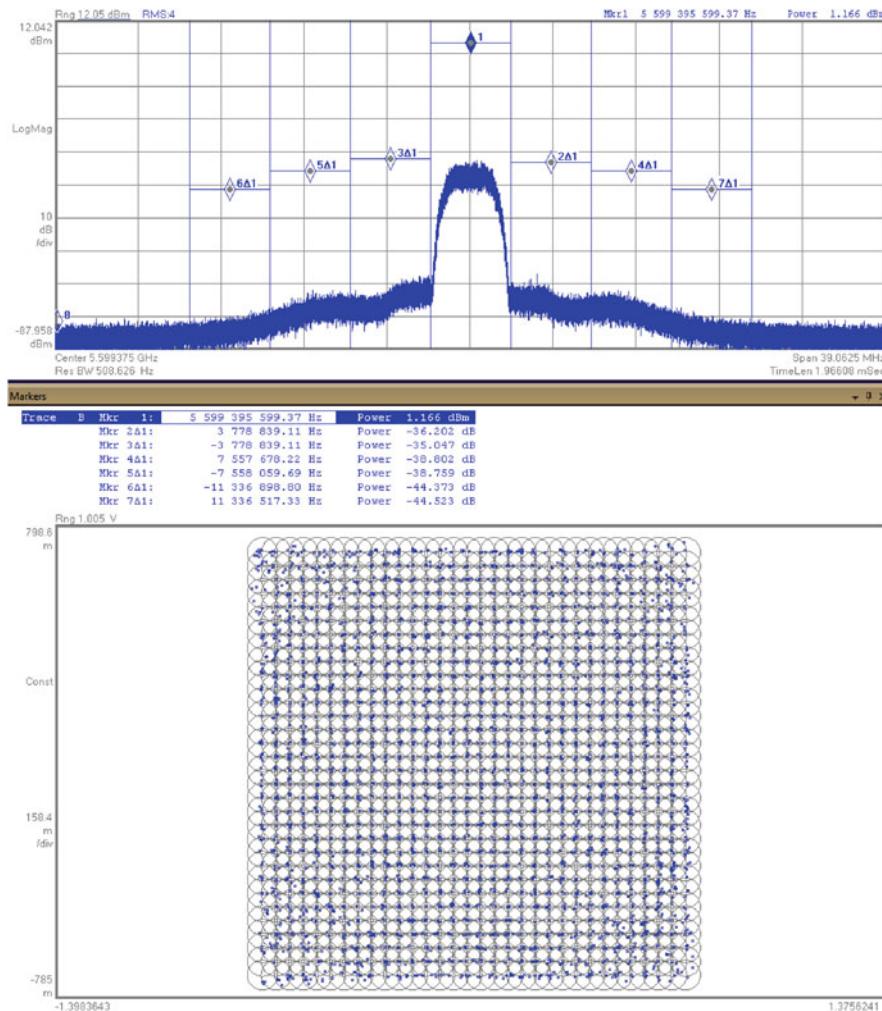


Fig. 4.27 SSPTX output at 1.2 dBm average power (1024 QAM). EVM = -40.1 dB

Table 4.1 EVM verification at 2.5 MHz modulation BW with increasing number of background calibration algorithms

	-2.7 dBm	1.2 dBm	3.7 dBm
Average power/constellation	1024 QAM	1024 QAM	256 QAM
No calibration	> -27.3 dB	> -26.1 dB	> -22.0 dB
+ AM-to-AM calibration	-35.4 dB	-32.4 dB	-31.7 dB
+ DTC INL calibration	-40.1 dB	-37.0 dB	-35.4 dB
+ DAC INL calibration	-41.3 dB	-40.1 dB	-37.1 dB

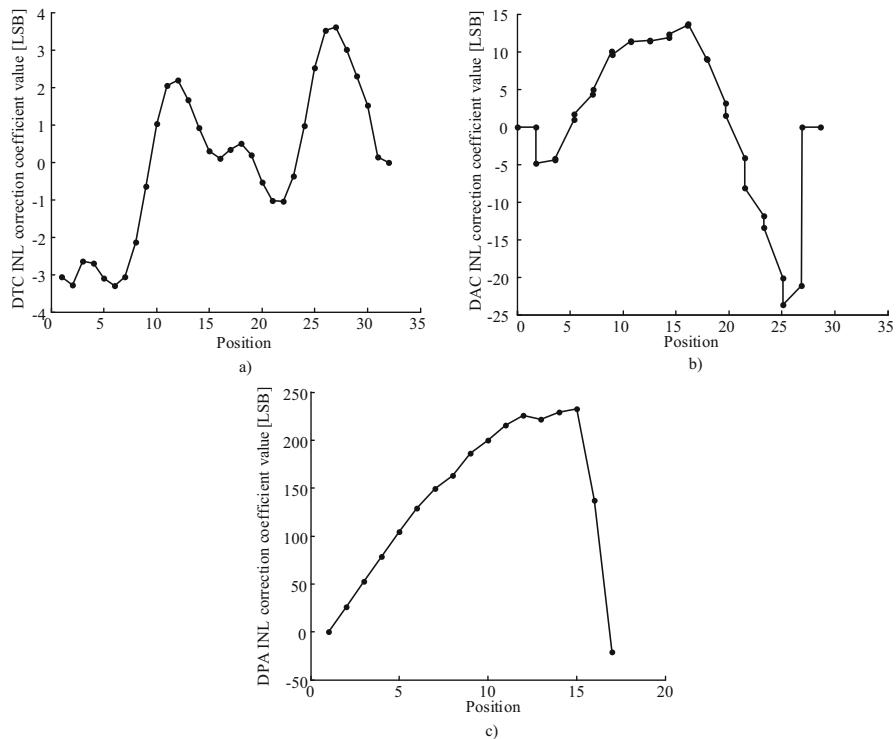


Fig. 4.28 Measured predistortion correction coefficients for (a) DTC INL; (b) DAC INL; (c) DPA INL calibration (after settling)

Fig. 4.29 Measured EVM at 1.2 dBm average output power as a function of PLL bandwidth

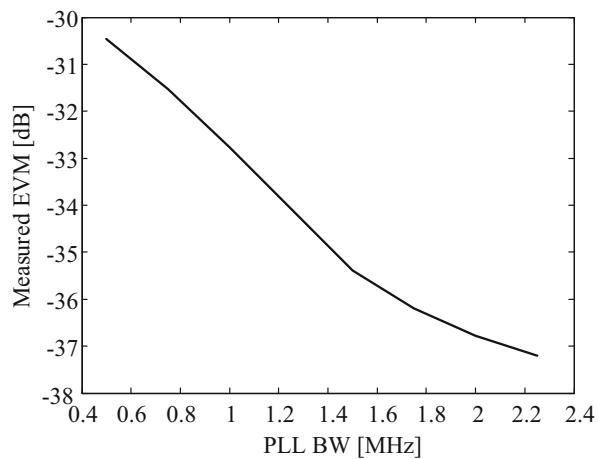


Fig. 4.30 Measured EVM with supply level variation: with and without background calibration

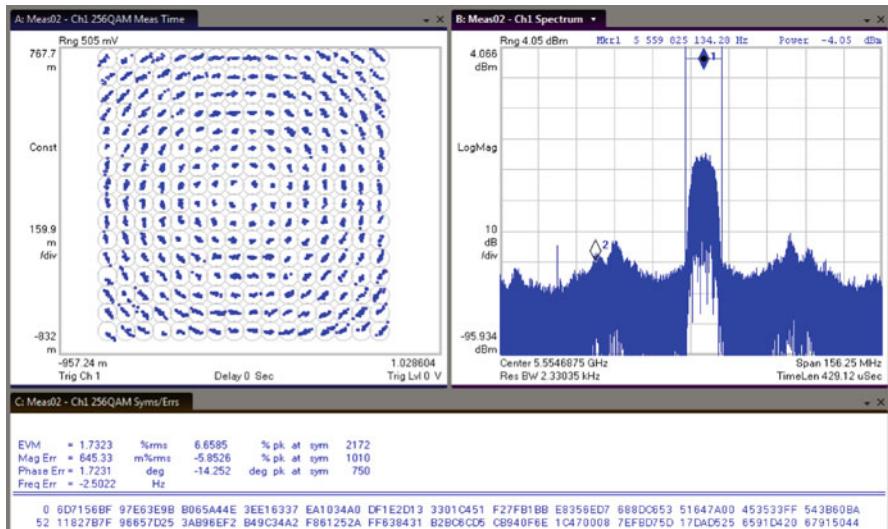
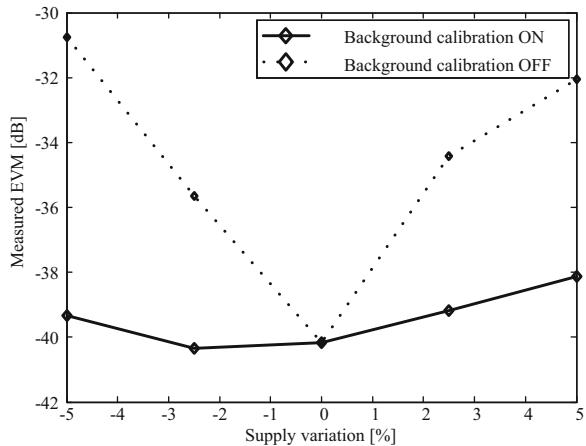


Fig. 4.31 SSPTX output at 10 MHz modulation BW

supply level shift. This problem is resolved with background error cancellation enabled by the SSPTX. Note that the continuous calibration does not disrupt the transmitted signal.

At modulation bandwidth much larger than the loop bandwidth, the SSPTX loses AM-to-PM filtering capabilities. Hence, to ensure a visible 256QAM constellation at 10 MHz bandwidth modulation, the SSPTX had to operate at very low average AM powers (-4 dBm). As shown in Fig. 4.31, the transmitter then produces -35.4 dB EVM. This degrades at higher output powers, for example, at -1.3 dBm average output power, the measured EVM is -32.3 dB which is enough only for a visible 64QAM constellation. Note that the calibration continuously operates

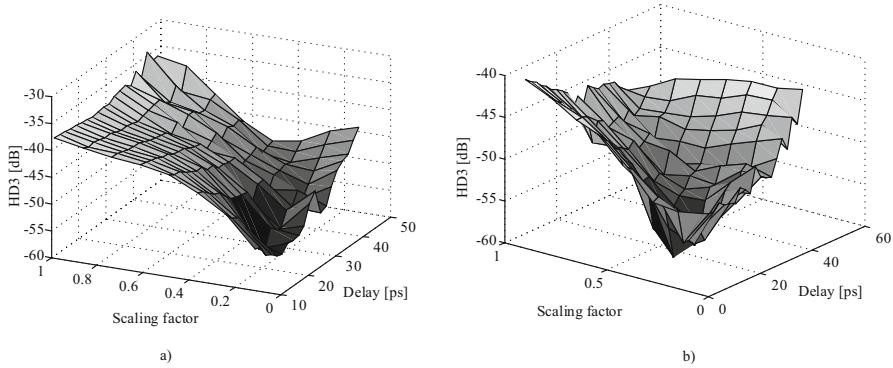


Fig. 4.32 HRM performance as a function of the AM scaling factor and relative phase delay between the 50% and 25% duty-cycled LO signal. **(a)** AM = 1023. **(b)** AM = 512

in the background, independently of modulation bandwidth. Figure 4.31 also shows aliases around the 40 MHz DAC clock. To improve the spectral profile, a higher clock rate should be employed (and better means for AM-to-PM distortion cancellation).

The HRM performance is evaluated by running the DPA at a constant output frequency (in phase lock) as a function of the scaling factor (i.e., the weight ratio between the code sent to the 50% (bank 1) and 25% (bank 2) LO duty-cycled switch bank) and the delay at which the 25% LO opening window appears with respect to the 50% one (see Sect. 4.5.2). The evaluated results are reported in Fig. 4.32 for suppression of the 3rd LO harmonic, at 2 separate full scale DPA output amplitudes (code 1023 and code 512). At maximum code (1023) sent to the bank 2, the best suppression (60 dB) is achieved with the amplitude scaling factor of 0.25 (bank 1 receives 0.25 times the full scale code of the bank 1) and a phase separation of approximately 22 ps. Similarly at code 512, the best suppression (59 dB) is achieved with the amplitude scaling factor of 0.35 (bank 1 receives 0.35 times the full scale code sent to the bank 1) and a similar phase separation. The measured results show that the HRM can be successfully achieved over the complete range of output powers with pure digital scaling of the AM codes. A negative consequence of scaling down the used bank 1 codes is output power degradation. In other words, less power delivered in the 50% duty-cycled window is not compensated by delivering more in the 25% duty-cycled window. This can be addressed in a redesign and was not under optimization in this proof-of-concept implementation. The SSPTX shows no degradation as long as the harmonic suppression is higher than approximately 30 dB (see Sect. 4.5.2), hence the measured quality of the transmitted signal (EVM) remains independent of HRM scaling used. The peak DPA output power is 11 dBm with the drain current consumption of 69 mA (20% drain efficient).

The DC current consumption from a 0.9 V supply during SSPTX transmission is as follows: digital consumes 10 mA, VCO consumes 7 mA, DACs consume 0.4 mA, and ADC, signal amplifier, and DTC consume 2.3 mA. The LO buffering and DPA

Table 4.2 SSPTX performance summary and comparison to other recent TX architectures

	This work	[Ba16]	[Fulde17]	[Boos11] ^a
Architecture	SSPTX	Polar TX	Polar TX	Polar TX
Technology [nm]	28	40	65	28
Supply [V]	0.9	1	1.2/2.5	1.3/1.1/1
Freq [GHz]	$5.5 \pm 10\%$	0.75–0.93	0.8–2.1 ^b	2.3–2.8
BW [MHz]	2.5 10	2	3.84	10
Pout _{average} [dBm]	–2.7/1.1/3.7 –4.0/–1.3	0	6	6
QAM	1024/1024/256 256/64	64	n.a.	n.a.
EVM [dB]	–41.3/–40.3/–37.2 –35.2/–32.2	–27	–29.1	–33.6
Consumption [mW]	–46.5/55/59 46	–7	72.5 ^c	n.a. ^d
Pout _{peak} [dBm]	11	8	8	12

^a At LTE10 high band

^b Paper indicated 2G/3G

^c Only RFDAC consumption at 6 dBm (no LO buffering or digital interface consumption)

^d 44 mA at –18 dBm unclear from which VDD

average consumption depends on the desired delivered output power. At –2.7 dBm average output power, 1024 QAM, the current consumption is 25 mA for the LO buffering and 7 mA for the DPA. At 1.2 dBm average output power, 1024 QAM, the current consumption is 25 mA for the LO buffering and 11 mA for DPA. And at the reported 3.7 dBm average output power, 256 QAM, the LO buffering consumes 29 mA while the DPA pulls 17 mA. Note that different power levels were achieved by digital scaling of the AM codes.

In summary, thanks to the linearized environment that allows high order constellation schemes, the delivered data throughput by the SSPTX is highly efficient: 1.86 nJ/bit (1024 QAM at 2.5 MHz BW and average DC power consumption of 46.5 mW). At 10 MHz BW and 256 QAM, we reach very low 0.58 nJ/bit. An overview of the SSPTX performance is shown in Table 4.2 with a brief comparison to other architectures. Note that this TX achieves lowest EVM at highest QAM complexity, while it operates at higher frequency band where integrated phase noise imposes larger limitations and power aided efficiency is normally smaller.

4.7 Conclusion

We have presented the first implementation of a subsampling polar transmitter that consists of a low-noise phase modulating digital subsampling PLL and an inverse-class-D digital power amplifier for amplitude modulation. The DPA is, unlike in a typical polar transmitter, placed within the PLL and the phase-error detection happens directly at the DPA output. The SSPTX is sensitive not only to phase errors, but also to modulation *amplitude*. The feature enables distortion filtering and accurate linearization. AM-to-PM distortion is suppressed by the loop operation and AM-to-AM, PM-to-PM distortion are detected and cancelled digitally in the

background, while the transmitter operates normally. Thanks to the highly linearized environment, the architecture enables use of complex constellation schemes (1024 or 256 QAM) with 2.5/10 MHz modulation bandwidth around a 5.5 GHz carrier, ensuring spectrally efficient data throughput that is ubiquitously desired in the existing and upcoming wireless links.

References

- [Ba16] A. Ba, Y.-H. Liu, J. van den Heuvel, P. Mateman, B. Büsze, J. Dijkhuis, C. Bachmann, G. Dolmans, K. Philips, H. De Groot, A 1.3 nJ/b IEEE 802.11 ah fully-digital polar transmitter for IoT applications. *IEEE J. Solid State Circuits* **51**(12), 3103–3113 (2016)
- [Boos11] Z. Boos, A. Menkhoff, F. Kuttner, M. Schimper, J. Moreira, H. Geltinger, T. Gossmann, P. Pfann, A. Belitzer, T. Bauernfeind, A fully digital multimode polar transmitter employing 17b RF DAC in 3G mode, in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International* (IEEE, San Francisco, 2011), pp. 376–378.
- [Chowdhury11] D. Chowdhury, L. Ye, E. Alon, A. M. Niknejad, An efficient mixed-signal 2.4-GHz polar power amplifier in 65-nm CMOS technology. *IEEE J. Solid State Circuits* **46**(8), 1796–1809 (2011)
- [Chowdhury12] D. Chowdhury, S.V. Thyagarajan, L. Ye, E. Alon, A. Niknejad, A fully-integrated efficient CMOS inverse class-D power amplifier for digital polar transmitters. *IEEE J. Solid State Circuits* **47**(5), 1113–1122 (2012)
- [Fulde17] M. Fulde, A. Belitzer, Z. Boos, M. Bruennert, J. Fritzin, H. Geltinger, M. Groinig, D. Gruber, S. Gruenberger, T. Hartig et al., 13.2 A digital multimode polar transmitter supporting 40MHz LTE Carrier Aggregation in 28nm CMOS, in *Solid-State Circuits Conference (ISSCC), 2017 IEEE International* (IEEE, San Francisco, 2017), pp. 218–219
- [Gao09b] X. Gao, E. Klumperink, M. Bohsali, B. Nauta, A low noise sub-sampling PLL in which divider noise is eliminated and PD/CP noise is not multiplied by N^2 . *IEEE J. Solid State Circuits* **44**(12), 3253–3263 (2009)
- [Gao09a] X. Gao, E. Klumperink, P. Geraedts, B. Nauta, Jitter analysis and a benchmarking figure-of-merit for phase-locked loops. *IEEE Trans. Circuits Syst. Express Briefs* **56**(2), 117–121 (2009)
- [Gao10] X. Gao, E. Klumperink, G. Soccia, M. Bohsali, B. Nauta, Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector. *IEEE J. Solid State Circuits* **45**(9), 1809–1821 (2010)
- [Gao15] X. Gao, E. Klumperink, B. Nauta, Sub-sampling PLL techniques, in *Custom Integrated Circuits Conference (CICC), 2015 IEEE* (IEEE, San Jose, 2015), pp. 1–8
- [Hershberg14] B. Hershberg, K. Raczkowski, K. Vaesen, J. Craninckx, A 9.1–12.7 GHz VCO in 28nm CMOS with a bottom-pinning bias technique for digital varactor stress reduction, in *European Solid State Circuits Conference (ESSCIRC), ESSCIRC 2014-40th* (IEEE, Venice Lido, 2014), pp. 83–86
- [Levantino14] S. Levantino, G. Marzin, C. Samori, An adaptive pre-distortion technique to mitigate the DTC nonlinearity in digital PLLs. *IEEE J. Solid State Circuits* **49**(8), 1762–1772 (2014)
- [Lin14] F. Lin, P.-I. Mak, R. P. Martins, A sine-LO square-law harmonic-rejection mixer-theory, implementation, and application. *IEEE Trans. Microwave Theory Tech.* **62**(2), 313–322 (2014)

- [Madoglio17] P. Madoglio, H. Xu, K. Chandrashekhar, L. Cuellar, M. Faisal, W.Y. Li, H.S. Kim, K.M. Nguyen, Y. Tan, B. Carlton et al., 13.6 A 2.4 GHz WLAN digital polar transmitter with synthesized digital-to-time converter in 14nm trigate/FinFET technology for IoT and wearable applications, in *Solid-State Circuits Conference (ISSCC), 2017 IEEE International* (IEEE, San Francisco, 2017), pp. 226–227
- [Malki16] B. Malki, B. Verbruggen, E. Martens, P. Wambacq, J. Craninckx, A 150 kHz–80 MHz BW discrete-time analog baseband for software-defined-radio receivers using a 5th-Order IIR LPF, active FIR and a 10 bit 300 MS/s ADC in 28 nm CMOS. *IEEE J. Solid State Circuits* **51**(7), 1593–1606 (2016)
- [Markulic14] N. Markulic, K. Raczkowski, P. Wambacq, J. Craninckx, A 10-bit, 550-fs step digital-to-time converter in 28nm CMOS, in *ESSCIRC 2014—40th European Solid State Circuits Conference (ESSCIRC)* (IEEE, Venice Lido, Sept 2014), pp. 79–82
- [Markulic16a] N. Markulic, K. Raczkowski, E. Martens, P.E.P. Filho, B. Hershberg, P. Wambacq, J. Craninckx, 9.7 A self-calibrated 10Mb/s phase modulator with -37.4dB EVM based on a 10.1-to-12.4GHz, -246.6dB-FOM , fractional-N subsampling PLL, in *2016 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, San Francisco, Jan 2016), pp. 176–177
- [Markulic16b] N. Markulic, K. Raczkowski, E. Martens, P.E. Paro Filho, B. Hershberg, P. Wambacq, J. Craninckx, A DTC-based subsampling PLL capable of self-calibrated fractional synthesis and two-point modulation. *IEEE J. Solid State Circuits* **51**(12), 3078–3092 (2016)
- [Marzin12] G. Marzin, S. Levantino, C. Samori, A. L. Lacaita, A 20 Mb/s phase modulator based on a 3.6 GHz digital PLL with -36 dB EVM at 5 mW power. *IEEE J. Solid State Circuits* **47**(12), 2974–2988 (2012)
- [Marzin14] G. Marzin, S. Levantino, C. Samori, A. L. Lacaita, 2.9 A background calibration technique to control bandwidth in digital PLLs, in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International* (IEEE, San Francisco, 2014), pp. 54–55
- [Mensink05] E. Mensink, E.A. Klumperink, B. Nauta, Distortion cancellation by polyphase multipath circuits. *IEEE Trans. Circuits Syst. Regul. Pap.* **52**(9), 1785–1794 (2005)
- [Raczkowski15] K. Raczkowski, N. Markulic, B. Hershberg, J. Craninckx, A 9.2–12.7 GHz wideband fractional-N subsampling PLL in 28 nm CMOS with 280 fs RMS jitter. *IEEE J. Solid State Circuits* **50**(5), 1203–1213 (2015)
- [Tasca11] D. Tasca, M. Zanuso, G. Marzin, S. Levantino, C. Samori, A. Lacaita, A 2.9–4.0-GHz fractional-N digital PLL with bang-bang phase detector and 560-fs RMS integrated jitter at 4.5-mW power. *IEEE J. Solid State Circuits* **46**(12), 2745–2758 (2011)
- [VDB13] A. Van Den Bosch, M. Steyaert, W. Sansen, *Static and Dynamic Performance Limitations for High Speed D/A Converters*, vol. 761. (Springer Science & Business Media, New York, 2013)
- [Weldon01] J.A. Weldon, R.S. Narayanaswami, J.C. Rudell, L. Lin, M. Otsuka, S. Dedieu, L. Tee, K.-C. Tsai, C.-W. Lee, P.R. Gray, A 1.75-GHz highly integrated narrow-band CMOS transmitter with harmonic-rejection mixers. *IEEE J. Solid State Circuits* **36**(12), 2003–2015 (2001)
- [Working Gr17] Working Group of the 802 Committee and others, Draft Standard for Information Technology Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific Requirements, *IEEE P802.11ax/D1.4* (August 2017)
- [Yoo11] S.-M. Yoo, J.S. Walling, E.C. Woo, B. Jann, D.J. Allstot, A switched-capacitor RF power amplifier. *IEEE J. Solid State Circuits* **46**(12), 2977–2987 (2011)

- [Yoo13] S.-M. Yoo, J.S. Walling, O. Degani, B. Jann, R. Sadhwani, J.C. Rudell, D.J. Allstot, A class-G switched-capacitor RF power amplifier. *IEEE J. Solid State Circuits* **48**(5), 1212–1224 (2013)
- [Zheng15] S. Zheng, H.C. Luong, A WCDMA/WLAN digital polar transmitter with low-noise ADPLL, wideband PM/AM modulator, and linearized PA. *IEEE J. Solid State Circuits* **50**(7), 1645–1656 (2015)

Chapter 5

Conclusion and Future Outlook



5.1 Summary

Thanks to the extreme connectedness of today's world, it is easier than ever to transfer information and to communicate.¹ Wireless links normally use transceivers with a local oscillator in their heart that is typically implemented as a PLL. Reliable and pristine PLL output frequency beat is of crucial importance for the efficient spectrum usage—and higher-order modulation schemes (at large bandwidths) are required to satisfy the end-users growing hunger for fast data throughput. The PLL's phase noise and spurious content indeed impose the fundamental limit to the *density* with which the information can be transmitted. Besides the exclusive LO generation, a PLL can be the basis for compact and power-efficient *polar* transmission. Polar TX architecture is an attractive, digitally intensive solution that simultaneously comes with a set of severe design challenges (such as bandwidth limitation and linearity) that need to be carefully tackled when targeting high-speed communication. This book offers contributions to the state of the art in fractional frequency synthesis and polar transmitter design. The presented material is built around three 28-nm bulk CMOS IC prototypes that investigate and push the boundaries of the field.

The first design (Chap. 2) explores opportunities behind a subsampling PLL architecture. This divider-less loop makes use of a phase-error detection mechanism (subsampling) with extremely high gain that efficiently compresses in-band PLL noise. The loop therefore permits extension of the PLL filtering bandwidth, which is beneficial for the VCO's phase noise filtering profile. This in turn leads to minimization of the overall integrated noise and very low jitter. In fact, the integer-N subsampling PLL proves to be a leading architecture for low noise and for power vs. performance tradeoff. Unfortunately, this synthesizer had originally no phase

¹The quality of information content is out of this book's scope.

modulating capabilities, i.e., no functionality for fractional synthesis and hence it had very little use for modern wireless standards. The VCO's output zero-crossing detection by subsampling at the reference rate (for detection of phase mismatch) is indeed only functional for integer ratio between the input and the output frequency. We present how to enable fractional-N operation using a subsampling PLL by introduction of a *digital-to-time converter* (DTC) in the reference path of the PLL. The DTC exploits *time* rather than *voltage*-domain processing, which is a trend repeatedly seen in the recent time-to-digital (TDC) based digital PLLs. A DTC, in contrast to a classical TDC, easily reaches fine resolution (low quantization noise) that is fundamental for high performance. Its purpose is to dynamically compensate for the fractional residue between the input and the output frequency, by instantaneously delaying the reference sampling edge. We propose one of the first DTC designs suitable for digitally intensive frequency synthesis, with special care drawn to analysis of its linearity, phase noise, and supply noise rejection. The goal is naturally to maintain all the benefits of subsampling while operating in the fractional-N multiplication mode. The newly developed FNSSPLL operates in the 9.2–12.7-GHz range with 280 fs of RMS jitter in the worst-case fractional-N modes. With a 40-MHz crystal oscillator reference, the measured in-band performance of the PLL is -104 dBc/Hz. This measure stands as a proof of the concept. At the moment of publication, no other fractional-N PLL design reported better in-band noise. The reported FOM of the PLL is -240 dB, at 13-mW power consumption.

In the second design (Chap. 3), we address the fundamental limitations of the newly proposed loop to minimize the performance gap between integer-N and fractional-N subsampling modes. To achieve this, the FNSSPLL must operate without nonlinearities in the phase comparison path (and without large or data-dependent quantization noise excess). Note that for example, in typical analog PLL's these problems originate within the linearity of the charge pump, while in digital loops, the TDC becomes typically the major culprit for the performance degradation (because of its nonlinearity and/or limited resolution). In a DTC-based subsampling FNSSPLL, it is of course the performance of the DTC that sets the bottleneck for the PLL's output spectral purity. In the first design, we worked towards analog resilience of the DTC to different degradation mechanisms, such as code-dependent supply noise or its nonlinearity. Similarly, we optimized the DTC's quantization error, keeping it below the environmental noise. In this design we explore digital DTC background calibration and randomization techniques to improve inherit performance of the DTC. We develop algorithms that rely on a simple single-bit (bang-bang) error extraction at the output of the PLL's PD. The digital information about the error can be correlated to the instantaneously used DTC input codes to continuously track its behavior. Simultaneously, we use this information to determine optimal code predistortion that nulls the undesired nonlinearity in its origin. Thanks to the digital calibration that tracks PVT in the background, the second prototype shows almost no performance gap between the integer and fractional modes. The designed PLL operates around 10.1–12.4 GHz with best-/worst-case RMS jitter of 176/197 fs at 5.6 mW power, achieving FOMs that exceed -247 dB and challenge state of the art in PLL design. Moreover, within

the same IC prototype we extend the basic loop and propose the first subsampling PLL-based phase/frequency modulator. To achieve wide-band phase modulation, we make use of two-point modulation. To avoid signal degradation, the injection paths need to be accurately matched, both in injection gain and in the time instant at which the modulation data is injected. The linearization techniques proposed earlier are similarly adopted in modulation mode (DTC and modulating DAC background calibration) to achieve high accuracy in terms of EVM. The mixed mode phase modulator operates at 10-MHz GMSK modulation bandwidth with up to -40.5 dB EVM, which surpasses similar work within the state of the art.

The third IC prototype (Chap. 4) of this book introduces a *Digital Subsampling Polar Transmitter* (SSPTX) as a novel architecture that offers several advantages with respect to a classical polar TX. The new system build-up starts by digitization of the FNSSPLL. We make use of an ADC that converts information about sampled phase error into digital data and feeds it to a digital loop filter. This stands as a point of contrast to an analog loop in which the sampled error biases a transconductor that feeds current into a (area consuming) filtering capacitor. The loop filter produces a control signal for the DCO that is built from a reference clocked resistive DAC and a varactor. The Chap. 4 describes how the added quantization noise can be handled to avoid performance degradation with respect to the analog loop. The digital subsampling PLL can operate as a phase modulator in a fully background-calibrated environment that was build based on the previously developed techniques. To enable polar data signaling, the phase modulator needs to be complemented by an amplitude modulator, typically implemented as a digital power amplifier (DPA) in polar transmitters. The DPA is normally the most power hungry part of a transmitter, hence maintaining its power efficiency is utmost importance. Highly power-efficient DPAs are unfortunately prone to distortion. They typically exhibit AM-to-AM nonlinearity and AM-to-PM cross-talk that can severely degrade the output signal quality, increasing obviously the measured EVM. We make use of such a DPA (inverse class-D DPA) within the devised Polar TX, however, instead of simply cascading the DPA with the PLL, the DPA is in our system part of the closed subsampling loop. This means that the phase-error detection happens directly at the DPA output, instead of at the VCO output. The phase-error detector becomes in this way additionally sensitive to the modulation *amplitude*. This feature is exploited in two directions. First, since the AM signal accuracy can be tracked, we propose mechanisms that enable background AM-to-AM distortion cancellation by digital data predistortion. The second benefit is compression of AM-induced PM since the DPA effectively becomes part of the phase lock that filters output-referred errors. Thanks to the proposed environment which is capable of complete and independent self-linearization in continuously changing environments (PVT tracking), we demonstrate one of the first wireless polar architectures capable of synthesizing a 1024-QAM constellation with <-41 -dB EVM. The chip operates with 2.5/10-MHz modulation bandwidths consuming approximately 50 mW (the exact number depends on the desired output power) from a single 0.9-V supply and from a single 40-MHz input clock reference. The newly proposed environment

offers great spectral efficiency and an attractive solution for a low-cost bulk CMOS wireless digital transmitter.

5.2 Contributions

This book is built around three IC prototypes in 28-nm bulk CMOS that push limits of the art in digital frequency synthesis and polar transmitter design. Book contributions are summarized below.

The first IC prototype (Chap. 2) implements a Digital-to-Time converter (DTC)-based fractional-N subsampling PLL.

- We present how to enhance the low-jitter, divider-less PLL with a high-gain subsampling phase-error detection core for fractional-N synthesis using a DTC. A DTC is used as a phase modulator in the reference path of the PLL. The block exploits time-domain rather than voltage-domain signal processing, with very low quantization error excess and permits high (in-band) phase noise performance.
- We propose one of the first DTC implementations suitable for high-performance frequency synthesis in advanced digital nodes. Special care has been dedicated to the DTC linearity, phase noise performance, and supply noise rejection.

In the second IC prototype (Chap. 3), we propose a background-calibrated environment that ensures pristine fractional-N PLL output spectrum.

- We analyze and reveal typical distortion mechanisms in a DTC-based fractional-N frequency synthesizer.
- We develop digital algorithms that can run in the background as the PLL operates normally, to compensate for the nonlinearities in the phase comparison path (DTC) by pure digital predistortion. In this way, the analog subsampling PLL operates with virtually no performance gap between integer-N and fractional-N modes.
- We present the first subsampling PLL-based frequency/phase modulator. Highly accurate two-point phase injection is achieved thanks to the background-calibrated environment.

In the third IC prototype (Chap. 4), we present a digitized DTC-based subsampling PLL as part of a novel transmitter architecture, a digital subsampling polar TX.

- We present how to digitize an analog DTC-based subsampling PLL without compromise in phase noise performance and with all the benefits (area savings, and loop control) of a digital loop.
- We develop a Polar TX that uses an amplitude modulator (DPA) within the phase lock. The subsampling phase-error detection path becomes thus sensitive to both, amplitude and phase errors. This enables inherent AM-to-PM cross-

distortion compression. Moreover, we develop algorithms that use this property for automatic background calibration of AM and PM distortion within the loop. The developed system offers highly linear environment that permits very complex modulation schemes and pristine output spectrum.

5.3 Future Outlook

There are several areas that can be explored to improve the presented material and to build on top of it. A small summary of those is presented next.²

5.3.1 *Lock Time Optimization in a Subsampling PLL*

When the analog (Chaps. 2 and 3) or digital (Chap. 4) fractional-N subsampling PLL operates in lock, it does not effectively discriminate between channels that are integer-N apart. To cope with the issue, the systems make use of a frequency acquisition loop with a dead zone. This loop runs in parallel to the subsampling loop, ensuring that the VCO settles to the desired frequency before the phase lock (subsampling) takes over. The dead zone is in our designs set to $T_{REF}/2$, where T_{REF} stands for the period of the reference clock. This means that the frequency acquisition loop reacts only if the equivalent amount ($T_{REF}/2$) of phase error between the input and the output gets accumulated. The gain of the frequency error detector is set to *higher* values than the subsampling PD gain, so that it overrules any information from the subsampling loop in the presence of a frequency error. Once when the frequency is locked, the acquisition loop can be disabled. Locking process has not however been optimized in the proposed designs. The ratio of gains for optimal acquisition time, or dead-zone size optimization (potential power vs. performance trade-off) are still areas that can be explored. Moreover, an interesting aspect is PLLs behavior in the presence of a *lost lock*. Would it be possible to find an alternative to the power-consuming frequency acquisition loop so that the subsampling loop can operate independently? In the presented standalone ICs, the frequency acquisition loop can be disabled since there is no *pulling*. This scenario may change in situation where the PLL operates in the same environment with, e.g., an instantaneously switching power circuit nearby. Resilience to these effects (without necessity of running the frequency acquisition loop in parallel) could be a great added feature to the described PLLs.

²Summarizing future work hopefully helps with *letting go* of the book writing, for the curse of *it could always be better...* is not an easy one to confront. I hope that the paragraphs written here fall into a view of someone who can still get the opportunity to walk the paths that we had no time for.

5.3.2 Towards Higher Modulation Bandwidths and Better Out-of-Band Noise Rejection

The presented phase modulators of Chaps. 3 and 4 operate at 10-MHz highest modulation bandwidth. In an ideally matched two-point injection scheme, the phase modulation bandwidth (and the highest data rate) is only limited by the modulating DAC clock rate (Nyquist criteria), i.e., by the update speed with which the VCO changes its operating frequency. In the presented designs, phase/frequency (and later amplitude) modulation path operate at a 40-MHz reference clock rate. The 10-MHz bandwidth is used as a value safely below the theoretical 20-MHz limit ($F_{CLK}/2$).

The ratio between F_{CLK} and F_{MOD} (where F_{CLK} stands for the modulating DAC clock rate and F_{MOD} is the highest frequency component of the modulation data, respectively) is also relevant with respect to the out-of-band noise rejection. As the ratio increases, the rejection is improved, since higher order digital filtering (higher number of the FIR shaping coefficients) can be enabled. Moreover, just as in the case of any data converter, clocking at the higher rate pushes the (sinc filtered) aliases of the original data to higher frequencies.

To increase the modulation speed, or to improve the out-of-band noise rejection, the modulating DACs necessitate higher clock rate. This clock can be generated by division from the VCO operating frequency or completely independently, by reference frequency multiplication. In both cases, one can expect a raise in power consumption (and potentially a drop in system efficiency), since more circuitry needs to operate at higher speeds. Fortunately, since most of the described consumption is digital, the degradation is expected to lower with advance of technology. The former approach with VCO frequency division is easier with respect to added design effort, however, special care needs to be taken since a *modulated* VCO output is used to create the data clock. This lies in contrast to a simple situation when the clock is derived from the constant reference input (by frequency multiplication), as in the latter case. This approach has other negative aspects, for example, generating an additional frequency reference comes at a cost of power and area consumption, especially if extremely low jitter is targeted. Simply put, an additional PLL might be needed to clock the original PLL.

In context of the proposed SSPTX architecture (see Chap. 4), an additional issue that arises with increased modulation rate is the more pronounced AM-to-PM distortion. The SSPTX is indeed mostly resilient to modulation bandwidth increase since the principles of PM-to-PM and AM-to-AM calibration remain intact. This unfortunately does not hold for AM-to-PM cross-distortion suppression, since it mostly has effect within the PLL bandwidth (approximately 2.5 MHz). To overcome this, in theory, similar principles as proposed in Sect. 4.3.2.2 for AM-to-AM predistortion could be applied to overcome AM-to-PM. As discussed there, the SSPTX uses the amplitude-modulating block within the loop. This means that any AM-induced PM appears as a phase error within the system, at the PD output,

correlated to the instantaneously used AM code.³ This property can be used to systematically track influence of the AM on the phase error. Similarly as in the case of AM-to-AM distortion correction, this can be utilized to populate a predistortion LUT that modifies the original PM data sent to the VCO so that it compensates for the phase errors that are induced by the AM. With this algorithm enabled, the SSPTX becomes a fully calibrated environment that tracks PVT variations, even at large modulation bandwidths.

5.3.3 *Towards Higher TX Efficiency*

In Chap. 4, we present a novel polar TX architecture that is fully background calibrated with respect to AM-to-AM, PM-to-PM, and potentially AM-to-PM distortion (as discussed above). These benefits should enable a highly efficient system that exploits a nonlinear, switched mode DPA. In the presented design, we make use of a nonlinear, inverse-class D DPA for amplitude modulation. This architecture can, in theory, be 100% efficient; however, it had not been optimized in this proof of concept system for that purpose. An eventual redesign should involve efficiency optimization that exploits the fact that the nonlinearities are background calibrated at the system level. The overall optimization target can thus be the energy efficiency of the system, i.e., the value in Joule's per bit (that also grows with greater modulation complexity).

5.3.4 *Towards Other Modulation Schemes*

The developed algorithms for digital-to-transmitted signal calibration are not bound to any modulation principle and can be (in some form) applied to other modulators. For example, the FNSSPLL-based phase-modulator could be used as basis for a Frequency-Modulated Continuous Wave (FMCW) radar. This example is highlighted on purpose since the radar applications are recently gaining on importance [Yeo16, Wu14] with the extremely fast developments in the automotive industry (and environment-aware or autonomous driving). In FMCW applications, it is necessary to generate highly linear, several hundreds of MHz wide, slowly changing (period in microseconds) frequency chirps. This is a very challenging task in nonlinear capacitance-to-frequency conversion environments introduced by the typical LC-based VCOs. In this case, it becomes useful to predistort the modulation data, instead of investing effort into analog linearization techniques. The proposed mechanisms (see Sect. 4.3.1) can readily be applied in this context—

³The effect becomes increasingly complex in the presence of memory effects, i.e., if the induced error is not instantaneous but it depends on certain amount of previous AM samples.

potentially reducing the necessary calibration time (background-calibrated instead of measurement-based calibration) in changing environments (automotive is especially demanding with respect to PVT coverage). A note that must be taken here is that the developed algorithms manipulate the digital data to achieve linear *phase* response of the signal, while FMCW radars necessitate linear *frequency* response. The modeling and predistortion principles should therefore be accordingly adjusted.

References

- [Wu14] W. Wu, R.B. Staszewski, J.R. Long, A 56.4-to-63.4 GHz multi-rate all-digital fractional-N PLL for FMCW radar applications in 65 nm CMOS. *IEEE J. Solid State Circuits* **49**(5), 1081–1096 (2014)
- [Yeo16] H. Yeo, S. Ryu, Y. Lee, S. Son, J. Kim, 13.1 A 940MHz-bandwidth 28.8 μ s-period 8.9 GHz chirp frequency synthesizer PLL in 65nm CMOS for X-band FMCW radar applications, in *2016 IEEE International Solid-State Circuits Conference (ISSCC)* (IEEE, Piscataway, 2016), pp. 238–239

Index

A

All-Digital Phase Locked Loop (ADPLL), 7–8
Analog-to-digital converter (ADC), 107

B

Background calibration process, *see*
DTC-based subsampling PLL
(FNSSPLL)

C

Charge Pump (CP)-based phase/frequency
detector, 7–8

D

DAC, *see* Digital-to-analog converter
Delay-spread cancellation algorithm, 72–74
Design of time quantizers (TDCs), 8
Digital linearization techniques, 93
 AM-to-AM distortion
 G_{PD} estimation, 96–98
 nonlinear DPA, 98–100
 AM-to-PM distortion
 amplitude modulating block, 101
 DPA-produced phase deviations,
 102–103
 Matlab simulation environment,
 102–104
 measurement-based predistortion, 102
 phase-domain Matlab simulations, 100–101
 PM-to-PM distortion
 calibration speed, 96

compensation values, 95–96
detected phase error, 92, 94
INL, 94–96
input DTC/DAC codes, 92, 94
LUT correction coefficients, 94–96

Digitally-Controlled Oscillator (DCO), 7
Digital power amplifier (DPA)

 detection slope, 97
 phase-error detection principle, 87, 90
 PLL phase noise floor, 90–91

Digital subsampling PLL
 die micrograph, 113–114
 integer-N spur level, 114–115
 measured phase noise, 114
 VCO setting modes, 113–114

Digital subsampling polar transmitter
(SSPTX), 129–130
 accurate frequency acquisition, 91–93
 amplitude modulation, 90–91
 bandwidth modulation, 120–121
 block diagram, 86–87
 DC current consumption, 116, 121–122
 digital linearization techniques (*see* Digital
 linearization techniques)
 EVM, 116–120
 fractional-N subsampling PLL
 ADC quantization error transfer
 function, 88–89
 DTC quantization noise, 88
 high frequency VCO output, 86
 measured spectrum, 114–115
 overview of, 89, 92
 phase-error detection, 86–88
 sinusoidal DPA output, 86, 88

Digital subsampling polar transmitter (SSPTX)
 (cont.)

HRM performance, 121

inverse class-D DPA

- AM-to-AM background calibration algorithm, 110
- controlled parallel array, 110
- HRM, 112–113
- secondary switch bank, 111–112
- square wave driven switching pair, 110–111
- switching PA structure, 109
- zero voltage switching, 109–110, 112

measured predistortion curves, 116, 119

modulation bandwidth, 132–133

performance of, 122

phase/frequency modulation, 89–90

prototype's target, 91–92

simulated PLL performance, 93

subsampling path

- ADC, 107
- block diagram, 104
- DAC, 108–109
- performance of, 114–115
- phase-error detection process, 104–105
- sampler, 105–106
- signal amplifier, 106–107
- timing diagram, 104

TX output spectrum, 116–118

Digital-to-analog converter (DAC), 108–109

Digital-to-time converter (DTC), 18–19, 130–131

- code-dependent supply noise, 128
- comparator and output buffer, 39–40
- delay control block
 - architectural overview, 36–37
 - complementary delay element, 37–39
 - flicker noise contribution, 37–38
 - input buffers, 37
 - mismatch-based errors, 39
 - nMOS transistor, 37–38
 - output node, 37
 - sampling clock, 36
 - white noise contribution, 37–38

FNSSPLL, 36

FoM, 52–54

fractional spur, 45, 49

high-speed oscilloscope, 49–50

in-band phase noise level, 51–52

INL and DNL characteristics, 44–45

linearity, 36

nonlinearity, 31–32

offset and gain error, 30–31

performance characterization, 50

phase noise, 32

quantization noise, 28–29, 88

RMS jitter, 51

synthesized fractional frequency, 50–51

TDC, 30, 128

tunable regulated supply, 39–41

worst fractional spur level, 50–51

DNL, 44–45

DPA, *see* Digital power amplifier

DTC, *see* Digital-to-time converter

DTC-based subsampling PLL (FNSSPLL)

- alternative implementation, 70
- comparator offset compensation, 75–76
- digital background compensation, 69
- EVM, 79–80
- vs. fractional-N CMOS PLLs, 76, 78
- GMSK spectrum and constellation, 77, 79
- in-loop phase noise generators, 59
- measured spurious performance, 76–77
- operation of, 60–61
- output phase noise profile, 76
- prototype IC, 74–75
- Random-Jump mechanism
 - frequency synthesizer, 74–75
 - INL (*see* Integral nonlinearity)
 - quantization noise randomization, 62–63
- raw Error Sign output, 69
- sign extracting node *Vextract*, 68
- system overview, 74
- two-point phase modulator, 60
 - delay-spread cancellation, 72–74
 - fDAC calibration, 71–72
 - Nyquist–Shannon sampling theorem, 70–71
 - phase errors, 70
 - time-domain modulation operation, 70–71
 - wide bandwidths, 59
 - zero phase-offset condition, 69

DTC delay calculation method, 62–63

E

Error vector amplitude (EVM), 79–80

SSPTX bandwidth setting, 116, 119

supply level variation, 117, 120

verification, 116, 118

F

Figure of merit (FoM), 14, 17, 52–54
 Fractional-N subsampling phase-locked loop (FNSSPLL), 128
 ADC quantization error transfer function, 88–89
 addition of divider, 26
 analog core, 23
 architecture of, 32–33
 characteristics, 14–15
 digital computation, 27
 divider-less loop, 127
 DTC (*see* Digital-to-time converter)
 FoM, 14, 17, 52–54
 frequency-acquisition loop, 27, 43
 general system, 24
 high frequency VCO output, 86
 high-performance, 17
 implementation of, 26, 34–36
 LDO, 43
 lock time optimization, 131
 loop digitization, 23
 loop filter, 15
 LPF, 33
 measured phase noise
 $\Delta\Sigma$ modulator order, 47–48
 DTC gain mismatch, 47–48
 RMS jitter and integer- N jitter, 46–47
 Rohde & Schwartz FSQ26 spectrum analyzer, 47, 49
 worst-case fractional- N scenario, 46
 measured spectrum, 114–115
 operation of, 15
 overview of, 89, 92
 phase-error detection, 15–16, 86–88
 phase modulating element, 24–25
 phase noise generation, 15–16
 power-efficiency improvements, 17
 prototype chip, 43–44
 reference clock, 26
 sinusoidal DPA output, 86, 88
 time-domain analysis, 24–25
 type-two PLL's zero-pole separation, 16
 VCO, 41–42, 44–45
 voltage amplitude, 15
 zero-crossings, 27
 Frequency-Modulated Continuous Wave (FMCW) radar, 133–134

H

Harmonic-rejection-mixing (HRM), 110, 112–113, 121

I

Integral nonlinearity (INL)
 accurate predistortion, 64–65
 convergence time, 67–68
 digital gates, 67
 DTC, 44–45
 32-entry piece-wise-linearized equivalent, 67
 error-induced phase noise, 63–64
 input code and detected phase error, 65–66
 linear gain error, 68
 LMS algorithm, 65–66
 LUT, 64–65
 PM-to-PM distortion, 94–96
 PRBS, 62, 64
 simulation of, 66–67
 Inverse class-D DPA
 AM-to-AM background calibration
 algorithm, 110
 controlled parallel array, 110
 HRM, 112–113
 secondary switch bank, 111–112
 square wave driven switching pair, 110–111
 switching PA structure, 109
 zero voltage switching, 109–110, 112

L

Least mean square (LMS) algorithm, 30, 65–66
 Linearization techniques, 129
 Local oscillator (LO)
 Cartesian transceiver
 noise effect, 2–3
 receive mode, 1–2
 specifications, 2–3
 transmit mode, 1–2
 PLL (*see* Phase-locked loop)
 polar transmitter, 3–4
 Look-up table (LUT), 64–65
 Low-dropout regulator (LDO), 43
 Low Noise Amplifier (LNA), 1–2
 Low-pass filter (LPF), 24–25, 33

M

Metal–oxide–metal (MOM) capacitors, 34

N

Near-to-zero voltage sampling, 61

O

Open loop gain (G_{OL}), 10

P

Peak-to-Average Power Ratio (PAPR), 90
 Phase-error detector (PD), 5, 58
 Phase/Frequency Detection (PFD) mechanism, 91–93
 Phase-locked loop (PLL)
 ADPLL, 7–8
 background-calibrated phase/frequency modulator, 19
 classical analog, 5–6
 history of, 17–18
 mixed-signal environments
 division modulus, 6
 fractional- N frequency synthesis, 6–7
 frequency granularity, 7
 fundamental bandwidth *vs.* phase noise trade-off, 7
 pseudo-random dithering, 7
 VCO output frequency, 6
 phase noise
 digital architecture, 11
 dominant input referred noise, 12
 dominant output referred noise, 12
 integer- N and fractional- N frequency synthesis, 14
 noise generation, 13–14
 open loop reference, 11
 output phase noise spectra, 11–12
 phase-error detection gain, 12–13
 polar TX, 19
 PVT variations, 4
 small signal phase-domain model, 9–10
 subsampling loop
 characteristics, 14–15
 DTC, 18–19
 FOM, 14, 17
 high-performance, 17
 loop filter, 15
 operation of, 15
 phase-error detection, 15–16
 phase noise generation, 15–16

power-efficiency improvements, 17
 type-two PLL's zero-pole separation, 16
 voltage amplitude, 15

PLL-based phase modulation (PM)

in-loop and quantization noise, 57
 time-to-digital conversion, 59
 two-point injection, 58–59

Polar transmission, 4, 127, 133

Power Amplifier (PA), 1

Process-temperature-voltage (PVT) variations, 4

Proportional-integral (PI) characteristic, 9

Pseudo-random binary sequence (PRBS), 62–63

R

Radio Frequency (RF) carrier, 1

S

SSPTX, *see* Digital subsampling polar transmitter
 Successive Approximation Register (SAR), 106

T

Time-to-digital converter (TDC), 7–8, 30

V

Voltage-controlled oscillator (VCO)
 bottom-pinning bias technique, 42
 class-B structure of, 41–42
 fractional spur, 49
 measured tuning range, 44–45
 out-of-band noise rejection, 132
 phase noise filtering profile, 127
 setting modes, 113–114