OULU 2008
C 306

ACTA

UNIVERSITATIS OULUENSIS

Maciej Borkowski

DIGITAL Δ-Σ MODULATION

VARIABLE MODULUS AND TONAL BEHAVIOUR IN A FIXED-POINT DIGITAL ENVIRONMENT

FACULTY OF TECHNOLOGY,
DEPARTMENT OF ELECTRICAL AND INFORMATION ENGINEERING,
INFOTECH OULU,
UNIVERSITY OF OULU

ACTA UNIVERSITATIS OULUENSIS
C Technica 306

MACIEJ BORKOWSKI

DIGITAL Δ-Σ MODULATION

Variable modulus and tonal behaviour in
a fixed-point digital environment

Academic dissertation to be presented, with the assent of the Faculty of Technology of the University of Oulu, for public defence in Raahensali (Auditorium L10), Linnanmaa, on November 7th, 2008, at 12 noon

OULUN YLIOPISTO, OULU 2008

Copyright © 2008
Acta Univ. Oul. C 306, 2008

Supervised by
Professor Juha Kostamovaara

Reviewed by
Professor Michael Peter Kennedy
Professor Olli Vainio

ISBN 978-951-42-8909-5 (Paperback)
ISBN 978-951-42-8910-5 (PDF)
http://herkules.oulu.fi/isbn9789514289105/
ISSN 0355-3213 (Printed)
ISSN 1796-2226 (Online)
http://herkules.oulu.fi/issn03553213/

Cover design
Raimo Ahonen

OULU UNIVERSITY PRESS
OULU 2008

Borkowski, Maciej, Digital Δ-Σ Modulation. Variable modulus and tonal behaviour in a fixed-point digital environment

Faculty of Technology, Department of Electrical and Information Engineering, University of Oulu, P.O.Box 4500, FI-90014 University of Oulu, Finland; Infotech Oulu, University of Oulu, P.O.Box 4500, FI-90014 University of Oulu, Finland

Acta Univ. Oul. C 306, 2008
Oulu, Finland

Abstract

Digital delta-sigma modulators are used in a broad range of modern electronic sub-systems, including oversampled digital-to-analogue converters, class-D amplifiers and fractional-N frequency synthesizers.

This work addresses a well known problem of unwanted spurious tones in the modulator's output spectrum. When a delta-sigma modulator works with a constant input, the output signal can be periodic, where short periods lead to strong deterministic tones. In this work we propose means for guaranteeing that the output period will never be shorter than a prescribed minimum value for all constant inputs. This allows a relationship to be formulated between the modulator's bus width and the spurious-free range, thereby making it possible to trade output spectrum quality for hardware consumption.

The second problem addressed in this thesis is related to the finite accuracy of frequencies generated in delta-sigma fractional-N frequency synthesis. The synthesized frequencies are usually approximated with an accuracy that is dependent on the modulator's bus width. We propose a solution which allows frequencies to be generated exactly and removes the problem of a constant phase drift. This solution, which is applicable to a broad range of digital delta-sigma modulator architectures, replaces the traditionally used truncation quantizer with a variable modulus quantizer. The modulus, provided by a separate input, defines the denominator of the rational output mean.

The thesis concludes with a practical example of a delta-sigma modulator used in a fractional-N frequency synthesizer designed to meet the strict accuracy requirements of a GSM base station transceiver. Here we optimize and compare a traditional modulator and a variable modulus design in order to minimize hardware consumption. The example illustrates the use made of the relationship between the spurious-free range and the modulator's bus width, and the practical use of the variable modulus functionality.

Keywords: delta-sigma modulation, digital circuits, fixed point arithmetic, frequency synthesis, limit cycle

Acknowledgements

This thesis is based on research work carried out at the Electronics Laboratory of the Department of Electrical and Information Engineering, University of Oulu, during the years 2003-2008.

I would like to express my deepest gratitude to my supervisor, Professor Juha Kostamovaara, for his guidance and great example. I would like to thank Mr Tom Riley for introducing me to frequency synthesis and delta-sigma modulation, and Dr. Juha Häkkinen for his excellent guidance during my first years at the laboratory. Professor Timo Rakhonen has often offered me his expert's insight into many difficult problems in signal processing and telecommunication. I would like to thank my friends Sami Karvonen, and Harri Rapakko for their personal support, and also Professor Sergei Vainshtein for his friendship and valuable insights into the scientific world.

I wish to thank professors Olli Vainio and Michael Peter Kennedy for their exceptional work in examining this thesis, and Mr. Malcolm Hicks for revising the English of the manuscript. I would also like to thank to Kaveh Hosseini and Marko Neitola for their valuable comments.

I express my appreciation to all my co-workers and friends in the Electronics Laboratory for the great working atmosphere there, especially to Janne Aikio, Mikko Loikkanen, Ilkka Nissinen, Jussi Nissinen, Zheng Shu Feng and Jani Pehkonen, who was also a great companion in learning LaTeX. Finally I would like to thank our fantastic secretary Päivi and all my other colleagues and friends at the department and on the department staff.

I would like to thank my wife Paula for her love, support and encouragement during all the years I have spent in PhD research. I would like to thank my parents Joanna and Marian for being the best parents on earth, and my sister Ola and brother Tomek for their positive energy and support. I would also like to thank my wife's family who accepted me from the start and became my second family, here in Finland.

This work was supported financially by the Infotech Oulu Graduate School, the Nokia Foundation, the Riitta and Jorma J. Takanen Foundation, the Ulla Tuominen Foundation, Tekniikan edistämissäätiö and the Tauno Tönning Foundation, all of which are gratefully acknowledged.

This thesis has been written in LaTeX and compiled under MikTeX. All the line

figures were drawn using Xcircuit.

Abbreviations and symbols

AC	alternating current, in the context of a digital modulator denotes a varying input signal
ADPLL	all-digital phase-locked loop
Att	attenuation; a spectrum analyzer measurement parameter
BS	base station
BW	bandwidth
CMOS	complementary metal-oxide-semiconductor
CMQ	classical model of quantization
DAC	digital-to-analogue converter
DC	direct current, in the context of digital modulator denotes a constant input signal
DCO	digitally-controlled oscillator
DCR	direct-conversion receiver
DFS	discrete Fourier series
DFT	discrete Fourier transform
DSE	delta-sigma encoder
DSM	delta-sigma modulator (modulation)
DT	discrete time
EFM	error-feedback delta-sigma modulator architecture
FF	flip-flop
FPGA	field-programmable gate array
FSM	finite-state machine
GPRS	general packet radio service
GSM	global system for mobile communications
LC	limit cycle
LO	local oscillator
LSB	least significant bit
LSR	linear shift register
LUT	lookup table
MASH	multistage noise shaping delta-sigma modulator
MS	mobile station
MSB	most significant bit
Nsamp	number of samples; a spectrum analyzer measurement parameter
NTF	noise transfer function
PD	phase detector

PDF	probability density function
QTII	second quantization theory (formulated by B. Widrow)
RBW	resolution bandwidth; a spectrum analyzer measurement parameter
Ref	reference level; a spectrum analyzer measurement parameter
RF	radio frequency
rPDF	rectangular probability density function (dither)
SFR	spurious-free range
SNR	signal-to-noise ratio
SOC	system on a chip
SWT	sweep time; a spectrum analyzer measurement parameter
TI	Texas Instruments
tPDF	triangular probability density function (dither)
VBW	video bandwidth; a spectrum analyzer measurement parameter
VCO	voltage-controlled oscillator
VMDSM	variable modulus delta-sigma modulator
VMQ	variable modulus quantizer
WCDMA	wideband code division multiple access; a type of third generation cellular network
WLAN	wireless local area network

A_i	gain element within an error-feedback modulator
b	bus width
e	quantization error
f	frequency
f_k	discrete frequency
f_o	output frequency
f_r	reference frequency
f_s	sampling frequency
G	quantizer gain
I	integer part of the frequency division ratio in fractional-N frequency synthesis
I_i	initial condition within an error-feedback modulator
L	data path delay due to logic
L_s	sequence length
$M^{\mathrm{l}}(f_k)$	one-sided, discrete power spectrum based on the linear model; scaled for comparison with a measurement
$M^{\mathrm{s}}(f_k)$	one-sided, discrete power spectrum based on a simulation; scaled for comparison with a measurement

N	a coefficient used in frequency synthesis to obtained multiples of a reference frequency
n	discrete time instant
P	average power
$P(f_k)$	two-sided, discrete power spectrum
$P^{\mathrm{l}}(f_k)$	one-sided, discrete power spectrum based on a linear model
$P^{\mathrm{s}}(f_k)$	one-sided, discrete power spectrum based on a simulation
Q	modulus; denominator of the digital quantizer gain factor
q	quantizer output
R	data path delay due to routing
r	feedback signal in the error-feedback architecture
S	average power
$S(f)$	two-sided power spectral density
t	time
u	quantizer input
x	delta-sigma modulator input signal
y	delta-sigma modulator output signal

Δf	absolute accuracy of a frequency source
δf	relative accuracy of a frequency source
Δf_k	tone spacing due to the modulator sequence length
Δf_o	step size of a frequency synthesizer
Δ	quantization step
Δ_{ch}	radio channel spacing
σ_e^2	quantization error variance

$\mathbb{Z}$	set of integers

List of original articles

I Häkkinen J, Borkowski MJ & Kostamovaara J (2003) A PLL-based RF synthesizer test system. In IFIP VLSI-SOC 2003, International Conference on Very Large Scale Integration. Darmstadt, Germany, December 2003: 211–214.

II Borkowski MJ, Häkkinen J & Kostamovaara J (2003) A sigma-delta modulator development environment for fractional-N frequency synthesis. In IFIP VLSI-SOC 2003, International Conference on Very Large Scale Integration. Darmstadt, Germany, December 2003:50–54.

III Borkowski MJ & Kostamovaara J (2004) Post modulator filtering in Δ-Σ fractional-N frequency synthesis. In MWSCAS-04, The 2004 IEEE International Midwest Symposium on Circuits and Systems, volume I. Hiroshima, Japan, July 2004: 325–328.

IV Borkowski MJ & Kostamovaara J (2005) Spurious tone free digital delta-sigma modulator design for DC inputs. In ISCAS 2005, IEEE International Symposium on Circuits and Systems. Kobe, Japan, May 2005: 5601–5604.

V Borkowski MJ & Kostamovaara J (2006) On randomization of digital delta-sigma modulators with DC inputs. In ISCAS 2006, IEEE International Symposium on Circuits and Systems. Kos, Greece, May 2006: 3770–3773.

VI Borkowski MJ, Riley TAD, Häkkinen J & Kostamovaara J (2005) A practical Δ-Σ modulator design method based on periodical behavior analysis. IEEE Trans Circuits Syst II 52: 626–630.

VII Borkowski MJ & Kostamovaara J (2007) Variable modulus delta-sigma modulation in fractional-N frequency synthesis. Electronics Letters 25: 1399–1400.

Contents

1 Introduction

The communication electronics industry has experienced very rapid development since Nokia introduced its first highly successful mobile phone, the Mobira Cityman 900, on the mass market in 1987. Weighing just under 5 kg, this now 20-year-old grandfather of our present-day mobile phones, showed that people can communicate instantly over any distance wherever they are, and brought about a revolution in communication technology that coincided with the rapid development of personal computers and integrated circuits. Now everybody carries a tiny box weighing a mere 80 grams in thier pocket, a device which is still called a mobile phone but is in fact a small mobile personal computer. These are available everywhere in the world, and it is the global markets and global needs for them which dictate the current trends in their development.

The large computer and memory markets drive a "state-of-the-art" deep sub-micron CMOS fabrication technology. Traditional analogue and mixed mode radio electronics is constantly being pushed towards implementation in digital CMOS. The trend is to pack in more and more transistors of ever smaller size, and with rather poorer analogue properties [1–4], so that the emphasis is on developing digital techniques, which scale better with technology and can be reused easily and ported from process to process. These trends converge in the concept of a system on a chip (SOC), where all the components of a computer system and various electronic sub-systems are integrated into a single chip [3, 4]. A flexible, multi-standard radio transceiver is likely to become a dedicated radio processor on board a SOC computer [5]. The trend is continuing towards digitizing radios, moving closer and closer towards the antenna, and this is exerting a growing pressure on the blocks which convert the signals from the analogue to the digital world. This is the domain of delta-sigma data converters in the field of wireless transceivers [6].

Oversampled delta-sigma modulators (DSM) were invented 45 years ago [7, 8], and were originally used in analogue-to-digital conversion [9–11], and later in digital-to-analogue conversion [12–14]. Either because they appeared historically earlier, or because they constitute a broader class of systems, the analogue DSMs have received much more attention in the literature [15–17]. The situation has changed fairly recently, however, with the introduction of digital DSMs for radio transceiver applications. Digital DSMs are now found in applications which contribute to greater integration and

digitization of the radio front-end, including oversampled digital-to-analogue converters (DACs), mismatch shaping converters and, most notably, delta-sigma fractional-N frequency synthesis [6, 18, 19].

This work studies two properties of digital DSMs which are directly related to fixed-point digital implementation. The first is tonal behaviour, which, due to the discrete, limited DSM state space, is an even more pronounced problem in digital DSMs than in analogue ones. Our analysis and suggested solution allow DSM spectrum quality to be traded for the number of bits required to implement the modulator. The second problem studied here is related to the foundations of fixed-point digital systems, in which all arithmetic operations are performed with the modulus being a power of two. We suggest a general modification to all DSM topologies which would allow them to operate with any integer as the modulus. This new, more general class of digital DSMs would be better suited for multi-standard radio transceiver applications.

The remaining part of this chapter will focus on applications of digital DSMs in wireless transceivers. Selected applications are characterized briefly in Sec. 1.2 with special attention paid to the types of DSM input signals. Digital DSMs with DC inputs can be represented as digital signal generators and considered separately from the more general class of modulators working with arbitrary inputs. The main application of interest, fractional-N frequency synthesis is considered in greater detail in Sec. 1.3, which will highlight the most common problems arising from the use of delta-sigma technology and the existing solutions. Finally, Sec. 1.4 will list the original contributions contained in this thesis and Sec. 1.5 will present an outline of the thesis. The literature review presented in this chapter, like all the other reviews included in this thesis, is based on scientific journals, conference proceedings and available books, which means that our view of the "state-of-the-art" is dominated by academic research. The achievements of industry are covered only to the extent to which information on them is available in the scientific literature.

1.1 A short introduction to digital delta-sigma modulation

Digital DSMs are very useful at the interface between the digital and analogue worlds. An interpolator followed by a digital DSM, a low resolution DAC and an analogue low pass filter together form an oversampled DAC. The same DSM controlling a

multi-modulus divider in a phase-locked loop (PLL) forms a fractional-N frequency synthesizer, which can also be regarded as a digital-to-phase or digital-to-frequency converter. Delta-sigma data converters achieve high signal-to-noise ratios (SNR) without using high precision analogue circuits. The conversion is performed at data rates much higher than the signal bandwidth, but with reduced resolution. This makes delta-sigma data converters particularly well suited for VLSI technology optimized for high-speed digital circuits [6].

A digital DSM is a nonlinear system which converts a high resolution discrete time signal x into a low resolution discrete time signal y. Strictly speaking the digital DSM coarsely re-quantizes the input signal x, although the literature on the subject uses a common term *quantization* in the context of both digital and analogue DSMs. Quantization alone unavoidably reduces the SNR, and in order to preserve a high SNR the DSM filters the quantization noise from the neighbourhood of the input signal in a process called *noise-shaping*. The principle of operation of a digital DSM is presented in Fig. 1. The digital DSM is applied to a signal which is already *oversampled* and occupies only a small bandwidth (BW) between zero and half of the sampling frequency. Proper preparation of the input signal x is the task of an interpolation filter, usually preceding the DSM. The signal path in a digital DSM data converter is described in greater detail in [20, Ch. 9, Ch. 10] and [17, p. 221].

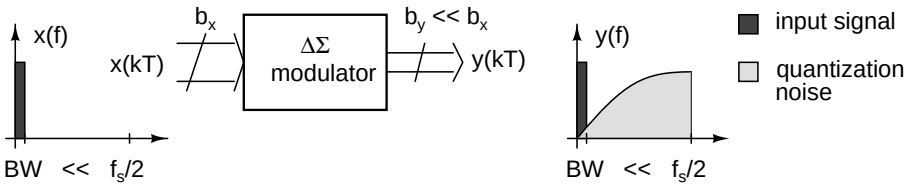

Fig 1. Principle of operation of a digital delta-sigma modulator.

A *stable* DSM is characterized by regular behaviour, as presented in Fig. 1, where the output signal is the sum of the input signal and properly shaped quantization noise. Many DSMs are only conditionally stable, however, in that stable behaviour occurs when the input signal and the DSM state variables are bounded within certain ranges. One of the external signs of instability is a rapid decrease in the SNR within the bandwidth of interest (BW). In the field of analogue DSMs the appearance of short limit cycles is also considered a sign of instability. If a DSM is operating in a short limit cycle, the

quantization noise power will be concentrated into a few strong spurious tones. DSM stability is a broad and important subject which is left beyond the scope of this thesis. A good overview of the problem of DSM stability is presented in [16, Ch. 4] and [21, Ch. 5]. This thesis makes two contributions to the field of digital DSMs, which are outlined in Sec. 1.4 and presented in Ch. 3 and Ch. 4.

1.2 Digital delta-sigma modulation in wireless transceivers

Data converters provide an interface between the analogue and digital worlds. Conversion is not limited to the voltage or current domains, but particularly in radio applications it includes the time, frequency and phase domains. A typical example of this kind of converter is a delta-sigma fractional-N frequency synthesizer, which converts a baseband digital signal into a phase or frequency-modulated RF signal. Overviews of delta-sigma data converters in wireless transceivers are presented in [6, 22]. This section is focused on applications of purely digital modulators and provides a short overview that emphasizes the type of DSM input signal: constant or arbitrary. It also considers typical solutions to the tonal problems found in wireless transceivers.

1.2.1 Classical transceivers based on a local oscillator

The first group of radio systems which benefit from delta-sigma techniques are classical radio topologies. The classical radio transmission and reception scheme based on the use of a local oscillator (LO) and a mixer dates back to 1914, when the first superheterodyne receiver was invented [23, 24] and the principle has remained in common use in RF microelectronics up to the present day [22, 25]. The direct-conversion transceiver, which uses just one local oscillator, has become a popular choice in mobile communication as it combines a high level of integration with low power consumption [24, 26]. Direct conversion receivers and transmitters have been used successfully for constructing fully integrated Bluetooth, and most notably multi-standard GSM, WCDMA and WLAN radios [27–31].

A block diagram of a direct-conversion receiver (DCR) is depicted in Fig. 2. The RF signal coming from the antenna is amplified by a low-noise amplifier and downconverted directly to baseband in-phase (I) and quadrature (Q) signals. The low pass filters for

channel selection and gain control are therefore implemented in the baseband. The apparent simplicity of the block diagram does not indicate the real implementation difficulties, involving DC offset, $1/f$ noise, the generating of I and Q phases and gain matching. The design effort required to build a DCR is considerable [24–26], and therefore designers tend to choose other classical solutions in many cases, including double conversion or low IF, which in some cases also lead to high levels of integration [32–35].

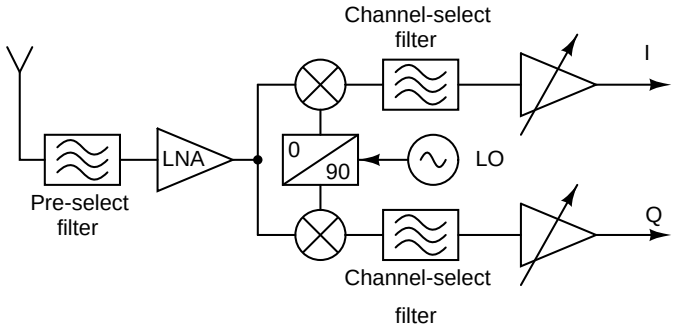

Fig 2. A direct-conversion receiver.

What is important from the point of view of this work is that the majority of integrated radios are built using the classical approach based on a LO and mixer. The LO is built using a PLL frequency synthesizer which generates a single, stable frequency. In highly integrated solutions, particularly in multi-standard radios the fractional-N frequency synthesizer is currently a common choice [28–31, 36]. Such synthesizers are based on a DSMs working with DC inputs, but this solution is particularly prone to producing unwanted, spurious tones [15–17, 19]. The most common approach to eliminating DSM spurious tones is based on dithering [29, 37–40]. A more thorough discussion of tonal behaviour and the available solutions is presented in Sec. 1.3.3.

1.2.2 A transmitter based on a modulated frequency synthesizer

A delta-sigma fractional-N frequency synthesizer with a DC input works as a LO in a direct conversion transceiver, and when the input to the DSM is a transmit data

stream, the synthesizer can be used as a stand-alone transmitter capable of phase and frequency modulation [41]. Such a transmitter can be regarded as a digital-to-phase or digital-to-frequency converter. This approach increases the level of digitization of the radio transmitter. Modulating a PLL synthesizer involves a built-in trade-off between the data rate and the bit-error rate. The data rate can be increased by increasing the loop bandwidth, but this automatically increases the noise and consequently the bit-error rate.

This inherent limitation can be addressed by introducing a pre-emphasis filter to compensate for the PLL transfer function [42–44]. As the delta-sigma modulator contributes to the overall noise, special techniques have been invented to minimize this source of noise. One possibility is to use a low-order DSM, which produces a minimum amount of the out-of-band quantization noise, which can later be subtracted from the PLL using special compensation techniques [45]. Another way of reducing the DSM-induced noise is to modulate the PLL division ratio at the level of fractions of the VCO cycle [46]. Among the difficulties encountered in fractional-N synthesizer design are fractional spurious tones [47] and high linearity requirements. Any nonlinearity within the PLL loop will increase the in-band noise through intermodulation mechanisms [48]. Despite the inherent design challenges, this technique has been successfully targeted toward various standards including: ISM [49, 50], bluetooth [40], GSM [39], multi-band GSM-GPRS [43] and GSM,GPRS and WCDMA [51].

From the point of view of this thesis, it is worth noting that this application relies on a digital DSM working with arbitrary AC inputs. Even though the modulator works with active inputs, some authors have addressed the problem of spurious tones by choosing dithering as a solution [39, 40]. The problem of the contribution of a DSM to in-band noise has also been addressed by means of a specially tailored noise transfer function [37, 52].

1.2.3 All-digital transmitters

A PLL transmitter can also be realized using exclusively digital components and, according to the classification introduced in [53], converted to all-digital PLL (ADPLL). A critical part for an ADPLL in the context of an RF transmitter is the digitally-controlled oscillator (DCO). One such DCO has been successfully developed by a research group at Texas Instruments (TI) [54, 55] and has led to fully digital transmitters for Bluetooth and GSM [56–58].

The TI DCO is constructed as an LC oscillator with digital tuning, where the total

capacitance of the oscillator is split into a bank of small, selectable capacitors [54, 55]. Despite the small size of the capacitors, of the order of tens of attofarads, the frequency step in the 2 GHz range is insufficient for mobile applications [5, 54]. Therefore, the DCO uses a digital DSM to increase the resolution when the ADPLL is operating in tracking mode. A group of capacitors is switched at high frequency using the DSM output signal. The mean value of the stream is in the range between zero and one, and effectively gives the DCO fractional resolution. The principle of operation is similar to fractional-N frequency synthesis, described in the following sections. The digital DSM works with arbitrary inputs and and has a resolution of up to 10-bits. Since it operates at a frequency as high as 600 MHz, it must be carefully optimized for speed and power consumption [54].

1.3 Fractional-N frequency synthesis

This section introduces the fundamental operation of a fractional-N frequency synthesizer, a technique belongs to the PLL frequency synthesis family. PLL frequency synthesizers are most commonly used as LOs in radio transmitters and receivers. A fractional-N frequency synthesizer can be used as a very flexible, fast-switching LO, or can form a very compact transmitter when directly modulated. A thorough description of fractional-N frequency synthesis and other frequency synthesis techniques can be found in the basic literature [18, 19, 47, 53]. This chapter provides a general introduction to the fractional-N technique with special emphasis on the role of the digital DSM.

1.3.1 Fundamental operation

A general block diagram of a PLL frequency synthesizer is depicted in Fig. 3. The basic PLL is composed of a VCO, phase detector (PD) and loop filter. The PLL is a synchronization system in which the output frequency f_o tracks the input frequency f_r, also known as the reference frequency in the context of frequency synthesis. The loop tries to synchronize the phases of the output and the reference signals, hence the name phase-locked loop. The block diagram in Fig. 3 also includes a frequency divider, which plays a key role in frequency synthesis. The output frequency f_o is divided by N before it is compared with f_r. The synchronization mechanism ensures that the output

frequency f_o becomes an integer multiple of the frequency f_r:

$$f_o = N \cdot f_r. \tag{1}$$

This is the fundamental principle of an integer-N PLL frequency synthesizer, which uses a stable reference frequency f_r and produces multiples which are synchronized in phase [18, 19, 47, 53].

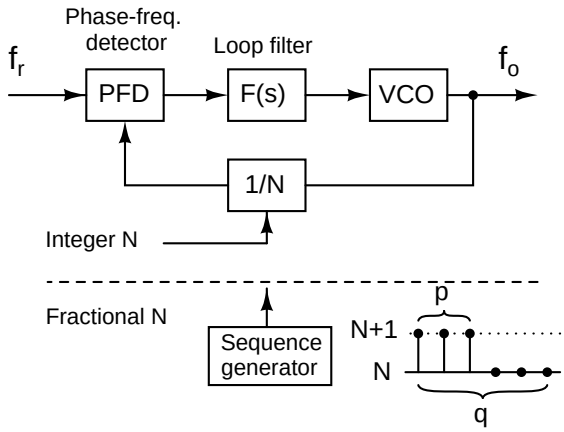

Fig 3. A PLL frequency synthesizer.

An integer-N frequency synthesizer has several fundamental limitations, however. First of all, it can only generate a set of frequencies which are separated by integer multiples of f_r. If a radio standard requires narrow channel spacing, the reference frequency must be low, which results in many undesirable effects. The PLL's loop bandwidth must be significantly lower than the reference frequency f_r to prevent the reference signal from feeding through to the VCO [19]. The smaller the loop bandwidth, the larger the settling time. Secondly, with low f_r and f_o in the GHz range, the division ratio N must be high. This adversely affects the synthesizer phase noise, as the in-band phase noise is proportional to $20\log(N)$ [19]. It has been recognized that the integer division ratio N is a fundamental bottleneck in this technique and methods have been invented to provide a fractional division ratio.

Fractional-N frequency synthesis is based on the idea of switching the division ratio between two or more integer values, so that the average ratio is a fraction [18, 47, 53]. This principle is depicted in Fig. 3. Suppose that the divider is capable of dividing by N,

or $N+1$ and the division ratio is controlled by a sequence generator. The generator produces a repeating sequence of length q which controls the divider so that division by $N+1$ occurs p times and division by N occurs $q-p$ times. The average output frequency can therefore be calculated as (2). This is the fundamental principle of fractional-N frequency synthesis.

$$f_o = f_r \cdot \left(\frac{N(q-p)+(N+1)p}{q} \right) = f_r \cdot \left(N + \frac{p}{q} \right) \tag{2}$$

Fractional-N frequency synthesis is similar to an earlier technique called digiphase [59]. Both techniques and the historical background to fractional-N synthesis are presented in [18, 47]. We will concentrate here on what has evolved to become the mainstream fractional-N architecture, the content of the sequence generator box in Fig. 3 and its impact on synthesizer performance.

1.3.2 Delta-sigma modulation in fractional-N

Although the theoretical foundations of fractional-N frequency synthesis, as presented in the previous section, appear to be simple in principle, this synthesis technique has been studied intensively up to the present day to find out exactly how the sequence generator depicted in Fig. 3 should be constructed, what is the best way to connect it to the PLL loop and how the negative side effects of modulating the frequency division ratio can be mitigated. In the early days of the method the sequence generator was implemented as a simple accumulator [60]. It is in fact a very important system, which may explain the majority of the problems and solutions and also the trends that have prevailed in the most recent research in this field. An accumulator with a constant input controls a frequency synthesizer and causes a periodic phase ramp at the divider output. Such a regular sawtooth signal contains most of its power in the first few odd harmonics of the fundamental frequency. This gives rise to unwanted spurious tones at the synthesizer output [18, 47]. Early solutions to the problem of spurious tones included compensation techniques designed to reduce the effects of the phase ramp. The integrated content of the accumulator, which is a measure of the phase ramp, can be used for compensation using a DAC [18, 47, 53, 60]. A similar technique, based on a specially designed PFD/DAC structure, has recently been reported in the context of direct modulation [45].

A major breakthrough in the fractional-N technique took place when it was noticed that an accumulator is in fact a first-order DSM [61, 62]. This suggested that instead of

compensating for the phase ramp using a DAC, it is possible to increase the order of the DSM. A higher-order DSM produces a more randomized output signal and therefore the troublesome fractional ramp within the PLL loop should never appear. From that point on, PLL-frequency synthesizers of the fractional-type were able to benefit from the theory of oversampled delta-sigma data conversion systems [15]. This marked the rise of the delta-sigma fractional-N frequency synthesizer family, a typical representative of which is depicted in Fig. 4.

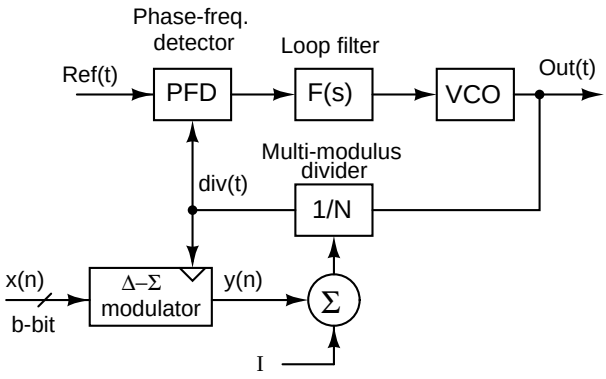

Fig 4. A delta-sigma fractional-N frequency synthesizer.

The fundamental benefit of the fractional-N technique lies in the improved resolution, which allows for a higher reference frequency f_r and a lower division ratio N. The fractional N is usually the sum of a fixed integer part I and a fractional part $x/2^b$, where x is the DSM input and b is the input bus width. The synthesizer output frequency can be expressed as:

$$f_o = f_r \cdot N = f_r \cdot \left(I + \frac{x}{2^b} \right). \tag{3}$$

This equation shows that the desired frequency can be approximated with an accuracy that depends on the DSM bus width b. The DSM used in fractional-N frequency synthesis operates with the reference frequency as a clock signal. The higher f_r is, the more power is consumed by the DSM. Power and area consumption are very important, and therefore compact, fast topologies such as multistage noise shaping architecture (MASH) have become very popular [19]. In a digital DSM, bus width truncation often plays the role of a quantizer. This is the easiest way of implementing the quantizer in a digital environment, as it requires no hardware. The truncation quantizer in fact

performs the fixed point division modulo 2^b, which is the reason why the fractional coefficient in (3) is based on the power of two.

This limitation can be removed by applying an accumulator of programmable size [19]. A variable-size accumulator is equivalent to a first-order variable-modulus DSM. In this work we generalize this concept to higher-order single and multi-bit DSMs. The concept of a variable modulus quantizer (VMQ), which converts any digital DSM into a variable modulus DSM (VMDSM) that is capable of generating any rational fraction, is introduced in Ch. 4. This type of modulator allows precise sequences of frequencies to be generated for a given telecommunication standard. A dynamically changed modulus allows a synthesizer to switch easily between different standards in a multi-mode transceiver. A fractional-N frequency synthesizer based on a VMDSM can generate exact frequencies instead of approximating them as in (3).

1.3.3 The spurious tones problem

As was stated in the short introduction in Sec. 1.1, a DSM quantizes the input signal and the resulting quantization noise is conveniently filtered away from the neighbourhood of that input signal. It is very important in radio applications that the spectrum of the filtered quantization noise should be smooth and free from dominant spurious tones. Unfortunately, all DSMs have the potential to produce highly correlated noise under certain conditions; this will be discussed in greater detail in Ch. 2. It has been observed that higher-order modulators produce better, more randomized quantization noise than low-order modulators. Spurious tones are also more likely to occur with slowly varying or DC inputs. Solutions to this problem vary depending on the DSM input, the particular application and whether the DSM is implemented in the digital or in the analogue domain [15–17, 19].

In transmitters using direct synthesizer modulation, as described in Sec. 1.2.2, the problem of spurious tones can be solved using compensation [40, 45, 63]. Such applications typically use a low-order DSM to limit excess out-of-band quantization noise and allow for a wider synthesizer loop bandwidth. A low order DSM and accumulator in particular will inherently produce highly correlated quantization noise, which can be measured and subtracted within the synthesizer loop through cancelation circuitry [40, 45, 63].

One of the most common ways of addressing the tonal problem without using special cancellation techniques is called dithering. This is based on injecting an additional

pseudo-random signal into the DSM loop or adding it to the DSM input signal [16, 17]. When the dither generator is turned off, even higher-order DSMs will produce very distinct tones [38, 39, 64]. The use of dithering has been reported in a number of transceiver designs [29, 37–40]. Criticisms of dithering are usually related to the higher noise floor, possible problems with DSM stability and the fact that it is difficult to guarantee that simple (hardware-efficient) dither generators will always remove the spurious tones.

A number of randomization methods exist which are applicable specifically to digital DSMs working with DC inputs. These are based on loading predefined initial conditions [18, 19, 41, 49, 65–69], setting the LSB of the input signal [61, 70] and using a prime modulus [71] or prime feedback [72]. The method of initial conditions is considered in detail in Ch. 3 of this thesis, where it is shown that applying predefined initial conditions allows precise control of the sequence length and consequently the spurious-free range (SFR). The approach proposed here quantifies the relationship between the SFR and the amount of DSM hardware.

1.4 Contribution

The literature and our general understanding of digital DSMs has been greatly affected by the legacy of the theory developed for analogue modulators. Until now, the theory concerning purely digital DSMs has been limited. Motivated by this fact, combined with their widespread use in wireless transceiver technology, we set out to study two fundamental properties of digital DSMs.

The first contribution of this thesis addresses DSM tonal behaviour in the case of the most problematic DC inputs. Such modulators are used in frequency synthesizers working as local oscillators in homodyne and heterodyne transceivers (see Sec. 1.2.1). A digital DSM working with a DC input is a finite state machine, and limit cycles and tonal behaviour can be regarded as being among its fundamental properties. It will be shown here that the lengths of the limit cycles can be controlled precisely by applying predefined initial conditions and scaling the DSM buses. As a result, the quantization noise power can be distributed over the required number of discrete tones in the DSM output spectrum. This contribution is important, because it shows a clear relationship between the amount of hardware necessary to implement the DSM and the quality of the spectrum.

The second contribution of the thesis addresses the use of digital delta-sigma

modulators in multi-standard transceivers. We explore the possibility of implementing DSMs in a digital environment with an arbitrary modulus quantizer. We propose a digital quantizer topology which converts any digital DSM into a variable modulus DSM having two inputs, one for the input data and the other for the modulus. A variable modulus DSM can generate arbitrary rational fractions defined by the data input and modulus input. A fractional-N synthesizer based on a variable modulus DSM can easily be adapted to various telecommunication standards. By selecting the modulus the synthesizer can be optimized for a chosen standard and can generate prescribed sequences of frequencies.

1.5 Overview of the thesis

The remaining part of the thesis is organized in the following way.

Chapter 2 reviews the classical theory of quantization. Under certain conditions quantization noise can be modelled as additive white noise that enables linear DSM analysis. We will review the conditions under which the classical model fails. When this happens the DSM reveals the negative effects of nonlinearity; namely limit cycles and tonal behaviour. This chapter reviews these two problems in the context of analogue and digital DSMs and summarizes the most common methods of addressing them in both domains. Digital DSMs working with DC inputs constitute a special case in which the spectrum is inherently tonal. Such modulators always generate limit cycles and the only unknown remaining is their length.

Chapter 3 explores the possibility of controlling the lengths of limit cycles in digital DSMs with DC inputs. Such modulators can be regarded as delta-sigma encoders which encode a digital input word into an output sequence of the desired length. This chapter presents a linear model for a delta-sigma encoder, taking into account the periodically correlated quantization error signal, and compares the output of this model with simulated and measured encoder spectra. A methodology is presented which allows the sequence length of a practical encoder to be controlled and a minimum hardware implementation (bus width) to be selected in order to satisfy given spectral requirements. The solutions suggested here lead to smooth and spurious tone-free spectra for all DC inputs.

Chapter 4 generalizes the concept of coarse quantization in the digital domain. A digital quantizer can be modelled as an arithmetic divider, which leads to the notion of a variable modulus quantizer (VMQ), having two inputs: one for the quantized signal and

the other for the modulus - which effectively sets the denominator for the arithmetic division. Practical implementations are proposed for single-bit and multi-bit variable modulus quantizers. The average output of a DSM with a VMQ can be any rational fraction, as set by the DSM DC input and the modulus. This functionality is useful in the context of multi-standard wireless transceivers.

Chapter 5 shows in practice how to use the knowledge presented in chapters 3 and 4. It presents selected aspects of DSM design for a frequency synthesizer intended to meet GSM base station specifications. The selected example illustrates how the required frequency resolution and spectral quality affect the selection of a traditional or variable modulus DSM, and how to select minimum hardware implementations in both cases.

2 Quantization noise in delta-sigma modulation

Delta-sigma modulators are widely used in ADCs [9–11], DACs [12–14] and radio communication [6]. They play an important role in data conversion systems as they eliminate the need for high amplitude resolution by using more bandwidth instead [16, 17, 73]. This is exactly in line with the current trend in microelectronics, where the accuracy of analogue components is constantly deteriorating while operating speed is increasing. DSMs constitute a very active research field that has grown to a respectable size[1] since the introduction of the technique 45 years ago [7, 8]. The DSM is built up around one or more coarse quantizers, see Fig. 5, which makes it a highly nonlinear dynamic system, and therefore one that is difficult to analyse in strict mathematical terms. Delta-sigma modulators inherently add quantization noise to the input signal and spread it over the increased bandwidth. One of the important recurring problems mentioned in the literature is the tendency for DSMs to generate unwanted tones in the frequency domain and limit cycles in the time domain.

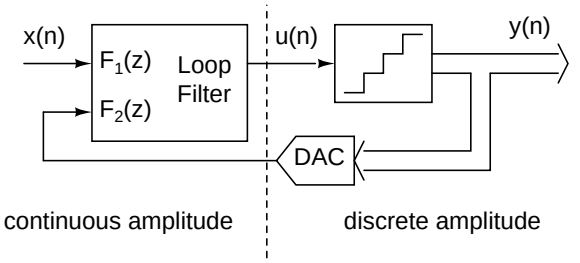

Fig 5. A discrete-time analogue DSM.

This chapter is devoted to quantization noise. The review of the literature will highlight differences between digital and analogue modulators, which are important from the point of view of the contribution of this thesis as presented in the subsequent chapters. The

[1] A literature search for DSMs in the IEEE Xplore database returned the following results. Search by title: 1278 papers, including 352 in journals. Search by index terms: 2842 papers including 806 in journals.
Search expression 1: ((delta<and>sigma<and>modul*)<in>ti),
Search expression 2: ((delta<and>sigma<and>modul*)<in>de)

term *analogue* will be taken here to refer to discrete-time (DT), continuous-amplitude DSMs, while the term *digital* will refer to DT, discrete-amplitude DSMs. Continuous-time DSMs lie beyond the scope of this thesis. This chapter is organized in the following way. We will first consider quantization alone without DSM feedback. Sec. 2.1 will review the classical model of quantization with special attention to the conditions governing its applicability. When these applicability conditions are met, the DSM can be subjected to simple linear systems analysis (see [16, 17, 73]), but if they are violated, the DSM will display unwanted tonal behaviour, as discussed further in Sec. 2.2.

2.1 Quantization

Although signals quantized both in time and magnitude have been known for a long time in the art of communication, it is the introduction of digital systems and pulse-code modulation that has triggered more intensive study of quantized signals [74]. A quantizer is a central point of any DSM, the main source of nonlinearity and the reason why DSM analysis is so difficult. The principles of quantization and the properties of quantization noise were studied long before DSMs appeared.

Quantization can be defined as the division of a quantity into a discrete number of small parts, often assumed to be integral multiples of a common quantity [75]. A quantizer is a central part of an analogue-to-digital converter, where it maps a continuous input amplitude signal into a discrete number of digital steps. Quantization, or strictly speaking re-quantization, can also be performed in the digital domain to reduce the signal resolution. Quantization in the digital domain is covered in greater detail in Ch. 4. Quantization is usually considered to be a memoryless, time-invariant, nonlinear operation. There are three main types of quantizer which are identified by the way in which the input-output characteristic crosses the zero point of a coordinate system. Fig. 6 show examples of multi-bit *mid-rise* and *mid-tread* quantizers. The third type, shown in Fig. 25, which is used primarily in the digital domain, is known as a *truncation quantizer*.

It is convenient to represent a quantizer as a combination of a linear gain G and an added quantization error e:

$$q = Gu + e. \tag{4}$$

This representation attributes all the nonlinearity to the quantization error e, while G is an important parameter from the point of view of stability and modulator dynamics. As

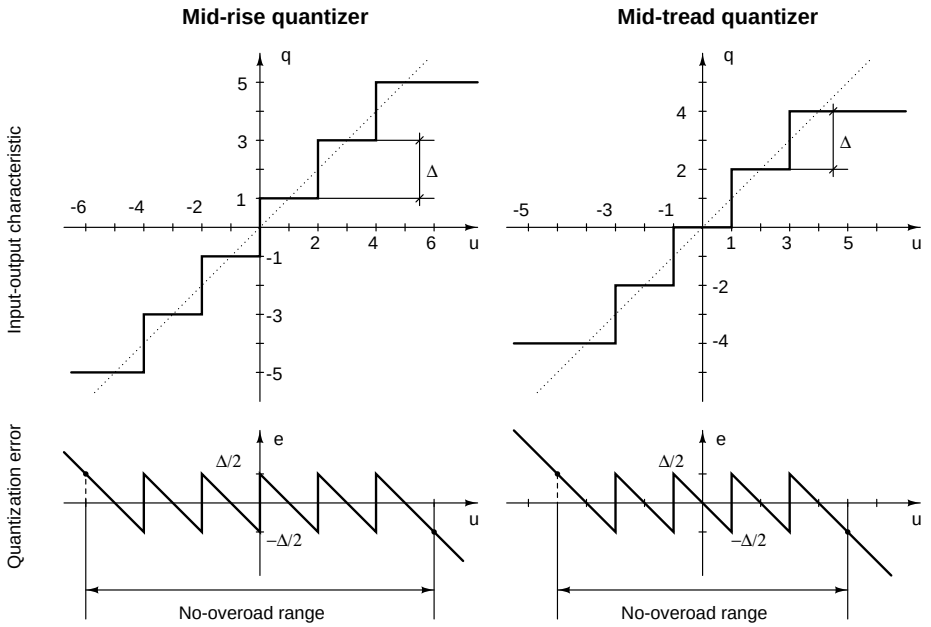

Fig 6. An example of a mid-tread and a mid-rise quantizer.

long as the quantizer input u remains in the *non-overload range* the quantization error e is bounded in the range from $-\Delta/2$ to $\Delta/2$ (see Fig. 6). The non-overload range criterion can be used to determine the maximum available range for the input signal u. The representation (4), commonly referred to as an additive noise model of quantization, as shown here in Fig. 7, is in fact not a model in the strict sense, but just a convenient way to define the quantization error e. Modelling begins only when simplifying assumptions are made about the statistical properties of e, as discussed in the following section.

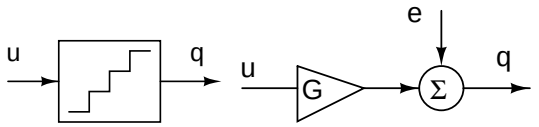

(a) Quantizer (b) Additive noise model

Fig 7. The additive noise model of quantization.

2.1.1 The classical model of quantization

Although the relationships between the quantizer input u, output q and quantization error e are purely deterministic, quantization often produces similar effects to those of an additive independent source of noise. This observation led Bennett in [74] to postulate what is now known as the classical model of quantization (CMQ). This simple and very useful model has three fundamental properties:

1. e is statistically independent of the input signal u
2. e is uniformly distributed in $[-\Delta/2, \Delta/2]$
3. e is stationary noise with a flat power spectral density (its autocorrelation is the Dirac delta function)

We have adopted the term CMQ in the following [76], but the same model has also been referred to as "additive white-noise approximation" [16] or "pseudo-quantization noise" in the statistical theory of quantization [77].

Property 1 allows the purely deterministic quantizer to be replaced with a stochastic linear system. This simplifies the analysis, as it permits power spectra analysis without the additional complexity of cross-terms between the input u and the quantization noise e [16, p. 47], and ensures that the quantizer's output power spectrum can be computed as the sum of the input power spectrum and the quantization error power spectrum.

Property 2, that the quantization noise error is uniformly distributed in the range $[-\Delta/2, \Delta/2]$, as shown in Fig. 8, effectively determines the average quantization error power. Any signal with the same amplitude distribution, regardless of its frequency content, will have the same variance. If the mean of the signal is zero, as in Fig. 8, then the average quantization noise power power S_e will be equal to the variance σ_e^2 [20, p. 755], and can be expressed by the familiar formula:

$$S_e = \sigma_e^2 = \int_{-\Delta/2}^{\Delta/2} e^2 p(e)\,\mathrm{d}e = \frac{1}{\Delta}\int_{-\Delta/2}^{\Delta/2} e^2\,\mathrm{d}e = \frac{\Delta^2}{12}. \tag{5}$$

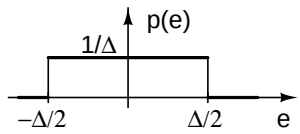

Fig 8. Probability density function of the quantization error.

Property 3 determines the time and frequency properties of the quantization error, stating that the quantization error signal has a flat power spectral density, as shown for a sampled quantization error in Fig. 9. In two-sided representation the power spectral density for the CMQ model is equal to:

$$S_e(f) = \frac{\Delta^2}{12 f_s}, \tag{6}$$

where f_s is the sampling frequency. When integrated from $-f_s/2$ to $f_s/2$, this gives the average quantization noise power:

$$S_e = \int_{-f_s/2}^{f_s/2} S_e(f)\,\mathrm{d}f = \frac{\Delta^2}{12}. \tag{7}$$

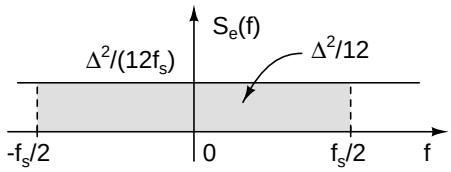

Fig 9. Power spectral density of the quantization error.

The property "white", according to the Wiener-Khintchine theorem [20], implies that the autocorrelation function of the quantization error signal is a Dirac delta. The CMQ effectively assumes that the quantization error signal is aperiodic. For a periodic signal with period p, the autocorrelation function would be also periodic, i.e. it would be a train of Dirac pulses repeating every p samples. The spectrum for such signals would be discrete. The periodicity of the quantization error is discussed further in Ch. 3 in the context of digital DMSs.

The CMQ has been widely adopted due to its simplicity and has become a basic tool for explaining the fundamental principle of the DSM operation, i.e. the noise-shaping [6, 16, 17]. Although use of the model is not justified in general, it is often justified when the input signal varies rapidly and is large compared with the quantization step Δ. The CMQ validity conditions are reviewed in greater detail in the following sections.

2.1.2 CMQ validity: formal conditions

This section will review the conditions under which the CMQ can be applied. These conditions indicate what the properties of the quantizer input signal u should be in relation to the quantization step Δ so that the CMQ will be valid. The CMQ validity conditions were first formulated by Bennett in [74], and have often been quoted in the literature since then [16, 75, 76, 78]:

1. The input signal u is in the non-overload region.
2. The input signal u has a smooth probability density function.
3. The number of quantization steps is asymptotically large.
4. The quantization step Δ is asymptotically small.

Condition 1 is fundamentally important and can easily be verified. When it is violated, the CMQ assumption immediately breaks down, as the quantization error e exceeds the permitted range $[-\Delta/2, \Delta/2]$. The remaining three conditions can serve only as a general guideline, as they are not very precise or specific from a practical point of view.

Necessary and sufficient conditions for the applicability of the CMQ to continuous-time systems can also be expressed in the form of one compact formula, known as quantization theory II (QTII) [77, 79]. This forms part of an interesting statistical theory of quantization which explores the similarity between quantization and sampling[2]. QTII highlights the fact that the conditions for the validity of the CMQ assumption can be expressed as a relation between the probability density function (PDF) of the input signal u and the quantization step Δ. It has been shown that when QTII is satisfied the input signal u and the quantization error e become uncorrelated and all the other properties of the CMQ described in Sec. 2.1.1 are valid. This theory also shows that if the input signal alone cannot satisfy the QTII, it is possible to add an independent dither signal.

Despite everything that was said above, a quantizer remains a simple, deterministic system. The correlation between the quantizer input u and the quantization error e gradually increases as the quantizer is constructed with fewer and larger steps relative to the input signal. The "no correlation" property of the CMQ occurs only under the limit conditions, and can be proven analytically using Bennet's original conditions

[2]The formula behind the QTII is intentionally not included in the thesis. While the formula itself is fairly short, its interpretation requires familiarity with a considerable amount of background information provided in [77, 79]

or the QTII. The vast theoretical literature devoted to quantization is reviewed in the monumental work of Gray and Neuhoff entitled “Quantization” [75].

2.1.3 CMQ validity: practical, approximate conditions in discrete time systems

An interesting attempt to assign concrete numerical values to CMQ validity has been made in [78]. This is based on an interesting and slightly unusual application of sampling theory. Under normal conditions, sampling is performed at a frequency that is higher than twice the signal bandwidth, in order to avoid overlapping of the spectral images. In the case of quantization we seek a contrary effect. The images of the quantization error spectrum should be allowed to overlap significantly in order to form an approximately flat spectrum. This reasoning is supported by the fact that sampling and quantization are commutative, which means that it does not matter in which order they are applied to the signal.

The condition proposed in [77, 78] imposes an upper bound on the sampling frequency f_s as a function of the quantization step Δ and the average value of the first derivative of the continuous time input signal $E\{|\dot{u}(t)|\}$ is:

$$f_s < K\frac{E\{|\dot{u}(t)|\}}{\Delta}. \tag{8}$$

The selection of the coefficient K depends on the type of input signal u. For noise-like input signals the condition (8) can be relaxed, while for sinusoidal and other deterministic signals it must be stricter. Sample values of K proposed in [78] are as follows: Gaussian $K = 3.2$, narrowband white $K = 1.2$ and sinusoidal $K = 0.4$. This result allows an appropriate f_s to be selected for a given signal but does not tell us how to select f_s for a class of signals. It is an approximate condition which relates the sampling frequency, the average growth rate of the input signal u and the quantization step Δ.

The condition given above is similar to a heuristic rule often quoted in the context of discrete time signals, stating that: *the CMQ can be applied when the quantizer input signal u traverses several quantization levels between two successive samples* [20, p. 755], [17, p. 24]. This way of stating the validity conditions appears to be useful in the field of oversampling converters, where there are usually only a few quantization steps. This approximation indicates that the CMQ is not valid for slowly varying or DC signals.

2.2 Tonal behaviour and limit cycles in delta-sigma modulation

The CMQ described in the previous section allows the quantizer to be replaced with an additive source of white noise. This enables linear analysis of the DSM and allows the signal and noise transfer functions to be calculated. Linear analysis, which is extensively covered by the fundamental DSM literature [16, 17], neglects the nonlinear element at the centre of the DSM and therefore describes the DSM noise behaviour only under ideal conditions. Only a nonlinear analysis will reveal the true complexity of DSM behaviour [80, 81]. There are many possible patterns of behaviour depending on the modulator input, initial conditions and topology. What is important from a practical point of view is that DSMs have a tendency to generate spurious tones which cannot be predicted by linear system analysis [16, 17].

The spurious tones often appear as individual tones superimposed on a noise-like spectrum and correspond to certain patterns in the DSM output signal which are likely to occur with slowly varying or DC input signals. These spurious tones are sometimes referred to as *idle channel tones* or *pattern noise* [16, 17, 82, 83]. The literature on the subject distinguishes a special case in which the DSM works with a DC input and the patterns in the output signal are strictly periodic, whereupon the DSM output signal and the quantization error signal are periodically correlated. A DSM output sequence repeating itself indefinitely is known as a **limit cycle** (LC) [16, 83][3]. The spectrum of the time waveform corresponding to a cycle consists of discrete tones, where the number of tones is equal to the length of the LC. Consequently, the entire quantization noise power will be concentrated into these few tones, and the SFR will be deteriorated [16, 82, 83].

It is very difficult to calculate analytically the exact spectra of the DSM's output signal, in view of the large variety of modulator topologies, involving one or more quantizers with one or more quantization steps. The available analytical solutions are applicable only to first-order DSMs and selected second-order topologies [10, 16, 84–90]. Even when it is possible to calculate the spectrum exactly (under certain simplified

[3] In the field of nonlinear dynamics a *limit cycle* is defined as an isolated periodic orbit [81]. We will nevertheless use the customary definition given above, as it has been adopted among engineers in the field of DSMs design [83, 84]. The term *limit cycle* has been used in a similar way to describe oscillations caused by nonlinearities due to finite precision arithmetics in discrete time recursive systems with DC inputs [20, p. 583], [84].

conditions), the formulae are rather difficult to apply in practice. It should be noted that many works published on the subject, including very recent ones, begin by emphasizing our very limited theoretical understanding of DSMs [82, 83, 91, 92]. Simulation is the basic tool for engineers working with delta-sigma modulation [16, 17, 93], and this has been used to determine the stability and noise performance of a DSM and to gain insight into LC behaviour.

2.2.1 Limit cycles in analogue DSMs with DC inputs

DC inputs constitute an important special case in DSM analysis. These can take the form of silence in audio applications, and long segments of DC levels can also appear in the case of a high oversampling ratio and a slowly-varying input. DSM behaviour under DC excitation can be conveniently described by state-space trajectories [80–82, 90, 94]. The DSM state vector, defined by outputs of the integrators, moves between two discrete points on the trajectory in each clock cycle[4]. If a trajectory alternates between a finite number of points it is called a periodic orbit [81], and under such conditions the DSM will generate a periodic pattern at its output [82–84]. A short periodic state-space orbit is shown in see Fig. 10.

It is not easy to locate LCs in continuous amplitude DSMs, as the state space comprises an infinite number of points. In recent works all output sequences up to a certain length have been examined to verify whether or not they could represent valid LCs [82, 83]. A sample result (see Fig. 11) shows how many different LCs of a particular length can occur within a given DSM topology. Locating particular limit cycles allows one to determine their impact on the SFR [82], the minimum amount of dither required to break them up and the optimum places for injecting the dither within the architecture [83]. The DC input case and related limit cycles are also important from the point of view of stability analysis, as stability under DC inputs is a necessary condition for stability under more general inputs [90]. Theoretical and simulation-based methods have been presented to determine stability bounds for state variables [90, 92, 94].

[4]This description is relevant for discrete time DSMs. Continuous time DSMs are left beyond the scope of this thesis. A short introduction to the field of asynchronous DSM and limit cycles is presented in [95].

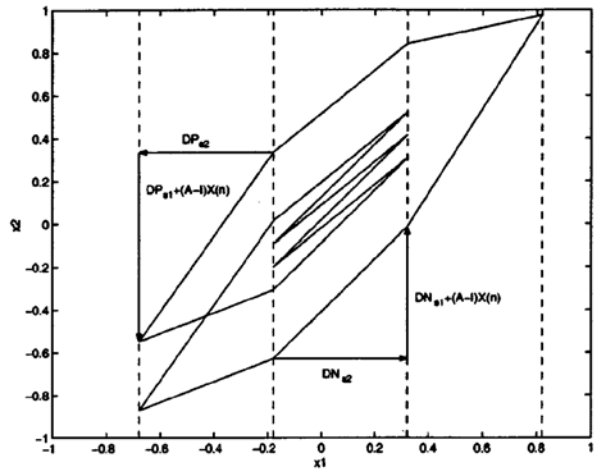

Fig 10. A short, periodic state-space orbit for a second-order system. The X and Y-axes represent the respective state variables. Figure reprinted from [82]. © [2002] IEEE.

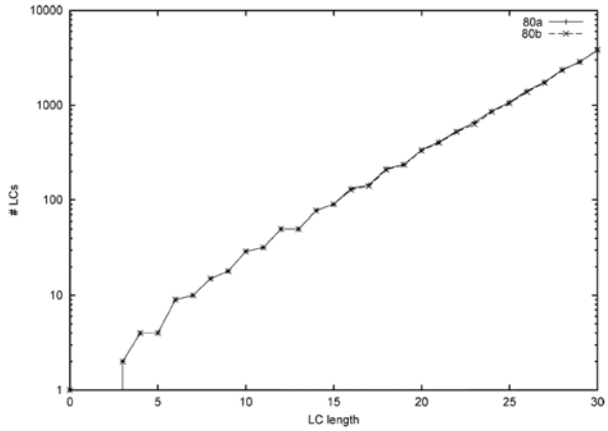

Fig 11. Occurrence of limit cycles (LCs) of particular length. Y-axis: number of LCs, X-axis: length of LCs expressed in a number of clock cycles. Figure reprinted from [83]. © [2005] IEEE.

2.2.2 Inherent periodicity of digital DSMs with DC inputs

What is characteristic of digital DSMs is their discrete amplitude. This can be considered a special case of more general continuous amplitude analogue DSMs. When digital delta-sigma modulators are implemented in fixed point digital systems all the signals are represented by rational numbers [16, Ch. 2]. It has been shown for a first-order DSM [10] and for a second-order DSM [84] that when the input $\frac{c}{d}\Delta$ is an irreducible rational number normalized with respect to the quantization step Δ the output is periodic, with a period that is a multiple of the denominator d. This result highlights the fact that digital and analogue modulators are very different. As seen in Sec. 2.2.1, limit cycles are difficult to locate in analogue DSMs and yet they occur spontaneously and worsen the SFR. On the other hand because of the finite discrete state space and deterministic dynamics, the limit cycles are inherently present in digital DSMs. When a digital DSM operates with a DC input it will always generate a limit cycle, the length of which may depend on the input and the initial conditions. The ways of handling limit cycles can therefore differ between digital and analogue DSMs.

2.2.3 Randomizing DSMs with arbitrary inputs

The goal of any randomization method is to guarantee that the quantization noise can be modelled as additive white noise. Probably the best-known and most widely used method is that called dithering (described in greater detail below), which addresses the CMQ applicability criteria by adding a source of pseudo-random noise [16, 76, 96, 97]. Other well-known methods include the use of chaos [81, 89, 91, 98–100]. Further older methods and those used primarily in the analogue domain are summarized in [16, Ch. 3]. Other methods that are specific to digital DSMs are described in greater detail in the following sections.

Dithering. Probably the most general and best-known solution to the problem of limit cycles and tonal behaviour that is applicable to all types of input (DC or varying) and all types of modulator is that called **dithering** [16, 96]. Dithering aims at improving the statistical properties of the quantization error by adding pseudo-random noise to a signal prior to quantization. This effectively "randomizes" the quantizer input and ensures that the CMQ validity conditions described in Sec. 2.1.2 and Sec. 2.1.3 are satisfied. A comprehensive overview of dithering is presented in [16, Ch. 3]. The

effects of dithering have been studied theoretically by a number of authors [76, 96, 97]. Adding an external noise to the signal inherently degrades the achievable dynamic range. This fundamentally negative side-effect can be addressed in a number of ways. One possibility is simply to subtract the dither signal from the quantizer output. This approach, called subtractive dithering [16, p.123], is not very commonly used in DSMs as it requires additional precise circuitry to remove the dither signal. A non-subtractive approach appears to be more popular in the DSM field, but it requires other means of achieving an adequate dynamic range.

Noise-shaped dithering. This is an approach in which the dither signal is filtered away from the neighbourhood of the input signal in the same way as the quantization error. The simplest way of achieving this is to add the dither immediately before the quantizer [16]. Optimum performance can be obtained in this configuration when the dither is spectrally white with a triangular PDF and a peak-to-peak amplitude of two least significant bits (LSB) of the quantizer [76][5]. A dither signal injected directly before the quantizer must have a significant amplitude and in certain cases may lead to instability in the modulator. Recent studies suggest that this is not a very effective way of removing strong limit cycles [83], [V]. A much smaller disturbance introduced into the DSM state variables (registers) or at the input may yield better results [83, 101, 102].

LSB dithering. It has been known for some time that adding a small disturbance to the input can help to break up the limit cycles [102, 103]. Low-level, or LSB dithering has been recommended recently in [97, 101]. In this case, LSB refers to the LSB of the input signal. For a b-bit input signal the LSB dither could be added to the $(b+1)$st LSB and remain below the noise floor of the input signal. If the extra hardware needed to extend the DSM bus by one bit is a concern, the dither could be added to the LSB of the input signal and be noise-shaped. For the dither to be effective, however, the order of the dither filtering should be matched to the transfer function of the DSM. This relationship has been studied for a MASH architecture, and the corresponding DSM and dither filter topologies are determined in [101].

[5]The term LSB is used in two ways in the context of dithering. Sometimes it refers to the LSB of the quantizer [76], but in digital DSMs it can also refer to the LSB of the input signal [101]. This distinction is important, because these LSBs differ dramatically in amplitude!

2.2.4 Randomizing digital DSMs with DC inputs

Once it has been recognized that limit cycles are a fundamental feature of digital DSMs, the problem of tonal behaviour can be reduced to the following two problems: the length of the LC, which determines the number of discrete tones among which the quantization error power is distributed, and how evenly the power is distributed between the tones. There are many methods available for randomizing digital DSMs working with DC inputs.

Precalculated initial conditions have been used for long time as a means of ensuring a good quality DSM spectrum [41, 49, 66–69]. The initial conditions can be loaded using a simple reset mechanism and therefore do not require any additional hardware, but the method does require resetting the DSM each time the DC input is changed. It has been shown in [67–69] that this approach results in a smooth spectrum.

The problem of short limit cycles has been addressed in a number of works [19, 61, 65, 70]. Initial conditions have been used explicitly as a way of avoiding short limit cycles [65], [19, p. 349]. It has also been noted that setting the input LSB to '1' increases the length of the limit cycles [61, 70]. At the same time, the quantization noise power is smoothly distributed over a large number of tones. This approach entails a small hardware penalty as it requires extending the DSM bus by one bit, but it does not require resetting the DSM and could possibly be used with arbitrary inputs.

Full attention is given to the limit cycle length in [II], [IV], [V], [VI]. Probably all of the preceding published works had aimed either at "removing" the limit cycles or lengthening them beyond the detection levels. It was shown empirically in [VI] and has been proved analytically in [104, 105] that the length of the limit cycles can be determined. As a result it is possible to establish a relationship between the amount of DSM hardware and the quality of the output spectrum. Thus it is shown in [VI] for classical MASH and error-feedback modulator (EFM) topologies that the sequence length, and consequently the number of tones, can be controlled in practice regardless of the DC input. The length of a limit cycle can be set by applying certain initial conditions [VI], or by setting the input LSB [IV] and selecting the appropriate bus width. A digital DSM working in such a regime can be regarded as a delta-sigma encoder. This approach is one of the main contributions of this thesis and is discussed further in Ch. 3.

The classical MASH and EFM topologies considered in [VI] do not use the hardware in a very efficient manner, as the longest orbits utilize only a fraction of the available state space points. This observation allowed the amount of hardware in MASH to be

reduced without any loss of sequence length [106]. In other works certain topology modifications have been found to allow better utilization of the remaining state-space points and to enlarge the sequence length dramatically regardless of the initial conditions and the DC input [72]. It has also been shown that using a prime modulus quantizer can lead to very long sequences and a smooth spectrum regardless of the DC input and initial conditions [71]. A general class of digital delta-sigma modulators based on a variable modulus quantizer is studied in Ch. 4 and constitutes the second main contribution of this thesis.

2.3 Summary

This chapter has briefly reviewed the classical model of quantization (CMQ) and its applicability conditions. When the model applies, the quantization error can be modelled as an input-independent, additive source of white noise. We have reviewed both the formal and the approximate conditions determining the applicability of the model. When CMQ applies, it allows a simple linear DSM analysis, and when it fails, the DSM behaves unpredictably and is likely to generate spurious tones and limit cycles.

It is very difficult to locate and study particular limit cycles in the context of continuous amplitude analogue DSMs, as stable analogue DSMs can be characterized by a bounded but infinite state space, a characteristic which distinguishes them from discrete amplitude, digital DSMs. Digital DSMs working with DC inputs constitute a particularly simple class from the point of view of limit cycle analysis, as such modulators always generate limit cycles, the only unknown that remains being their length.

For the above reasons the problem of randomization can be approached differently for digital and analogue DSMs . While dithering applies to both modulator types, there are several randomization techniques which apply specifically to digital DSMs. These include setting specific initial conditions, setting the input LSB to one, or using a prime modulus. The following chapter will explore the possibility of controlling the length of the limit cycles precisely for all DC inputs as a means of assuring good quality output spectra.

3 Delta-sigma encoders

Digital DSMs working with DC inputs and implemented in a fixed point arithmetic system constitute a special case among analogue and digital modulators, as under such conditions they inherently generate periodic sequences. Short sequences are undesirable, because they concentrate the DSM quantization noise power into a few dominant tones and in consequence worsen the SFR. This problem is usually addressed by lengthening the limit cycles in an uncontrolled way beyond any length that matters. This can be achieved by adding pseudo-random noise, called dither, or by several other techniques described in Sec. 2.2.

This chapter will consider a different approach to the problem of short limit cycles, that proposed in [IV], [V], [VI]. We will explore a method of controlling the lengths of the limit cycles precisely and will study the consequences of this approach. From such a perspective, the DSM can be seen as an encoder which transforms an input digital word into an output periodic sequence of prescribed length. In this thesis we will refer to the DSMs working in this way as delta-sigma encoders (DSEs), in order to distinguish this mode of operation from more general delta-sigma modulation.

This chapter begins with an analysis of an ideal DSE in Sec. 3.1. The classical model of quantization will be adapted to reflect the periodicity of the DSE's quantization error sequences in Sec. 3.1.1, which will allow us to observe the relationship between sequence length and the quantization noise spectrum in Secs. 3.1.2 and 3.1.3. Practical DSEs are considered in Sec. 3.2. After presenting a practical means for controlling the sequence length in Sec. 3.2.1, we will compare measured and simulated DSE spectra with the performance of the linear model in Sec. 3.2.3. The chapter will conclude with a comparison of dithered and non-dithered DSEs in Sec. 3.2.5.

3.1 An ideal delta-sigma encoder

A delta-sigma encoder is a specific type of delta-sigma modulator. A traditional DSM is a continuous amplitude feedback system in which the input and the internal state are both continuous amplitude signals, as shown in Fig. 5. Even the output digital signal, after conversion to the analogue domain, is affected by noise and can never be reproduced ideally. On the other hand, the input and the internal state of a digital DSM

are purely discrete in amplitude. All arithmetic operations are performed with perfect accuracy and all the signals inside and at the input and output interfaces are in the fixed point digital domain (see Fig. 12) and can be represented by integers or rational numbers. The difference between digital and analogue DSMs becomes particularly apparent when a digital DSM is implemented with a narrow bus width. Due to the limited internal state space, such a modulator is particularly prone to producing repeating patterns, which lead to unwanted tonal behaviour in the output spectrum.

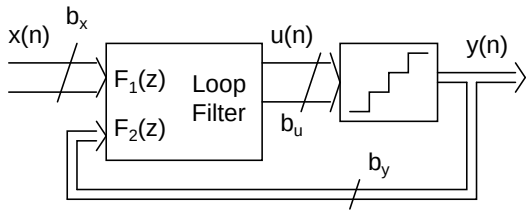

Fig 12. A digital delta-sigma modulator.

A digital DSM can be regarded as a deterministic finite state machine (FSM). It is a type of FSM in which there is one and only one transition to the next state for each pair of state and input. When a digital DSM works with a DC input and the trajectory returns to a previously visited state, it must follow the same state-space orbit again. Therefore such a system must always produce periodic sequences. A DSM producing short periodic sequences clearly cannot be modelled by a linear model based on the CMQ presented in Ch. 2, as the CMQ assumes an aperiodic quantization error signal. A DSE with a periodically correlated output and a quantization error signal requires certain modifications to be made to the conventional quantization noise model.

3.1.1 Quantization noise model for an ideal DSE

The classical model of quantization discussed in Ch. 2 assumes that the quantization error signal is statistically independent of the input signal, is uniformly distributed in the range $\pm\Delta/2$, is aperiodic and has a flat power spectral density. The DSE differs fundamentally from the traditional DSM in that the quantization error signal is inherently periodic. It is a periodic sequence with a period L_s:

$$e(n) = e(n + k \cdot L_s), \text{ where } k \in \mathbb{Z}. \tag{9}$$

An important consequence of the quantization error signal being periodic is that the Fourier spectrum is discrete. For a periodic sequence $e(n)$ of length L_s, the discrete Fourier transform (DFT) is another periodic sequence with L_s tones [20, 107]. The quantization noise power is distributed among exactly L_s tones between zero and the sampling frequency f_s and there is no power between the tones. Fig. 13 shows an idealized, two-sided discrete power spectrum $P_e(f_k, L_s)$ for the quantization error $e(n)$, where:

$$f_k = f_s \frac{k}{L_s}. \tag{10}$$

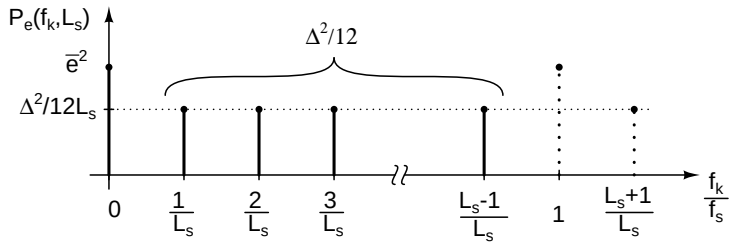

Fig 13. A discrete power spectrum of the quantization error.

The classical model of quantization assumes that the quantization error is uniformly distributed in the range $\pm\Delta/2$, which implies that the quantization noise variance is equal to $\Delta^2/12$, see Ch. 2. In the digital domain all signals are initially quantized and a coarse quantizer only re-quantizes its input signal. For this reason the distribution of the quantization error over the interval Δ is discrete. It has been shown in App. 2 that if the quantization error is uniformly distributed over a large number of discrete points the quantization noise variance can be approximated with a commonly assumed value:

$$\sigma_e^2 \approx \frac{\Delta^2}{12}. \tag{11}$$

It can be shown for a real-valued quantization error signal that the average power P_e is a sum of the variance σ_e^2 and the squared mean value $\bar{e}^2$ (squared DC value):

$$P_e = \sigma_e^2 + \bar{e}^2. \tag{12}$$

Parseval's relation for a periodic quantization error $e(n)$ with the fundamental period L_s

can be written as:

$$P_e = \frac{1}{L_s}\sum_{n=0}^{L_s-1} e^2(n) = \sum_{k=0}^{L_s-1} P_e(f_k, L_s) = \bar{e}^2 + \sum_{k=1}^{L_s-1} P_e(f_k, L_s). \tag{13}$$

The variance is a power of the active components and is therefore distributed over $L_s - 1$ discrete tones, excluding the DC component. Taking into consideration (13) and (12) we can write:

$$\sigma_e^2 = \sum_{k=1}^{L_s-1} P_e(f_k, L_s). \tag{14}$$

Assuming a flat spectrum $P_e(f_k, L_s)$ the power per tine is equal to:

$$P_e(f_k, L_s) = \begin{cases} \bar{e}^2 & k = 0 \\ \frac{\Delta^2}{12(L_s-1)} \approx \frac{\Delta^2}{12L_s} & k = 1, 2, \ldots, L_s - 1. \end{cases} \tag{15}$$

This result characterizes the quantization error spectrum for an ideal DSE, shown in Fig. 13. The model assumes a periodic quantization error $e(n)$, with a period L_s, a constant quantization error variance $\sigma_e^2 = \Delta^2/12$ and a flat quantization error spectrum. As a result the number of discrete tones in the spectrum and their level depends on the period length L_s. The relation between the average power, the mean value and the variance (12) and also Parseval's relation (13), are proved in App. 1.

3.1.2 Dependence of SFR on the sequence length

The distinguishing characteristic of a DSE is the ability to control the number of tones and their levels. This is a key difference between a delta-sigma modulator and encoder. As has been shown in Sec. 3.1.1 a DSE generates periodic sequences of a constant average power but varying length. Consequently, not only will the sequence length change, but the number of tones and their average levels will change as well. This principle is illustrated in Fig. 14, where $P_e(f_k, L_s)$ is a discrete power spectrum of the quantization error sequence. We can see in Fig. 14 that as the sequence length doubles, the amount of power per tone is reduced by a factor of two. The level $P_e(f_k, 16)$ is nearly half of $P_e(f_k, 8)$. In the first case we distribute the variance $\Delta^2/12$ over 15 tones, in the other over 7 tones, according to (15). As the sequence length increases, the ratio comes closer to $1/2$. This corresponds to a 3 dB difference in the level of tones each time the sequence length is doubled.

$$\lim_{L_s\to\infty} \frac{P_e(L_s)}{P_e(2L_s)} = \lim_{L_s\to\infty} \frac{L_s - 1}{2L_s - 1} = \frac{1}{2} \triangleq -3\,\text{dB} \tag{16}$$

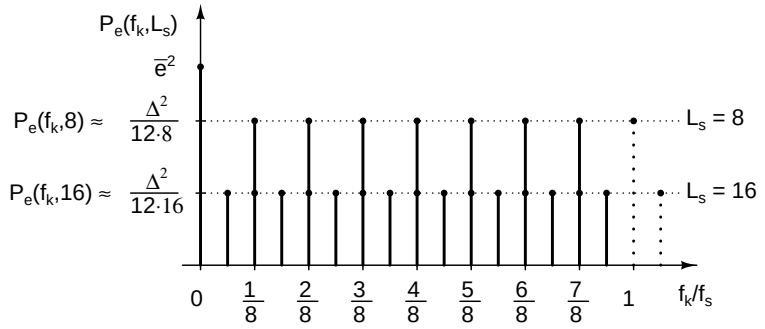

Fig 14. The level of tones as a function of DSE sequence length.

This effect can be used in the design of DSEs. We will show in Sec. 3.2 what conditions must be satisfied for a practical DSE to produce a flat spectrum and to behave like the ideal model described in this section.

3.1.3 Frequency domain model of a delta-sigma encoder

So far we have adapted the classical model of quantization to reflect the periodic nature of the DSE. We will now use these results to calculate an ideal spectrum for the DSE output signal y. Let us consider a simple third-order modulator with a noise transfer function defined by:

$$\mathrm{NTF}(z) = \left(1 - z^{-1}\right)^3. \tag{17}$$

From this point forward we will use a one-sided spectrum representation and assume that all the power is concentrated in the range between zero and half of the sampling frequency f_s. This convention is used in the classical book on delta-sigma modulation [16, p. 5] and is recommended in the IEEE standard on definitions of physical quantities for frequency and time metrology [108]. In order to obtain a one-sided spectrum we will multiply the discrete power spectrum $P_e(f_k, L_s)$ by a factor of two, neglecting the existence of the DC component in (15) for the sake of simplicity, as the NTFs of the architectures used in this work have zeros at DC. We can express the one-sided output spectrum of an ideal DSE $P_y^{\mathrm{I}}(f_k, L_s)$ by shaping a flat spectrum of the quantization error,

thus:

$$P_y^{\mathrm{l}}(f_k, L_s) = 2\frac{\Delta^2}{12L_s}\left|\mathrm{NTF}(z)\right|^2_{z=e^{j2\pi f_k/f_s}} = \frac{\Delta^2}{6L_s}\left|\left(1-e^{-j2\pi f_k/f_s}\right)^3\right|^2. \quad (18)$$

DSE output spectra calculated for two sequence lengths are shown in Fig. 15. The ratio of the sequence lengths is $2^9/2^5 = 2^4$, which corresponds to a $4 \cdot 3\,\mathrm{dB} = 12\,\mathrm{dB}$ difference on a logarithmic scale.

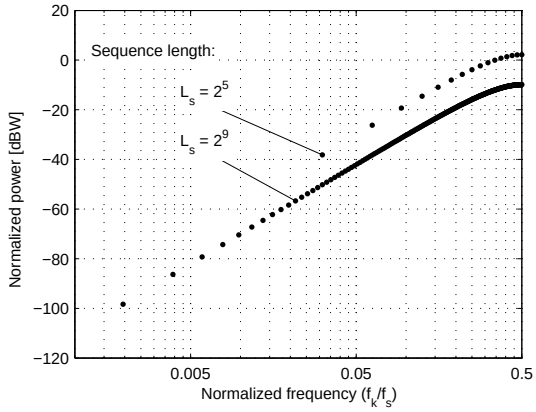

Fig 15. Noise-shaped spectrum of an ideal DSE $10\log(P_y^{\mathrm{l}}(f_k, L_s))$. The level of the tones depends on the sequence length L_s.

3.2 A practical delta-sigma encoder

In this section we will look at two classical modulator topologies, EFM and MASH, as shown in Fig. 16 and Fig. 17, and use them to encode a DC input value as a periodic sequence of prescribed length. Fig. 18 shows the content of the EFM1 block in MASH. The section will begin by reviewing practical methods for controlling the sequence length, after which we will compare the spectra of an ideal encoder with the simulated and measured spectra of the practical encoder. Finally we will compare the DSE performance with and without dithering.

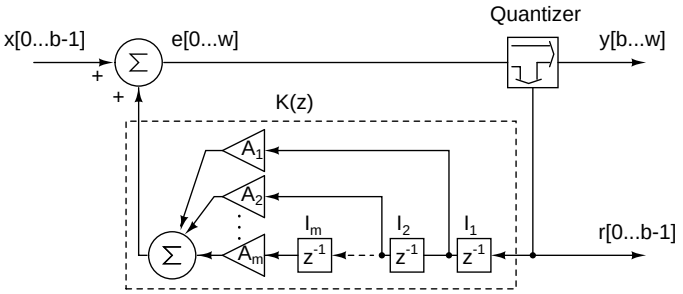

Fig 16. EFM architecture. Revised from [VI]. © [2005] IEEE.

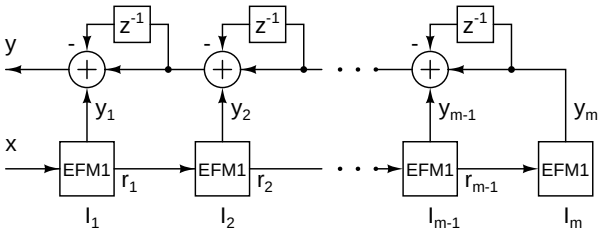

Fig 17. MASH architecture. Revised from [VI]. © [2005] IEEE.

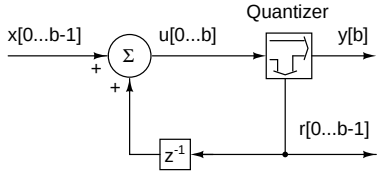

Fig 18. EFM1 - first-order error-feedback architecture. Revised from [VI]. © [2005] IEEE.

3.2.1 Sequence length control

A number of methods have been published to date for the precise control of DSM sequence length. The first study on this topic showed that the sequence length can be controlled for classical MASH and EFM topologies of order two to five by applying

predefined initial conditions each time the DC input changes [VI]. The second method required expansion of the input bus by one bit, but the initial conditions applied only at "power-up" [IV]. More recent methods based on a prime modulus quantizer [71] and prime feedback [72] require certain structural modifications, but allow sequence length control regardless of the initial conditions.

Reloading initial conditions. The error feedback architecture shown in Fig. 16 is composed of gain elements A_i, adders, a quantizer and a number of registers represented by unity delays z^{-1}. The method described here requires registers equipped with a standard reset or load mechanism. Each time the DC input changes all the registers will be loaded with the corresponding initial conditions I_i, which are stored in the hardware.

It has been shown in [IV], [V], [VI] that it is possible to control the sequence length of the DSE by applying predefined initial conditions. The results are collected together in two tables, of which Table 1 defines in a general way groups of favourable initial conditions applicable to the register in the EFM and MASH structures shown in Fig. 16 and Fig. 17, respectively. These conditions are favourable because when they are applied the DSMs produce very long sequences of known length.

Consider a third-order EFM as an example. As seen in Table 1, when the sum of the initial conditions I_1 and I_3 is odd, the sequence length can be calculated according to Table 2 as

$$L_s = 2^{b+1}, \tag{19}$$

where b is the DSM bus width. An example of an initial condition which can be stored in the hardware would be $I_1 = 1$, $I_2 = 0$, $I_3 = 0$. If this initial condition is applied when the DC input x is set, the sequence length will be known exactly and will depend only on the DSM bus width b and not on the input x. At the same time it is the longest sequence that this particular structure can ever generate.

Consider a second-order MASH as another example. Again as seen in Table 1, when the initial condition I_1 is odd, the sequence length will vary between the minimum guaranteed 2^{b-1} and the maximum available 2^{b+1}, depending on the input DC value x. Nevertheless, it is known that the sequence length will never be shorter than the minimum guaranteed value, and therefore we can base our SFR predictions on the lower of the two values and be sure that the DSM will not perform worse than that. The general rules formulated in Table 1 and Table 2 were derived as a result of extensive simulation [VI]. These results were later confirmed analytically for the MASH architecture [104].

Table 1. Initial conditions. Revised from [VI]. © [2005] IEEE.

Modulator order m	MASH	EFM $\sum I_i = \text{odd}$
2	I_1 odd	$I_1 + I_2$
3	I_1 odd	$I_1 + I_3$
4	I_1 odd	$I_1 + I_2 + I_3 + I_4$
5	I_1 odd	$I_1 + I_5$

Table 2. Sequence length. Revised from [VI]. © [2005] IEEE.

Modulator order m	Guaranteed sequence length L_s	Maximum sequence length
2	2^{b-1}	2^{b+1}
3	2^{b+1}	2^{b+1}
4	2^{b+1}	2^{b+2}
5	2^{b+2}	2^{b+2}

Applying the initial conditions only at power-up. The method described above can be extended so that the conditions are applied only at the moment when the DSM is powered up [IV]. After that there is no need to reload the initial conditions, the input can change in any way, but each time it remains at a given DC value the DSM will generate sequence lengths as described in Table 2.

In practice, this method requires extending the DSM bus width by one bit and connecting the LSB to zero. As a result, the input x can have only even values. It has been shown analytically in [IV] that once the DSM has started from the initial conditions defined in Table 1 and the input remains even, the contents of the registers always satisfy the conditions of Table 1 [IV]. In other words, when the DSM input is even, Table 1 defines a set of initial conditions which has the property that all subsequent states lie in the original set. Finding sets with this property of *positive invariance* has been used to guarantee DSM stability [94]. In this case it is used to guarantee that the lengths of the state space orbits can be controlled in accordance with Table 2.

Prime modulus and prime feedback. Recent research has shown that the use of a prime modulus quantizer allows the sequence length to be controlled irrespective of the DSM input and initial conditions [71]. For a MASH composed of cascaded first-order sections, the sequence length is always equal to the prime modulus Q of the first section. The role of the modulus in the context of digital delta-sigma modulation will be discussed in Ch. 4. Another way of controlling the sequence length independently of the DSM input and initial conditions is based on a prime feedback factor [72]. When applied to MASH, this technique allows the modulator to generate sequences of lengths corresponding to the coverage of almost the entire state space. A third-order 19-bit MASH, for example, produces a sequence length of $(2^{19}-1)^3$.

Both of the aforementioned methods allow sequence length control regardless of the DSM input and initial conditions. In particular, the prime feedback method utilizes hardware resources in the most efficient way and leads to sequence lengths that grow exponentially with the DSM order. It should be noted, however, that the method based on a prime modulus introduces a gain with the same prime value. The role of the gain is discussed in the context of a variable modulus quantizer in Sec. 4.2, and in the context of frequency generation in Sec. 5.1.2.

3.2.2 Dependence of SFR on the DSE bus width

It has been suggested in Sec. 3.1 that the spectrum of the DSE is inherently tonal and that the levels of the tones depend on the sequence length. This observation has been published earlier in [II], [III], [IV], [V], [VI] and has been supported by simulation results. In this section we further confirm it by means of measurement using a simple test system described in App. 4. We implemented a third-order EFM, shown in Fig. 16, with two different bus widths: 9-bits and 15-bits. This is the same architecture as was used previously in [VI]. According to (19) a 9-bit EFM produces a sequence of length $L_s = 2^{10}$. We will refer to this as EFM10. Correspondingly, we will refer to the 15-bit EFM, which generates a sequence of length $L_s = 2^{16}$, as EFM16. The EFMs operated at a clock frequency $f_s = 20\,\text{MHz}$, which resulted in a tone spacing of $f_s/2^{10} \approx 10\,\text{kHz}$ for EFM10 and $f_s/2^{16} \approx 300Hz$ for EFM16. The spectrum measurements are presented in Fig. 19, where the EFM10 spectrum is plotted in black and the EFM16 spectrum in grey.

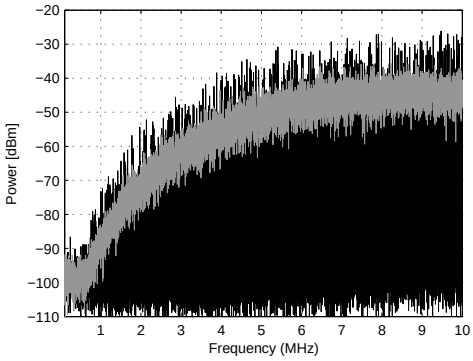

RBW	1	kHz
VBW	3	kHz
Nsamp	10001	
SWT	10	s
Att	5	dB
Ref	-20	dBm

(a) Spectrum measured from zero to half of the sampling frequency.

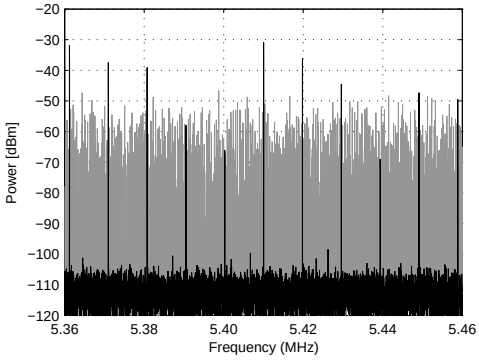

RBW	20	Hz
VBW	50	Hz
Nsamp	10001	
SWT	0.2	s
Att	20	dB
Ref	-20	dBm

(b) Spectrum magnified around 5.41 MHz.

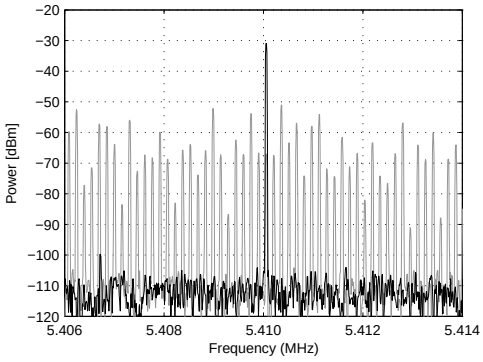

RBW	20	Hz
VBW	50	Hz
Nsamp	10001	
SWT	0.2	s
Att	20	dB
Ref	-20	dBm

(c) Spectrum further magnified around 5.410 MHz.

Fig 19. Measured spectrum of the DSE output signal. Black: third-order EFM implemented with a 10-bit bus; Grey: 16-bit bus.

The spectrum in the range from zero to half of the sampling frequency with the resolution bandwidth (RBW) set to 1 kHz is shown in Fig. 19(a). This resolution bandwidth is sufficient to discriminate between the 10 kHz-spaced tones of the DSM10, but is insufficient for the 300 Hz-spaced tones of the DSM16. In the latter case from 3 to 4 tones are averaged within the resolution bandwidth. As a result the spectrum of the DSM16 appears to be smoother in this measurement. It is important to note that discrete DSE spectra with tone spacings smaller than the RBW are practically indistinguishable in spectral measurements and appear as smooth curves.

Increasing the RBW to 20 Hz reveals the tonal nature of DSM16, as shown in Fig. 19(b) and Fig. 19(c). We see that, in accordance with our expectations, the tone spacing for DSM10 is equal to 10 kHz. Furthermore, the levels of the tones produced by DSM16 are clearly below those of DSM10. According to the estimate presented in Sec. 3.1.2, the average level of the tones should decrease by 3 dB for each doubling of the number of tones. In our case, this corresponds to a difference of $3\,\mathrm{dB} \cdot 2^{16}/2^{10} = 18\,\mathrm{dB}$. The differences in the levels of the tones in Fig. 19 appear to follow this principle.

3.2.3 Ideal, simulated and measured spectra

In this section we will compare the ideal, simulated and measured spectra more carefully in order to gain confidence in the correctness of our simulation-based analysis. Our measurement system is described further in App. 4. The exact spectrum computation using the DFT will be presented below and this will be compared later with the measured spectra.

Since the DSE produces repeating periodic sequences, it constitutes a special case for spectral analysis. In general, when studying aperiodic or very long sequences, the spectrum is only **estimated** from a fragment of the sequence, but in the DSE case the sequence length is known and therefore the spectrum can be computed **exactly** using DFT [20, 107]. The DFT[6] calculated from one period of a periodic output signal y can be expressed as:

$$Y(k) = \sum_{n=0}^{L_s-1} y(n) e^{-j2\pi kn/L_s} \qquad k = 0, 1, 2, \ldots, L_s - 1. \tag{20}$$

The power spectrum $P_y(f_k, L_s)$ of the DSM output signal y can be calculated using the

[6]The definition of the DFT used in this work has the same form as the definitions used in [20, 107] and in the Matlab® implementation.

DFT as:

$$P_y(f_k, L_s) = \frac{1}{L_s^2}|Y(k)|^2, \tag{21}$$

where:

$$f_k = f_s \frac{k}{L_s}. \tag{22}$$

$P_y(f_k, L_s)$ is a two-sided spectrum defined so that the sum of all the tones between zero and the sampling frequency f_s gives the average power P_y of the DSM output signal. Parseval's relation is proved in App. 1 but is repeated here for convenience:

$$P_y = \frac{1}{L_s}\sum_{n=0}^{L_s-1} y^2(n) = \sum_{k=0}^{L_s-1} P_y(f_k, L_s). \tag{23}$$

In order to comply with the spectrum representation standards [16, 108] and in order to compare the simulated and measured spectra, we will use a one-sided spectrum representation. In order to distinguish the one-sided power spectrum based on the DSE linear model $P_y^{\mathrm{l}}(f_k, L_s)$, introduced in Sec. 3.1.3, from a spectrum based on simulation, we will denote the later as $P_y^{\mathrm{s}}(f_k, L_s)$. The one-sided DSE output power spectrum can be calculated using the DFT, as:

$$P_y^{\mathrm{s}}(f_k, L_s) = \frac{2}{L_s^2}|Y(k)|^2. \tag{24}$$

There are a few additional steps required to convert $P_y^{\mathrm{s}}(f_k, L_s)$ into a practical measurable spectrum $M_y^{\mathrm{s}}(f_k, L_s)$. The digitally generated discrete signal y is converted into a continuous voltage waveform using a DAC, which performs a sample and hold operation which must also be taken into account. The first-order sample and hold operation is represented by a sinc function in the frequency domain:

$$\mathrm{sinc}(f) = \frac{\sin(\pi f/f_s)}{\pi f/f_s}. \tag{25}$$

The DAC also scales the DSE output signal y by a factor $\Delta = 75\,\mathrm{mV}$. As a result the computed power (24) will be multiplied by a scaling factor Δ^2. Finally we have to include the effect of the $50\,\Omega$ load, apply a logarithmic scale and add 30 dB in order to convert the units to the dBm as used by the spectrum analyzer. The final formula for the measured power includes all of the effects mentioned above, i.e.

$$\begin{aligned} M_y^{\mathrm{s}}(f_k, L_s) &= 10\log\left(\Delta^2 P_y^{\mathrm{s}}(f_k, L_s)\cdot \mathrm{sinc}^2(f_k)\right) + 30 - 10\log(50) \\ &= 10\log\left(\Delta^2 \frac{2}{L_s^2}\left|Y(k)\right|^2 \cdot \mathrm{sinc}^2(f_k)\right) + 30 - 10\log(50). \end{aligned} \tag{26}$$

Similar steps have been taken to adjust the ideal DSE spectra $P_y^{\mathrm{I}}(f_k, L_s)$ derived in Sec. 3.1.3 to the measurement results:

$$\begin{aligned} M_y^{\mathrm{I}}(f_k) &= 10\log\left(P_y^{\mathrm{I}}(f_k, L_s)\cdot \mathrm{sinc}^2(f_k)\right) + 30 - 10\log(50) \\ &= 10\log\left(\frac{\Delta^2}{6L_s}\left|\left(1-e^{-j2\pi f_k/f_s}\right)^3\right|^2 \cdot \mathrm{sinc}^2(f_k)\right) + 30 \\ &\quad -10\log(50). \end{aligned} \tag{27}$$

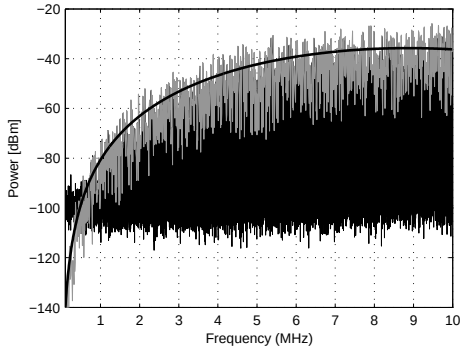

RBW	1	kHz
VBW	3	kHz
Nsamp	10001	
SWT	10	s
Att	5	dB
Ref	-20	dBm

(a) Spectrum measured from zero to half of the sampling frequency.

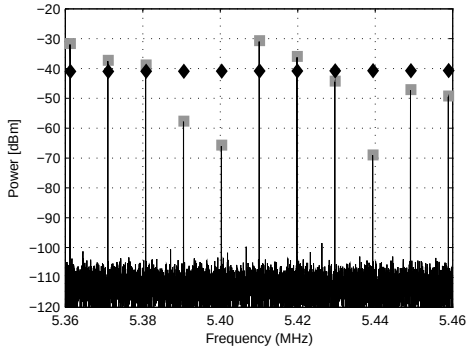

RBW	20	Hz
VBW	50	Hz
Nsamp	10001	
SWT	0.2	s
Att	20	dB
Ref	-20	dBm

(b) Spectrum magnified around 5.410 MHz.

Fig 20. Measured and calculated spectra of a DSE output signal. Third-order EFM implemented with a 10-bit bus. Black diamond: spectrum based on the linear model $M_y^{\mathrm{l}}(f_k, 2^{11})$; Grey rectangle: spectrum based on computer simulation $M_y^{\mathrm{s}}(f_k, 2^{11})$; Black line: measured spectrum.

The simulated power spectrum $M_y^{\mathrm{s}}(f_k, L_s)$, the ideal spectrum based on the linear model $M_y^{\mathrm{l}}(f_k, L_s)$ and the measurement results are shown in Fig. 20. As can be seen, the measured and simulated spectra overlap exactly. Each point of the spectrum $M_y^{\mathrm{s}}(f_k, L_s)$ based on the DFT corresponds exactly to one discrete tone and there is no power between the tones. The ideal spectrum $M_y^{\mathrm{l}}(f_k, L_s)$ approximates the varying levels of the tones. The simulated and measured DSE spectra generally follow the ideal prediction, in that all of them have the same tonal nature and equal average powers. The following sections will show that the simulated spectrum does not approach the ideal one, even when dithering is used. The experiment shows that the measured and simulated spectra can be approximated using the ideal DSE model introduced in Sec. 3.1. Moreover, the measured spectrum perfectly matches the simulated spectrum, which correctly predicts the tonal nature of the DSE.

3.2.4 Evaluating the worst-case performance

The concept of DSE simplifies the problem of evaluating the worst-case SFR in a reliable way. The fact that the DSE is a fully deterministic system producing repeating sequences of controlled length can be used to evaluate the spectrum under all possible practical conditions. The overall simulated spectrum for a third-order error-feedback DSE, as shown in Fig. 21, is relatively smooth and free from clearly dominant spurious tones for all input DC values in the available range. If we look at the three-dimensional image from the direction of the frequency axis, we will see the upper contour projected onto the power-frequency plane. This contour, known as the spectrum envelope, shows the highest level of tones that occurs for all DC inputs. The envelope quantifies the worst-case SFR which is valid for all DC inputs.

Three envelopes for different bus widths are shown in Fig. 22. The worst case SFR, represented by an envelope, scales very regularly with increasing bus width. Every additional bit doubles the sequence length, in accordance with (19), and this improves the SFR by 3 dB, according to (16). As has been shown in Sec. 3.2.3, the simulated and measured DSE spectra overlap completely. Therefore the worst case SFR for the DSE can be computed with full confidence using the DFT.

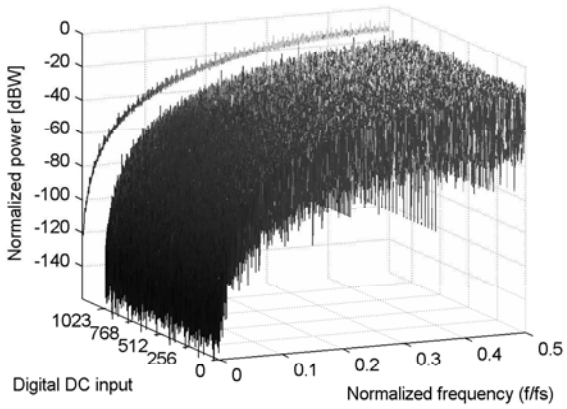

Fig 21. Discrete power spectrum $10\log(P_y^{\mathrm{s}}(f_k), 2^{11})$ with all digital DC inputs for a third-order EFM implemented with a 10-bit bus.

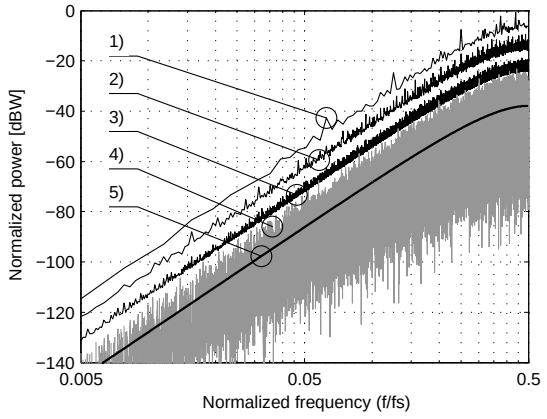

Fig 22. Spectrum envelopes and corresponding sequence lengths 1) $L_s = 2^8$, 2) $L_s = 2^{12}$, 3) $L_s = 2^{16}$. 4) Discrete power spectrum $10\log(P_y^{\mathrm{s}}(f_k, 2^{16}))$ for a single DC input, 5) Ideal model $10\log(P_y^{\mathrm{I}}(f_k, 2^{16}))$.

3.2.5 Dithering

We studied the impact of dithering on the performance of a third-order EFM in [V], where we examined the effects of dithering on the spectrum and on hardware consumption. In order to make a hardware comparison possible, the dither generators[7] were designed to produce a prescribed sequence length $L_s = 2^{16}$. The dithered DSM was then compared with a DSE designed according to [VI], producing another sequence of exactly the same length L_s. Three dither generators were used in the experiment: one based on a linear shift register (LSR) and two based on congruent random number generators producing rectangular probability density function (rPDF) dither and triangular (tPDF) dither. A comparison of the spectra is presented in Fig. 23, which shows that those with rPDF and tPDF dither and that of the DSE are very close. The spectra are presented as envelopes and therefore show the worst SFR for all DC inputs. The DSM based on the LSR dither performed significantly worse than the other three approaches. A classical noise-shaped dithering scheme was used in which the dither signal was added immediately prior to quantization, as shown in Fig.1 of [V]. It was shown that the DSE performs as well as the dithered DSM, while using less hardware resources.

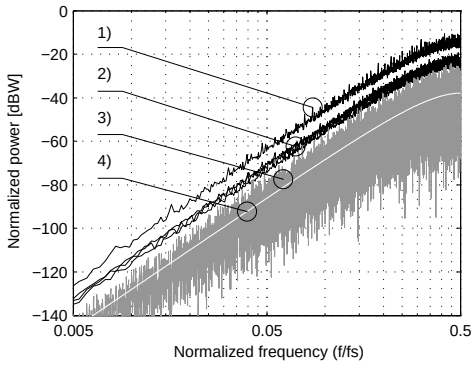

Fig 23. 1) Spectrum envelope for a DSM with shift-register dither. 2) Overlapping envelopes for DSMs with rPDF dither and tPDF dither. 3) Discrete power spectrum $10\log(P_y^{\mathrm{s}}(f_k, 2^{16}))$ for a single DC input without dither 4) Ideal DSE model $10\log(P_y^{\mathrm{I}}(f_k, 2^{16}))$. Revised from [V]. © [2006] IEEE.

[7] It is interesting to note that digital dither generators are also finite state machines. Consequently they, too, must produce finite length sequences. The lengths of the sequences depend on the amount of hardware used to build such a generator.

3.3 Summary

This section has examined the fundamental properties of ideal and practical DSEs, an interesting sub-class of DSMs. These encode a digital input number in repeating sequences of arbitrary length. The output sequence is composed of the DC component corresponding to the input DC value and filtered quantization noise. What is characteristic of delta-sigma encoders, and what distinguishes them from delta-sigma modulators, is the inherently tonal nature of the quantization noise spectrum. Since the density of the tones depends on the amount of hardware used to construct the DSE, it is possible to optimize DSEs for a given application and create a sufficiently dense spectrum that is practically indistinguishable from a continuous one. It has also been shown that setting the initial conditions in order to maximize the sequence length results in a flat spectrum that is free from dominant tones. Even a strong, triangular PDF dither cannot improve the spectrum any further, provided that the sequence length is guaranteed by the design.

4 Variable modulus DSM

This section will study a variable modulus DSM (VMDSM), as introduced in [VII]. A fractional-N frequency synthesizer used with the VMDSM can multiply a stable oscillator frequency by any rational number. Since a traditional DSM with a truncation quantizer can only have a modulus Q that is a power of two, i.e. $Q = 2^b$, the fractional ratios generated suffer from a limited choice for the denominator 2^b (see (3) in Sec. 1.3.2). Adding a certain amount of asynchronous logic will allow variable modulus quantizers to be constructed in which Q can take on any integer value.

DSMs capable of generating arbitrary ratios have been introduced in [62], in the patent literature [66, 109, 110] and more recently in [19]. These solutions are limited to accumulators with programmable size [19, 66] and DSMs with single-bit quantizers [62, 109, 110]. The literature to date does not provide any generally applicable theory governing the operation of variable-modulus DSMs, nor does it provide any examples of multi-bit variable modulus quantizers.

In this section we will present a way of modelling digital DSMs that takes into account the variable modulus possibility and will propose a family of variable modulus quantizers (VMQ) which can be used to convert any DSM topology into a variable modulus DSM. The quantizers presented here can have an arbitrary number of steps and are designed so that they can directly replace the traditional truncation quantizers. Practical applications of the variable modulus functionality will be discussed further in Ch. 5.

4.1 Quantization in the digital domain

4.1.1 Truncation quantizer

Coarse quantization of digital signals differs from the quantization of continuous amplitude analogue signals as discussed in Sec. 2.1. Since all digital signals are initially quantized, coarse quantization is just a way of reducing a signal's resolution. Quantization in the digital domain is most easily realized by truncating the digital bus. This is an attractive choice because it does not consume any hardware resources. Digital quantization, or re-quantization, can be performed simply by disregarding b LSBs, as shown Fig. 24. The remaining MSBs form the quantizer output q and the LSBs

disregarded form a signal r, which is a measure of the quantization error.

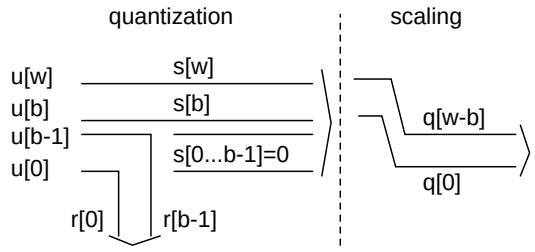

Fig 24. Truncation quantizer.

If the input u is a positive integer, the operation of the quantizer can be interpreted as division by 2^b, where r is the remainder after division. For all positive and negative integers u, the operation of the quantizer can be described by the following set of equations, where the "incomplete" brackets $\lfloor\ \rfloor$ denote the *floor* function:

$$\begin{cases} q &= \lfloor u \cdot 2^{-b} \rfloor \\ r &= u - 2^b \cdot q = u \bmod 2^b. \end{cases} \tag{28}$$

The input-output characteristic of the truncation quantizer is shown in Fig. 25. This quantizer can be implemented with the input u and output q, interpreted as two's complement signed numbers. The feedback signal should be interpreted as type unsigned or converted to two's complement representation by adding a zero at MSB.

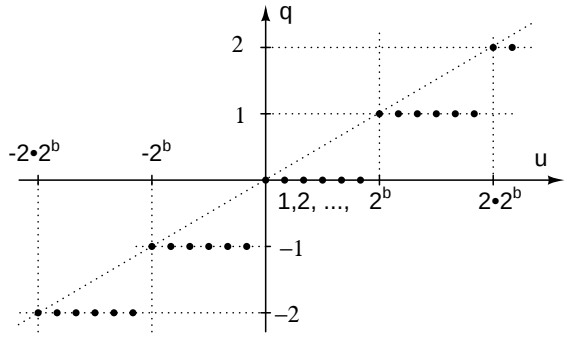

Fig 25. Input-output characteristic of the truncation quantizer.

The truncation quantizer can be described in terms of the traditional model introduced in Sec. 2.1, which assumes a linear coefficient G and the addition of quantization noise e:

$$\begin{aligned} q &= 2^{-b}(u-r) \\ &= 2^{-b}u - 2^{-b}\left(u \bmod 2^{b}\right) \\ &= G\cdot u + e. \end{aligned} \tag{29}$$

This is a convenient representation because it attributes all the nonlinearity to the properties of the quantization error e, while the quantizer itself is modeled by a linear gain. According to (29), the truncation quantizer introduces a linear gain:

$$G = 2^{-b}, \tag{30}$$

and the quantization error e is equal to:

$$e = -2^{-b}\left(u \bmod 2^{b}\right) = -r\cdot 2^{-b}. \tag{31}$$

The appearance of the non-unity gain and such a specific way of calculating the quantization error are consequences of interpreting the input u as an integer. This interpretation is a convention adopted here which allows simple generalization of the quantizer operation in Sec. 4.1.2 and has a very close and intuitive correspondence to the hardware implementation introduced in Sec. 4.3.

4.1.2 Arbitrary modulus quantizer

The output of the truncation quantizer increases by one each time the input reaches an integer multiple of 2^b. The operation of the quantizer can be generalized if we assume that the output increases each time the input reaches an integer multiple of an integer number Q. Direct modification of the truncation quantizer description (28) leads to the following:

$$\begin{cases} q &= \left\lfloor u\cdot Q^{-1}\right\rfloor \\ r &= u - Q\cdot q = u \bmod Q. \end{cases} \tag{32}$$

This modified quantizer can be interpreted as an arithmetic divider for positive integers, where r is the remainder from a division of u by Q. The modulus Q can take any positive integer value, and therefore this quantizer will be called the *arbitrary modulus quantizer*. The input-output characteristic of a five-step arbitrary modulus quantizer is shown in Fig. 26(a).

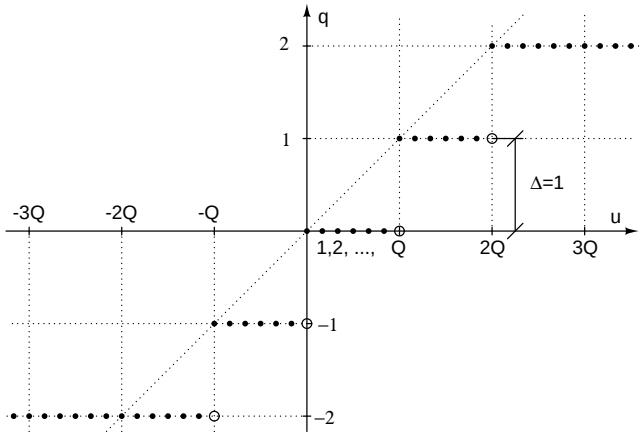

(a) Input-output characteristic.

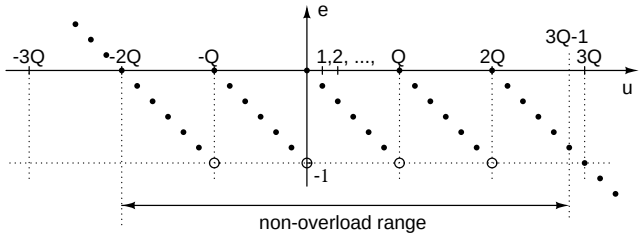

(b) Quantization error.

Fig 26. A five-step arbitrary modulus quantizer.

The arbitrary modulus quantizer can be described in terms of the traditional model, introduced in Sec. 2.1, which assumes a linear coefficient G and an addition of quantization noise e:

$$\begin{aligned} q &= Q^{-1}(u-r) \\ &= Q^{-1}u - Q^{-1}(u \bmod Q) \\ &= G\cdot u + e. \end{aligned} \tag{33}$$

According to (33), it introduces a linear gain:

$$G = Q^{-1}, \tag{34}$$

and the quantization error e is:

$$e = -Q^{-1}(u \bmod Q) = -Q^{-1}r. \tag{35}$$

This is the simplest generalization of the truncation quantizer. It could be converted to either a mid-tread or mid-rise quantizer by introducing an offset in the input-output characteristic to position it symmetrically with respect to the origin of the coordinate system, see [97]. This particular quantizer, however, has a simple mathematical description which serves well to demonstrate the role of the variable modulus.

4.2 Linear model of a DSM with an arbitrary modulus quantizer

A linear model of a first order error-feedback DSM with an arbitrary modulus quantizer is shown in Fig. 27. The quantizer is modelled by introducing a linear gain Q^{-1} with added quantization noise e, according to (33). The feedback signal r is calculated according to (32). The DSM output can be expressed in the z domain as:

$$y = \frac{x}{Q} + e\left(1 - z^{-1}\right). \tag{36}$$

This equation illustrates the principle of operation of a DSM with an arbitrary modulus quantizer. It shows that the DSM output signal y is composed of the input signal x divided by Q, with an added filtered quantization error e. According to the convention used in this work all signals within the DSM take on integer values.

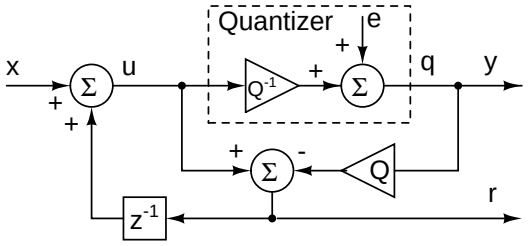

Fig 27. Linear model of a first-order DSM with an arbitrary modulus quantizer.

The effect of the variable modulus on the quantization noise is studied in greater detail in App. 2, where the quantization noise mean $\bar{e}$, variance σ^2 and power are derived on the assumption of a uniformly distributed quantization error. The results show that if the modulus Q is sufficiently large the quantization noise variance can be approximated

with the familiar value:

$$\sigma^2 \approx \frac{\Delta^2}{12}. \tag{37}$$

The quantization error is distributed among Q discrete values in the interval $\Delta = 1$, see Fig. 26(b). The quantization error variance already reaches 99 % of the limit value $\Delta^2/12$ with a DSM having a 4-bit bus ($Q = 2^4$). Therefore changing Q, in a case of a practical DSM, will not have any significant impact on the quantization noise variance. A further discussion of the selection of Q and the DSM bus width is presented in Sec. 5.3.2.

The relationship between the modulus and the sequence length was not studied in this work. A mathematical analysis of the sequence length for a MASH with prime modulus quantizers is presented in [71].

4.3 Quantizer implementations for variable modulus DSMs

DSMs can be designed either with hardwired or variable arbitrary modulus quantizers. This section will present practical implementations of single-step and multi-step variable modulus quantizers, as these are more versatile and can be easily converted to their hardwired equivalents. The example of a first-order DSM with a variable modulus quantizer shown in Fig. 28, has two inputs: one for the data signal x and the other for the modulus Q, which can be changed at any time during DSM operation.

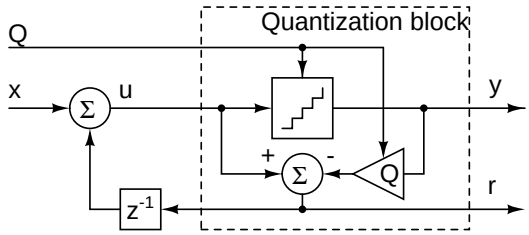

Fig 28. First-order DSM with an arbitrary modulus quantizer. Revised from [VII]. © [2007] IET.

The first-order DSM in Fig. 28 is presented in the error-feedback configuration. The quantization block, enclosed in a dashed-line box, contains both the quantizer and the additional hardware for computing the feedback signal r. The quantizer implementations presented in the subsequent sections provide the feedback signal r and can therefore

directly replace truncation quantizers in the error-feedback configurations. The proposed quantizers can be also adopted for other topologies such as output feedback, which do not require the signal r. For brevity we will refer collectively to the entire content of the quantization block as the "quantizer".

4.3.1 Single-bit quantizer

Truncation quantizer. Probably the simplest implementation of a first-order DSM with a single-bit truncation quantizer is that shown in Fig. 29. All the signals in the figure include their corresponding bus widths. The input signal $x[0...b-1]$, for example, is implemented on a b-bit bus with the LSB at position 0 and the MSB at position $b-1$.

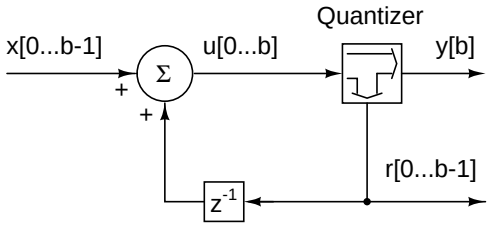

Fig 29. A first-order DSM with a single-bit truncation quantizer.

The quantizer is implemented as a bus-split and does not require any additional hardware resources. The DSM shown in Fig. 29 is equivalent to an accumulator where r can be interpreted as an arithmetic sum and y as the carry signal. This is a very popular choice for a hardware-efficient MASH topology. The quantizer in Fig. 29 can be described in the following way:

$$y = \begin{cases} 0 & u < 2^b \\ 1 & u \geq 2^b, \end{cases} \tag{38}$$

where the feedback signal r is equal to:

$$r = u - y \cdot 2^b. \tag{39}$$

This model uses only two levels of the quantizer input-output characteristic shown in Fig. 25, the level "zero" and the level "one", and consequently the non-overload range is $(0,1,\ldots,2\cdot 2^b-1)$. Since the feedback signal is in the range $r \in (0,1,\ldots,2^b-1)$, the input signal x has the same range.

Arbitrary modulus quantizer. The implementation of a single step variable modulus DSM, as depicted in Fig. 30 has two inputs: x for the data signal, and Q for the modulus. The variable modulus quantizer, enclosed in the dashed box, is designed so that it can directly replace the truncation quantizer shown in Fig. 29.

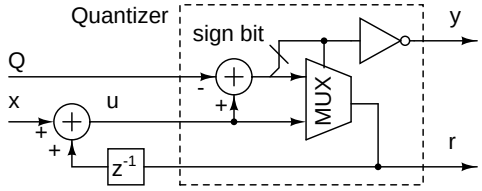

Fig 30. A first-order DSM with a single-bit variable modulus quantizer. Revised from [VII]. © [2007] IET.

The DSM output signal y is computed according to the rule:

$$y = \begin{cases} 0 & u < Q \\ 1 & u \geq Q. \end{cases} \tag{40}$$

The general rule for computing the feedback signal r is expressed by (32). Adapted to the notation used in Fig. 30, r can be expressed as:

$$r = u - Q \cdot y. \tag{41}$$

In order to avoid multiplication, this algorithm can be simplified in a practical implementation as follows:

$$r = \begin{cases} u & u < Q \\ u - Q & u \geq Q. \end{cases} \tag{42}$$

The algorithm described by (40) and (42) is implemented by means of a subtractor, an inverter and a multiplexer. The subtractor plays its nominal role and is also used as a comparator at the same time. The signals before and after the subtractor have two's complement representation. The sign bit interpreted as a comparison result controls the multiplexer and selects either the difference $u - Q \cdot 1$, or if $y = 0$, $u - Q \cdot 0$, simply u. The critical path for this implementation is the total delay introduced by the subtractor and multiplexer.

4.3.2 Multi-bit quantizer

Truncation quantizer. An arbitrary-order EFM with a multi-bit truncation quantizer is shown in Fig. 31 and its operation is described by (28). The number of bits selected for the output signal $y[b \dots w]$ is determined by the construction of the loop filter. Note that as long as the quantizer input u is in the non-overload range, the feedback signal r is bounded in the range $(0 \dots 2^b - 1)$. The feedback filter is typically implemented as an FIR which for a bounded input produces a bounded output. Assuming a bounded DSM input, the total bus width w can easily be determined by analysing the minimum sufficient bus widths for the arithmetic operations within the FIR filter. This reasoning shows that the DSM in this configuration produces a bounded output as long as the quantizer input remains in the non-overload range.

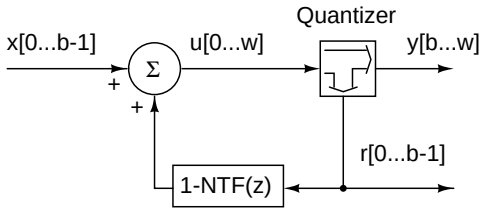

Fig 31. An EFM of arbitrary order with a multi-bit truncation quantizer.

Arbitrary modulus quantizer. Fig. 32 shows a five-step variable modulus quantizer which can directly replace the truncation quantizer in Fig. 31. This is an exact realization of the input-output characteristic shown in Fig. 26, with the same non-overload range. The algorithm for computing the output signal y and the feedback signal r is presented below:

$$y = \begin{cases} 2 & 2Q & \leq u & & d_{0,\dots,3} = 0,0,0,0 \\ 1 & Q & \leq u < & 2Q & d_{0,\dots,3} = 0,0,0,1 \\ 0 & 0 & \leq u < & Q & d_{0,\dots,3} = 0,0,1,1 \\ -1 & -Q & \leq u < & 0 & d_{0,\dots,3} = 0,1,1,1 \\ -2 & & u < & -Q & d_{0,\dots,3} = 1,1,1,1, \end{cases} \tag{43}$$

$$r = \begin{cases} u - 2Q & 2Q \leq u \\ u - Q & Q \leq u < 2Q \\ u & 0 \leq u < Q \\ u - (-Q) & -Q \leq u < 0 \\ u - (-2Q) & u < -Q. \end{cases} \tag{44}$$

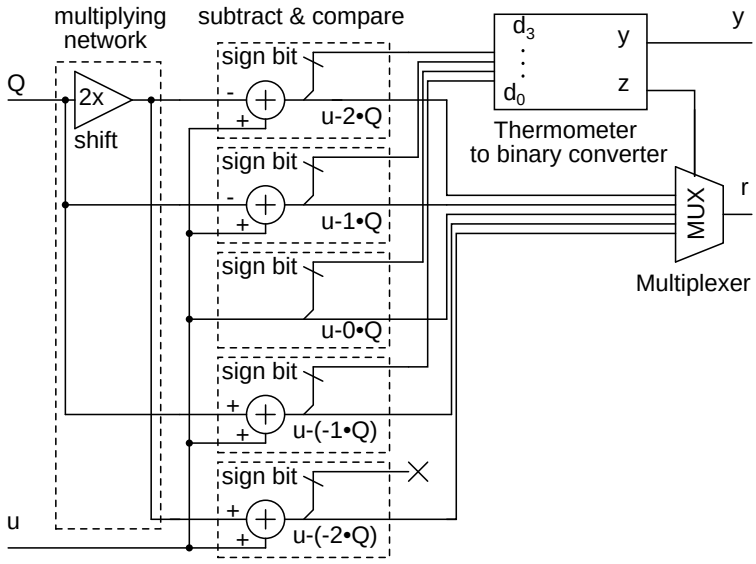

Fig 32. A five-step variable modulus quantizer. Revised from [VII]. © [2007] IET.

The quantizer shown in Fig. 32 is a natural extension of the single-bit quantizer shown in Fig. 30. With every additional quantization level there is one additional adder/comparator connected in parallel. This structure can be easily extended to provide more quantization steps. This construction ensures that the critical path does not increase significantly with each additional quantization level. The minimal implementation in Fig. 32 requires only a routing network to provide the multiples of Q. In any case, signals propagate through the multiplying network only when Q is changed. Therefore, although the multiplying network increases the hardware consumption, it does not contribute to the DSM's critical path.

4.3.3 Adapting quantizers to the output-feedback DSMs

The quantizers shown in Fig. 30 and Fig. 32 can easily be adapted to the output-feedback DSM topologies. In order to illustrate the required changes, the first order error-feedback DSM shown in Fig. 27 has been rearranged and presented in an equivalent output-feedback configuration in Fig. 33. As we can see the gain Q in the feedback must remain in order to compensate for the gain of the quantizer Q^{-1}.

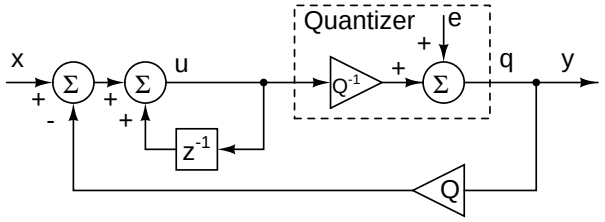

Fig 33. Linear model of a first-order DSM with an arbitrary modulus quantizer.

Consider the multi-bit quantizer implementation shown in Fig. 32. In the current configuration the multiplexer is connected to the adders and subtractors, so that the multiplexer's output is $u - y \cdot Q$. In order to adapt it to an output-feedback DSM the inputs to the multiplexer must be connected directly to the multiples of Q, so that the output will be equal to $y \cdot Q$. Since the computation of $u - y \cdot Q$ is no longer required, the adders and subtractors can be replaced with simple comparators. With these changes in place, the proposed quantizers can be applied to the output feedback DSM topologies.

4.4 Summary

We have introduced in this chapter a linear model of an arbitrary modulus DSM and have shown that for a DSM with an input DC value x and modulus Q the mean of the output is a rational number x/Q. In practice this means that by changing the modulus we can generate any rational mean with perfect accuracy. On the other hand, DSM implementations based on the modulus 2b can only approximate rational fractions with an accuracy that is related to the number of bits b. Thus arbitrary modulus DSMs have a potential for hardware savings, and represent a better choice for applications requiring perfect accuracy. The simple single-bit and multi-bit quantizer implementations

presented here illustrate the use of power-of-two and arbitrary moduli. A more exact comparison of the hardware consumption of the two quantizer types will be presented in the following chapter.

5 Practical design example

The frequency domain PLL model presented in [111] can be used to evaluate the contributions of various noise sources to the overall noise performance of a synthesizer. The model can be used to find the required noise transfer function and a suitable DSE order for a specific application. When the DSE topology and order have been selected based on a linear model, we have a certain expectation of how the noise will behave in the frequency domain. The examples presented in this section are based on a third-order MASH, which has been used previously in a frequency synthesizer for GSM 900 [39]. Once the topology and order have been selected, the DSE bus width should be determined. In this chapter we will show how to find the smallest possible bus width in order to minimize the area and power consumption.

This chapter focuses on the problem of DSE scaling with a truncation quantizer and variable modulus quantizer. It will show how to determine the smallest possible DSM bus width in order to meet the GSM BS requirements for frequency resolution and SFR, and will demonstrate the applicability of the results presented in Ch. 3 and Ch. 4. This chapter will conclude with a comparison of hardware consumption in MASH and EFM DSMs with truncation and arbitrary modulus quantizers.

5.1 Meeting RF channel accuracy specifications

A frequency synthesizer for a narrowband transceiver generates a set of frequencies corresponding to the available transmit and receive bands. It is usually sufficient for the synthesizer to generate the required frequencies with a certain finite accuracy, which can be used as the first indication of DSM scaling. The fundamental operation of a fractional-N frequency synthesizer has been described in Sec. 1.3, but we will repeat the main formulae here for the sake of clarity in the following analysis. The synthesizer output frequency f_o is a multiple of the reference frequency f_r and a number N which is composed of an integer part I and a fractional part x/Q. Thus,

$$\begin{aligned} f_o &= f_r \cdot N \\ &= f_r \cdot \left(I + \frac{x}{Q} \right). \end{aligned} \tag{45}$$

The fractional part is the DC component of the DSM output signal y and depends on the DSM input x and the DSM modulus Q (see Ch. 4). In a digital fixed point implementation all the numbers I, x and Q are integers and the fractional multiplier N can be any rational number. When the DSM input x is increased by one, the synthesizer output frequency increases by the smallest available quantity, called a *step size*, which can be expressed as:

$$\Delta f_o = \frac{f_r}{Q}. \tag{46}$$

The following sections will present two approaches to the selection of Q and the resulting step size, in order to satisfy the strict frequency accuracy specifications of the GSM system, as discussed further in App. 3. We will assume the following parameters: channel spacing $\Delta_{\text{ch}} = 200\,\text{kHz}$, frequency tolerance $\delta_f = 0.05\,\text{ppm}$, and a commonly assumed reference frequency value $f_r = 13\,\text{MHz}$ [28, 39, 51, 52].

5.1.1 Approximating RF channels

The most straightforward and most common approach to realizing a fractional-N frequency synthesizer is to use a high-resolution DSM and to approximate the required fraction with sufficient accuracy [39, 67]. This approach is based on the most hardware-efficient truncation quantizer, as discussed in Sec. 4.1.1. For such an implementation, the DSM resolution is a power of two, $Q = 2^b$, and the step size of the synthesizer can be expressed as:

$$\Delta f_o = \frac{f_r}{2^b}. \tag{47}$$

The tolerance of the output frequency for GSM is equal to $\delta_f = 0.05\,\text{ppm}$. This is a relative measure, and therefore the strictest requirements for an absolute frequency deviation expressed in Hz concern the lowest GSM bands. The smallest absolute frequency deviation for all bands is $\Delta f = 25\,\text{Hz}$ as derived in App. 3. The frequency step in a synthesizer approximating the required frequency should be smaller than or equal to twice the allowed frequency deviation, thus:

$$\Delta f_o \leq 2\Delta f. \tag{48}$$

This leads to a minimum required bus width for the DSM of:

$$b \geq \left\lceil \log_2 \left(\frac{f_r}{2\Delta f} \right) \right\rceil = \left\lceil \log_2 \left(\frac{13\,\text{MHz}}{50\,\text{Hz}} \right) \right\rceil = 18\,\text{bit}. \tag{49}$$

We will assume for the present purposes that the DSM bus width $b = 20$ bits, corresponding to the frequency step of the synthesizer $\Delta f_o \approx 12\,\text{Hz}$, which is closer to the actual values used in multi-band GSM transceivers [28, 39, 43]. The modulus Q will therefore be:

$$Q = 2^{20}. \tag{50}$$

In this case, the DSM input would be a 20-bit number in the range $0, 1, 2, \ldots, 2^{20} - 1$.

5.1.2 Generating RF channels with perfect accuracy

Instead of approximating the required frequencies, the RF channels can be generated with perfect accuracy[8]. This approach requires a certain amount of additional hardware to realize the arbitrary modulus quantizer described in Sec. 4.3. Using such a quantizer, we can match the synthesizer's step size Δf_o to the required channel spacing Δ_{ch}:

$$\Delta f_o = \Delta_{\text{ch}}. \tag{51}$$

According to (46), this leads to a condition for the DSM modulus Q:

$$Q = \frac{f_r}{\Delta_{\text{ch}}} = \frac{13\,\text{MHz}}{200\,\text{kHz}} = 65. \tag{52}$$

With the modulus $Q = 65$ and the reference frequency 13 MHz, the synthesizer output frequency will be increased by exactly 200 kHz each time the modulator input is increased by one. The output frequency of the synthesizer will be exactly

$$f_o = I \cdot 13\,\text{MHz} + x \cdot 200\,\text{kHz}. \tag{53}$$

Since the decimal number 65 is represented by 7 bits, i.e. $65_{10} = 1000001_2$, we need only 7-bits of DSM resolution to implement a quantizer with the modulus $Q = 65$. In this case the DSM input will be a 7-bit number in the range $0, 1, 2, \ldots, 64$, and we will have a 7-bit DSM with a certain amount of extra hardware for implementing the quantizer, as compared with a 20-bit DSM with a hardware-free quantizer as described in Sec. 5.1.1.

Generating the channels with perfect accuracy has one noteworthy consequence. If the generated frequency is not exactly equal to the nominal frequency, the phase error will accumulate, as the phase is an integral of the instantaneous frequency [108]. In

[8] This reasoning does not take into account the inherent instability and possible drift of the reference frequency generated by a quartz oscillator. This section focuses on the phase drift due to the DSM approximation alone.

other words, if the channel frequency is generated with an approximation error, however small, we will observe a phase drift. On the other hand, if the frequency is generated with perfect accuracy, as in this case, no phase drift will occur. An arbitrary modulus quantizer is the only choice for applications which do not tolerate phase drift.

5.2 Meeting spectrum specifications

The spectral requirements for a GSM transmitter are defined in the form of spectral masks that define the maximum permitted power levels both inside and outside the transmitting band (see Fig. 37 in App. 3). These spectral masks are obtained by measuring the power in a given bandwidth at various frequencies. The minimum measurement bandwidth $\mathrm{BW}_{\mathrm{min}}$ for qualifying the transmitter can be used to determine the minimum sufficient tone spacing due to the finite DSE output signal period, and in consequence the minimum sufficient DSM bus width. The minimum bandwidth used when measuring spurious emissions from a GSM BS transmitter is 30 kHz.

As has been shown in Ch. 3, the spectrum of a digital DSM working with a DC input is inherently tonal, where the tone spacing Δf_k is proportional to the reference frequency and inversely proportional to the DSM sequence length, i.e.

$$\Delta f_k = \frac{f_r}{L_s}. \tag{54}$$

We can control not only the spacing of the tones, but also their average level. As shown in Fig. 22, the average level of the tones decreases by 3 dB each time the DSE bus width is increased by one bit. The spectral envelope is regular, and the total output power is the same as that for the ideal model using the CMQ. The spectrum level control phenomenon can be observed in a spectrum analyzer only when the tone spacing is greater than the RBW of the instrument. In the case of a GSM BS there is no sense in increasing the DSE bus width any further once the tone spacing reaches 30 kHz, so that beyond this point the tones are only averaged within the resolution bandwidth. This principle is illustrated in Fig. 34, which shows idealized DSM quantization error spectra for two sequence lengths compared with the resolution bandwidth $\mathrm{RBW} = 30\,\mathrm{kHz}$, and assuming a reference frequency $f_r = 13\,\mathrm{MHz}$.

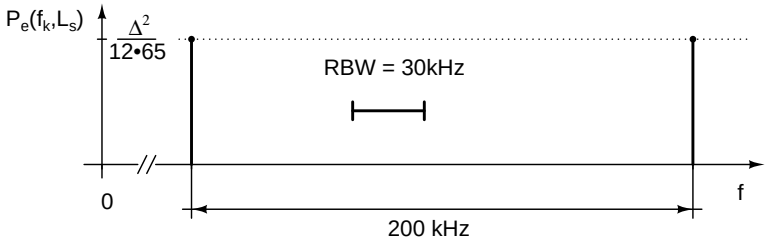

(a) Insufficient tone spacing.

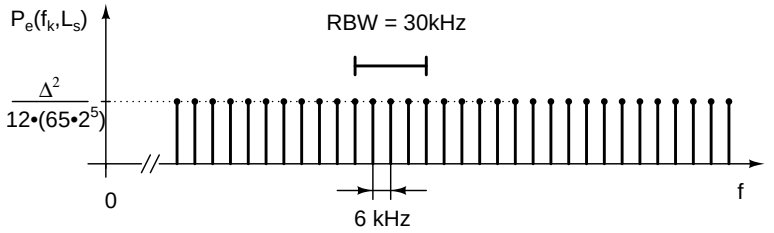

(b) Sufficient tone spacing.

Fig 34. Selecting a sufficient number of tones for the specification of GSM spurious emissions. The RBW is set to the minimum bandwidth used when measuring spurious emissions from a GSM BS transmitter.

The first example, illustrated in Fig. 34(a), shows a tone spacing that is too large. The sequence length $L_s = 65$ is very short and therefore the tone spacing Δf_k is much larger than the resolution bandwidth: $\Delta f_k = 13\,\text{MHz}/65 = 200\,\text{kHz} \gg \text{RBW}$. The tones will be detected clearly and their levels will be $\Delta^2/(12 \cdot 65)$, as suggested in Sec. 3.1.1. The second example illustrated in Fig. 34(b) shows a sufficiently small tone spacing for this application. The sequence length $L_s = 65 \cdot 2^5$ is sufficiently large and results in a tone spacing that is smaller than the resolution bandwidth $\Delta f_k = 13\,\text{MHz}/(65 \cdot 2^5) = 6\,\text{kHz} < \text{RBW}$. The absolute levels of the tones are smaller than in the previous example, being $\Delta^2/(12 \cdot 65 \cdot 2^5)$. No tonal spectrum will be detected in the measurement in this case. In fact, a further increase in the sequence length would not bring about much improvement. The power in the tones is simply averaged in the RBW, and consequently the spectrum will not be distinguishable from a truly continuous one.

We assumed here that the DSE is properly controlled and that the spectrum is as

near as possible to that of the idealized model described in Ch. 3. The power does not concentrate in the largely dominant tones and the spectrum tonality results from the finite sequence length. If a certain number of tones was averaged within the RBW, the variation seen in Fig. 20 would disappear and the practical DSE would closely approximate the ideal DSE spectrum. Based on the discussion in the previous paragraph, we will arbitrarily assume that $\Delta f_k = 6\,\text{kHz}$ is the maximum sufficient tone spacing for the 30 kHz GSM measurement bandwidth. According to (54), this leads to the shortest acceptable sequence length for this application:

$$L_{s\,\min} = 65 \cdot 2^5 = 2080. \tag{55}$$

5.3 Scaling DSE

In this section we will calculate the minimum required bus width for a DSE with a truncation quantizer or a variable modulus quantizer based on the channel accuracy and spectral requirements for GMS as detailed in Sec. 5.1 and Sec. 5.2.

5.3.1 DSM with truncation quantizer

A DSM with a truncation quantizer can be used to approximate the required RF channel frequencies. In order to meet the frequency accuracy specification, the bus width should be at least 20 bits, as shown in Sec. 5.1.1, which corresponds to a modulus $Q = 2^{20}$. If the initial conditions are set appropriately, a third-order MASH or EFM implemented with this bus width will generate repeating sequences of length $L_s = 2^{b+1} = 2^{21}$, as shown in Sec. 3.2.1, as a result of which the spectrum will be composed of 2^{21} discrete tones with a spacing of $\Delta f_k = 13\,\text{MHz}/2^{21} = 6\,\text{Hz}$. This means that over 5000 tones will be averaged within the measurement bandwidth of 30 kHz. This is easily a sufficient sequence length and the measurement will not reveal the tonal nature of the DSM output signal. This rationale shows that when a DSM is implemented in the traditional way, the strict GSM frequency accuracy specification will predominate in the question of bus width selection. The DSM will automatically meet the spectral resolution requirements.

It is interesting to note that when the channels are approximated the exact mean is achieved over the entire sequence of 2^{21} samples. Therefore, in order to use the full resolution, one has to wait for $2^{21}/13\,\text{MHz} = 0.16\text{s}$. From this point of view, extremely long sequences might not be desirable. If it takes an excess amount of time to convey

information about the exact mean value, “fine resolution” may never be achieved in practice. This observation is only intuitive and should be verified further.

5.3.2 An arbitrary modulus DSM

If an arbitrary modulus DSM is based on $Q = 65$, it can generate 200 kHz spaced GSM channels with perfect accuracy. A DSM with this modulus would require a 7-bit bus, as shown in Sec. 5.1.2. A third-order EFM with a 7-bit bus will generate output sequences of exactly 65 when implemented according to the methodology described in [VI] and [VII], resulting in a tone spacing $\Delta f_k = 200\,\text{kHz}$. Discrete tones with such a large spacing can easily be seen on a spectrum analyzer with a 30 kHz resolution bandwidth. In order to meet the spectral requirements the modulus must be increased to at least $Q = 65 \cdot 2^5$, see (55) in Sec. 5.2, which can be achieved by adding 5 LSBs. Following the sequence length control procedure described in [VI] we obtain a sequence length of $L_s = Q = 65 \cdot 2^5$, which results in 6 kHz spaced tones, as illustrated in Sec. 5.2. Overall, we obtain a $7 + 5 = 12$-bit DSM which generates the required channels with perfect accuracy and has a sufficiently dense spectrum, which will be indistinguishable from a continuous spectrum in a spectrum analyzer with a 30 kHz resolution bandwidth. In addition, the sequence length is relatively short and information on the perfect-precision mean is transferred in $(65 \cdot 2^5)/13\,\text{MHz} = 160\,\mu s$, 1000 times faster than in the case of a truncation quantizer DSM for the same application. The differences between two DSMs implemented with a variable modulus quantizer and truncation quantizer are summarized in Table 3.

Table 3. Comparison of the performance of two quantizers in a DSE for use in fractional-N frequency synthesis.

DSE implementation	Bus width	Max. freq. error	Tone spacing Δf_k	Sequence Length L_s
Truncation Quantizer	20-bit	6 Hz	6 Hz	160 ms
Variable Modulus Quantizer	12-bit	0	6 kHz	160 μs

5.4 DSE implementation

This section will evaluate the hardware consumption and speed of the third-order MASH and EFM architectures introduced in Ch. 3 when implemented with truncation and variable modulus quantizers. Both modulators were designed with the Xilinx® System Generator for Matlab® and targeted for a Xilinx® Virtex-4 FPGA (see App. 4). The implementation results for both topologies are gathered together in Table 4 and Table 5. Both modulators were scaled according to the principles outlined in Sec. 5.3.

Table 4. Hardware and speed costs of an Error Feedback DSM with different types of quantizer.

Q	Bus width [bits]	FFs	4-input LUTs	Data Path Delay [ns]			Max. Clock Frequency [MHz]
				Logic	Routing	Total	
2^{20}	20	60	181	6.9	5.9	12.8	78.0
65	20	60	501	10.0	11.1	21.0	47.4
65	12	36	327	8.8	10.2	19.0	52.8

Table 5. Hardware and speed costs of a MASH DSM with different types of quantizer.

Q	Bus width [bits]	FFs	4-input LUTs	Data Path Delay [ns]			Max. Clock Frequency [MHz]
				Logic	Routing	Total	
2^{20}	20	64	74	5.9	3.8	9.7	102.0
65	20	64	200	11.3	9.1	20.4	48.8
65	12	40	128	10.3	8.5	18.8	53.3

Both tables are organized in a similar way. Let us consider Table 4 as an example. The first row shows the implementation details for a DSM implemented with a 20-bit truncation quantizer. Such a modulator approximates the desired fractions with an accuracy that is related to the bus width (see Sec. 5.3.1). The second row shows the

same 20-bit DSM but with a variable modulus quantizer. This is the most flexible implementation, as it allows any modulus to be used within the 20-bit range, but the price of this flexibility is a three-fold increase in the asynchronous logic, expressed in the number of lookup tables (LUTs), and a 40 % decrease in the maximum achievable frequency. If such flexibility in the selection of the modulus were not required, the DSM bus width could be optimized as discussed in Sec. 5.3.2, whereupon the DSM would be capable of generating a series of fractions based on a selected denominator with perfect accuracy. Reducing the bus width will naturally lessen the number of flip-flops (FFs) and the number of LUTs and improve the maximum achievable frequency and the power consumption. Similar results, for the more hardware-efficient MASH are gathered together in Table 5.

The implementation results show that the relationship between the SFR and bus width introduced in Ch. 3 can be used to optimize the hardware resources, speed and power consumption.

5.5 Summary

This section has discussed the problem of DSE sizing in fractional-N frequency synthesis. After choosing a topology according to the model assuming white quantization noise, we have to select an appropriate DSE bus width. It is desirable to minimize the bus width in order to save area and minimize energy consumption. The reasoning behind selection of the appropriate bus width is presented for a DSE with either a truncation or arbitrary modulus quantizer, both sized to meet the specifications of a GSM transmitter.

The two DSE types scale differently in a system which requires high accuracy in the frequencies generated. A frequency synthesizer based on a truncation quantizer DSE only approximates the desired frequency with a prescribed accuracy which depends on the DSE bus width, and therefore the scaling is dominated by the required accuracy. For systems requiring high frequency accuracy, such as GSM, the DSE bus width will be large and consequently the tonal DSE spectrum will be very dense. On the other hand, a frequency synthesizer with an arbitrary modulus DSE can generate any desired frequency spacing with perfect accuracy. The bus width of an arbitrary modulus DSE must be large enough to implement the required modulus, but also large enough to produce a sufficiently dense spectrum. In the example presented in this section the minimum measurement bandwidth used in the GSM was taken as the predominant specification for scaling the arbitrary modulus DSE. As a consequence the variable

modulus quantizer can be implemented with a smaller bus width than the traditional DSM based on the truncation quantizer.

The narrow bus of an arbitrary modulus DSE might suggest a smaller area, but the example shows that this is not necessarily the case. The use of flip-flops depends directly on the DSE bus width and is indeed smaller for the arbitrary modulus DSE. On the other hand, a substantial amount of asynchronous logic is required to facilitate the arbitrary modulus, which limits the maximum clock frequency and means that an arbitrary modulus quantizer is a speed bottleneck for a DSE. An arbitrary modulus DSE can nevertheless be particularly hardware-efficient in output-feedback topologies with a limited number of quantization steps. In general, an arbitrary modulus DSE is irreplaceable in applications requiring perfect frequency accuracy and zero phase drift.

6 Overview of the contributing papers

The first two publications included in this thesis present a frequency synthesizer test platform composed of a PLL ASIC [I] and an FPGA hosting a delta-sigma modulator [II]. The platform was intended for testing the synthesizer's custom switching speed-up subsystems and various DSM structures in fractional-N mode. Measurements revealed that the tonal behaviour of a DSM is a significant challenge in fractional-N mode, although our research also revealed that digital DSMs with DC inputs display very regular behaviour with respect to the DC input and initial conditions. Such behaviour had not been discussed in the literature prior to our work and therefore inspired a more thorough study, which laid the foundations for Ch. 3 of this thesis.

The initial study [II] had already shown that when a third-order EFM is started from particular initial conditions the sequence length and the number of discrete tones in the spectrum will depend on the bus width and not on the modulator input. We have also shown for this architecture that each time the bus width is extended by one bit the SFR is improved by 3 dB. The performance of a DSM designed for a constant sequence length can be completely verified for all inputs.

The following work [III] extended the DSM cycle length analysis to a third-order MASH-1-1-1 and considered the effects of the DSM out-of-band noise on the performance of the synthesizer. We showed that a simple filter at the DSM output with one zero at the sampling frequency can improve the noise performance of the synthesizer, since, although it increases the in-band noise by 6 dB, it achieves the desired narrower phase noise distribution and reduces the charge pump ON time by 20%. As a result, the synthesizer is more immune to nonlinearities in the charge pump and to noise coupling during the ON time.

The DSM sequence length control method initiated in [II] was extended to cover MASH and EFM on a broader scale in [IV]. Groups of initial conditions that led to desired sequence lengths were identified in a general way and equations for sequence length were provided for architectures from orders two to five. Furthermore, a method was demonstrated that allows loading of the initial conditions at the system's start time and avoids reloading each time the DC input changes. As a consequence, the internal state of the DSM always remains in the desired sub-space. This leads to very long output sequences. These two practical methods of controlling the sequence length are

discussed in greater detail in Sec. 3.2.1.

The outcome of papers [II], [III], [IV] was a method for coping with the well-known DSM spurious tone problem, which is traditionally addressed using a method employing whitening noise, called dithering. Our proposed method based on the initial conditions was compared with a dither signal added to the EFM immediately before the quantizer in [V]. Three types of dither generation were considered, and it was shown that the spectrum obtained by applying predefined initial conditions is as good as that obtained using the best-known form of dither, while the proposed solution consumes less hardware.

A complete practical DSM design method based on the behaviour of limit cycles was published in [VI], which also presented the previously published results [II], [III], [IV], [V] in a more complete and unified manner and proposed a complete solution to the tonal problem for two classical DSM topologies. This paper shows that, by loading specific initial conditions and scaling the DSM bus width it is possible to achieve precise control over the level of DSM noise. At the same time the noise spectrum will remain very regular and free from dominant spurious tones.

The second contribution to digital DSMs, covered in Ch. 4 of this thesis, was introduced in [VII], in which we showed that coarse quantization in the digital domain is equivalent to arithmetic division and proposed practical means for changing the divisor, referred to as modulus. As a result the DC value of the DSM output can be any rational fraction generated with absolute accuracy. When such a DSM is used for frequency synthesis it allows individual frequencies to be generated without any approximation, so that they will not suffer from constant phase drift.

7 Conclusions

7.1 Summary

Two techniques are described here that improve the performance of digital delta-sigma modulators and extend their capabilities. Both techniques are suitable for modulators implemented in a fixed-point digital environment and are applicable to fractional-N frequency synthesis.

The first technique addresses the problem of quantization noise in digital DSMs with DC inputs. Our proposed approach ensures a smooth spectral envelope and allows the noise level to be traded off against DSM hardware consumption. We have shown for MASH and error feedback that the output sequence length can be controlled by loading pre-defined initial conditions and by scaling the bus width. As a result, it is possible to distribute the quantization noise power among a controlled number of tones. In effect, the SFR depends on the DSM bus width and not on the DC input. This technique enables one to avoid over-designing and can guarantee the desired performance under all practical conditions.

The second contribution is a proposal for a general family of digital quantizers with variable modulus. A DSM with such a quantizer will have two inputs, one for the data and the other for the modulus. A variable modulus will expand the DSM capability for generating the mean value as any rational fraction with perfect accuracy, within reasonable limits of practical implementation. This functionality is desirable in frequency synthesis applications as it allows sets of frequencies to be generated with perfect accuracy. Since the frequencies are generated exactly, without any approximation, the commonly observed phase drift can be avoided. The proposed practical implementation is general in nature, compact, scalable and applicable to many digital DSM topologies with any number of output bits.

7.2 Discussion

Sequence length control. The first contribution of this thesis addresses DSM's tonal behaviour. The idea of using predefined initial conditions as a means of ensuring a spurious tone-free DSM spectrum is not new, and was proposed in [41, 49, 66–69]. What

is new is the observation that such means allow the sequence length to be controlled precisely for all DC inputs. We have shown for the first time in [III], [IV], [V], [VI] that it is possible to control the sequence length and that this entails a number of benefits.

First of all, controlling the sequence length opens up a new perspective on noise shaping in the digital domain. Noise can not only be shaped, but it can also be distributed among a well-defined number of tones. The quantization noise power is always constant, and it is therefore desirable to maximize the number of tones in order to minimize their average level.

Sequence length can be used as a benchmark for DSMs. The longer the sequences that can be generated with a given amount of hardware, the better. The desire to generate very long sequences has led to the construction of novel, hardware-efficient topologies which produce significantly longer sequences than traditional MASH and error-feedback structures [71, 72]. These topologies score very highly in terms of the sequence length "benchmark", as they have the ability to utilize almost the entire available state space. The same line of thought has also allowed the amount of hardware in MASH to be reduced without sacrificing the achievable sequence length, and in consequence without compromising on performance [106]. Controlling the sequence length is a form of deterministic approach to controlling the spectral performance of a DSM, i.e. it is a form of solution "by design" [VI], [72].

It should be admitted, however, that the relationship between long sequences and a good-quality spectrum is not a simple one. A long sequence yields a large number of discrete tones, but it does not necessarily guarantee how the quantization noise power will be distributed among those tones. All the available studies have nevertheless shown that when the DSM is forced to produced long sequences the desired power distribution follows [VI], [71, 72]. Moreover, the solution based on controlled sequence length is entirely deterministic and the spectrum can be tested beyond any doubt under all practical conditions [VI].

Variable modulus DSM. Fractional-N frequency synthesis has a built-in capacity to generate any rational multiples of the reference frequency, see (2) in Sec. 1.3.1. It is simply a matter of maintaining proper control over the frequency divider. One known solution is based on accumulators of programmable size [19, 66], but the same effect can be obtained with a DSM that controls the division ratio. The known solutions are limited to DSMs with single-bit quantizers [62, 109, 110].

This work proposes a unified means of analysis and a practical solution that is

applicable to many DSM topologies. It is argued that the desired functionality can be achieved by appropriately constructing a digital quantizer. This new approach points to a unified way of modelling and interpreting the operation of a variable modulus DSM, and consequently leads us to propose a family of variable modulus quantizers which are easily scalable and can have any arbitrary number of output bits. It is also noted that as a by-product these quantizers can compute the quantization error.

Our proposed family of arbitrary modulus quantizers is based on truncation quantizers characterized by a DC offset in the quantization error, but can easily be converted to more traditional mid-tread or mid-rise quantizers.

7.3 Future work

Although the methodology described in Ch. 3 is applicable to DSMs working with DC inputs, in which the limiting factor is the requirement to reload the initial conditions each time the DC input changes, it has an interesting extension which is described in greater detail in [IV]. This second approach requires one to set the initial conditions only at power-up. The DSM can work with arbitrary inputs, but whenever the input has a DC-like segment, the DSM spontaneously generates a very long output sequence. This approach could potentially be used to address the tonal problem in DSMs working with arbitrary inputs. Further studies would be required to prove the usefulness of this method and to determine whether it can be applied to a broader class of digital DSMs.

When we forced the digital DSMs to produce very long periodic sequences in this work, the spectrum associated with the sequences was smooth and tone-free. Sequence length and the distribution of spectral power are two different things, however, and the relationship between a long sequence and the desired uniform power distribution has only been observed empirically. This relationship should be studied further.

References

1. Larson LE (2003) Silicon technology tradeoffs for radio-frequency/mixed-signal ”systems-on-a-chip”. IEEE Trans. Electron Devices 50(3): 683–699.
2. Bennett HS, Brederlow R, Costa JC, Cottrell PE, Huang WM, Immorlica AA, Mueller JE, Racanelli M, Shichijo H, Weitzel CE & Zhao B (2005) Device and technology evolution for Si-based RF integrated circuits. IEEE Trans. Electron Devices 52(7): 1235–1258.
3. Abidi AA (2003) RF CMOS come of age. IEEE Microwave 4(4): 47–60.
4. Upton A & Steel V (2006) The current state of technology and future trends in wireless communications and applications. Microwave Journal 49(9): 22–38.
5. Muhammad K, Staszewski R & Leipold D (2005) Digital RF processing: toward low-cost reconfigurable radios. IEEE Commun. Mag. 43(8): 105–113.
6. Galton I (2002) Delta-sigma data conversion in wireless transceivers. IEEE Trans. Microwave Theory Tech. 50(1): 302–315.
7. Spang H III & Schultheiss P (1962) Reduction of quantizing noise by use of feedback. IEEE Trans. Commun. 10(4): 373–380.
8. Inose H & Yasuda Y (1963) A unity bit coding method by negative feedback. Proc. IEEE 51(11): 1524–1535.
9. Candy J (1974) A use of limit cycle oscillations to obtain robust analog-to-digital converters. IEEE Trans. Commun. 22: 298– 305.
10. Gray RM (1987) Oversampled sigma-delta modulation. IEEE Trans. Commun. 35: 481–489.
11. Boser BE & Wooley BA (1988) The design of sigma-delta modulation analog-to-digital converters. IEEE J. Solid-State Circuits 23(6): 1298–1308.
12. Candy J & Huynh AN (1986) Double interpolation for digital-to-analog conversion. IEEE Trans. Commun. 34(1): 77–81.
13. Naus P, Dijkmans E, Stikvoort E, Holland D & Bradinal W (1987) A CMOS stereo 16-bit D/A converter for digital audio. IEEE J. Solid-State Circuits 22(3): 390–395.
14. Uchimura K, Hayashi T, Kimura T & Iwata A (1988) Oversampling A-To-D and D-To-A converters with multistage noise shaping modulators. IEEE Trans. Acoust., Speech, Signal Processing 36(12): 1899–1905.
15. Candy JC & Temes GC (1992) Oversampling Delta-Sigma Data Converters: Theory, Design, and Simulation. IEEE Press.
16. Norsworthy SR, Schreier R & Tames GC (1997) Delta-Sigma Data Converters. IEEE Press, New York.
17. Schreier R & Temes GC (2004) Understanding Delta-Sigma Data Converters. 0471465852. Wiley-IEEE Press, 1 edition.
18. Egan WF (2000) Frequency Synthesis by Phase Lock. John Wiley and Sons, 2 edition.
19. Rogers J, Plett C & Dai F (2006) Integrated Circuit Design for High-Speed Frequency synthesis. Artech House, Inc.
20. Proakis JG & Manolakis D (1996) Digital Signal Processing - Principles, Algorithms and Applications. Prentice Hall, 3 edition.
21. van Engelen J (1999) Stability analysis and design of bandpass sigma delta modulators. monograph, Technische Universiteit Eindhoven.
22. Abidi AA (2004) RF CMOS comes of age. IEEE J. Solid-State Circuits 39(4): 549–561.

23. Lessing L (1956) Man of High fidelity: Edwin Howard Armstrong. J.B. Lippincott Company, Philadelphia.
24. Abidi AA (1995) Direct-conversion radio transceivers for digital communications. IEEE J. Solid-State Circuits 30(12): 1399–1410.
25. Razavi B (1998) RF Microelectronics. Prentice-Hall Communications Engineering & Emerging Technologies. Prentice Hall.
26. Loke A & Ali F (2002) Direct conversion radio for digital mobile phones - design issues, status, and trends. IEEE Trans. Microwave Theory Tech. 50(11): 2422–2435.
27. Darabi H, Khorram S, Chien HM, Pan MA, Wu S, Moloudi S, Leete JC, Rael JJ, Syed M, Lee R, Ibrahim B, Rofougaran M & Rofougaran A (2001) A 2.4-GHz CMOS transceiver for bluetooth. IEEE J. Solid-State Circuits 36(12): 2016–2024.
28. Magoon R, Molnar A, Zachan J, Hatcher G & Rhee W (2002) A single-chip quad-band (850/900/1800/1900 MHz) direct conversion GSM/GPRS RF transceiver with integrated VCOs and fractional-N synthesizer. IEEE J. Solid-State Circuits 37(12): 1710–1720.
29. Perraud L, Recouly M, Pinatel C, Sornin N, Bonnot JL, Benoist F, Massei M & Gibrat O (2004) A direct-conversion CMOS transceiver for the 802.11a/b/g WLAN standard utilizing a cartesian feedback transmitter. IEEE J. Solid-State Circuits 39(12): 2226–2238.
30. Song E, Koo Y, Jung YJ, Lee DH, Chu S & Chae SI (2005) A 0.25-μm CMOS quad-band GSM RF transceiver using an efficient LO frequency plan. IEEE J. Solid-State Circuits 40(5): 1094–1106.
31. Kaczman D, Shah M, Godambe N, Alam M, Guimaraes H, Han L, Rachedine M, Cashen D, Getka W, Dozier C, Shepherd W & Couglar K (2006) A single-chip tri-band (2100, 1900, 850/800 MHz) WCDMA/HSDPA cellular transceiver. IEEE J. Solid-State Circuits 41(5): 1122–1132.
32. Rudell JC, Ou JJ, Cho TB, Chien G, Brianti F, Weldon JA & Gray PR (1997) A 1.9-GHz wide-band IF double conversion CMOS receiver for cordless telephone applications. IEEE J. Solid-State Circuits 32(12): 2071–2088.
33. van Zeijl P, Eikenbroek JWT, Vervoort PP, Setty S, Tangenberg J, Shipton G, Kooistra E, Keekstra IC, Belot D, Visser K, Bosma E & Blaakmeer SC (2002) A Bluetooth radio in 0.18-μm CMOS. IEEE J. Solid-State Circuits 37(12): 1679–1687.
34. Cho TB, Kang D, Heng CH & Song BS (2004) A 2.4-GHz dual-mode 0.1 8-pm CMOS transceiver for bluetooth and 802.11b. IEEE J. Solid-State Circuits 39(11): 1916–1926.
35. Ahola R, Aktas A, Wilson J, Rao KR, Jonsson F, Hyyrylainen I, Brolin A, Hakala T, Friman A, Makiniemi T, Hanze J, Sanden M, Wallner D, Guo Y, Lagerstam T, Noguer L, Knuuttila T, Olofsson P & Ismail M (2004) A single-chip CMOS transceiver for 802.11a/b/g wireless LANs. IEEE J. Solid-State Circuits 39(12): 2250–2258.
36. Filiol N, Birkett N, Cherry J, Balteanu F, Cojocaru C, Namdar A, Pamir T, Sheikh K, Glandon G, Payer D, Swaminathan A, Forbes R, Riley T, Alinoor S, Macrobbie E, Cloutier M, Pipilos S & Varelas T (2001) A 22 mW bluetooth RF transceiver with direct RF modulation and on-chip IF filtering. In: Solid-State Circuits Conference, 2001. Digest of Technical Papers. ISSCC. 2001 IEEE International, pp. 202–203,447.
37. Muer BD & Steyaert MSJ (2002) A CMOS monolithic Δ-Σ-controlled fractional-N frequency synthesizer for DCS-1800. IEEE J. Solid-State Circuits 37: 835–844.
38. Rhee W, Bisanti B & Ali A (2002) An 18-mW 2.5-GHz/900-MHz BiCMOS dual frequency synthesizer with $<$10-Hz RF carrier resolution. IEEE J. Solid-State Circuits 37: 515–520.
39. Hegazi E & Abidi A (2003) A 17-mW transmitter and frequency synthesizer for 900-MHz

GSM fully integrated in 0.35-um CMOS. IEEE J. Solid-State Circuits 38(5): 782–792.
40. Pamarti S, Jansson L & Galton I (2004) A wideband 2.4-GHz delta-sigma fractional-N PLL with 1-MB/s in-loop modulation. IEEE J. Solid-State Circuits 39(1): 49–62.
41. Riley T & Copeland M (1994) A simplified continuous phase modulator technique. IEEE Trans. Circuits Syst. II 41(5): 321–328.
42. Perrott MH, Tewksbury TL & Sodini CG (1997) 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.
43. Lee ST, Fang SJ, Allstot D, Bellaouar A, Fridi A & Fontaine P (2004) A quad-band GSM-GPRS transmitter with digital auto-calibration. IEEE J. Solid-State Circuits 39(12): 2200–2214.
44. McMahill D & Sodini C (2002) A 2.5-Mb/s GFSK 5.0-Mb/s 4-FSK automatically calibrated Σ-Δ frequency synthesizer. IEEE J. Solid-State Circuits 37(1): 18–26.
45. Meninger S & Perrott M (2006) 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.
46. Riley TAD & Kostamovaara J (2003) A hybrid Δ-Σ fractional-n frequency synthesizer. IEEE Trans. Circuits Syst. II 50: 176–180.
47. Crawford JA (1994) Frequency Synthesizer Design Handbook. Artech House, Norwood, MA.
48. Riley TAD, Filol N, Du Q & Kostamovaara J (2003) Techniques for in-band phase noise reduction in Δ-Σ synthesizers. IEEE Trans. Circuits Syst. II 50: 794–803.
49. Filiol NM, Riley TAD, Plett C & Copeland MA (1998) An agile ISM band frequency synthesizer with built-in GMSK data modulation. IEEE J. Solid-State Circuits 33: 998–1008.
50. Quinlan P, Crowley P, Chanca M, Hudson S, Hunt B, Mulvaney K, Retz G, O'Sullivan C & Walsh P (2004) A multimode 0.3-200-kb/s transceiver for the 433/868/915-MHz bands in 0.25-μm CMOS. IEEE J. Solid-State Circuits 39(12): 2297–2310.
51. Lee HI, Cho JK, Lee KS, Hwang IC, Ahn TW, Nah KS & Park BH (2004) A Σ-Δ fractional-N frequency synthesizer using a wide-band integrated VCO and a fast AFC technique for GSM/GPRS/WCDMA applications. IEEE J. Solid-State Circuits 39(7): 1164–1169.
52. Fahim A & Elmasry M (2003) A wideband sigma-delta phase-locked-loop modulator for wireless applications. IEEE Trans. Circuits Syst. II 50(2): 53–62.
53. Best RE (2003) Phase-Locked Loops. McGraw-Hill, 5 edition.
54. Staszewski RB, Leipold D, Muhammad K & Balsara PT (2003) Digitally controlled oscillator (DCO)-based architecture for RF frequency synthesis in a deep-submicrometer CMOS process. IEEE Trans. Circuits Syst. II 50: 815–828.
55. Hung CM, Staszewski R, Barton N, Lee MC & Leipold D (2006) A digitally controlled oscillator system for SAW-less transmitters in cellular handsets. IEEE J. Solid-State Circuits 41(5): 1160–1170.
56. Staszewski RB, Muhammad K, Leipold D, Hung CM, Ho YC, Wallberg JL, Fernando C, Maggio K, Staszewski R, Jung T, Koh J, John S, Deng IY, Sarda V, Moreira-Tamayo O, Mayega V, Katz R, Friedman O, Eliezer OE, De-Obaldia E & Balsara PT (2004) 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.
57. Staszewski RB & Balsara PT (2005) Phase-domain all-digital phase-locked loop. IEEE Trans. Circuits Syst. II 52(3): 159–163.

58. Staszewski R, Wallberg J, Rezeq S, Hung CM, Eliezer O, Vemulapalli S, Fernando C, Maggio K, Staszewski R, Barton N, Lee MC, Cruise P, Entezari M, Muhammad K & Leipold D (2005) All-digital PLL and GSM/EDGE transmitter in 90nm CMOS. In: Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International, pp. 316–600Vol.1.
59. Gillette G (1969) Digiphase synthesizer. In: 23rd Annual Symposium on Frequency Control, pp. 201–210.
60. Reinhardt V, Gould K, McNab K & Bustamante M (1986) A short survey of frequency synthesizer techniques. In: Frequency Control, 40th Annual Symposium on. 1986, pp. 355–365.
61. Miller B & Clonley RJ (1991) A multiple modulator fractional divider. IEEE Trans. Instrum. Meas. 40: 578–583.
62. Riley TAD, Copeland MA & Kwasniewski TA (1993) Delta-sigma modulation in fractional-N frequency synthesis. IEEE J. Solid-State Circuits 28: 553–559.
63. Temporiti E, Albasini G, Bietti I, Castello R & Colombo M (2004) A 700-kHz bandwidth Σ-Δ fractional synthesizer with spurs compensation and linearization techniques for WCDMA applications. IEEE J. Solid-State Circuits 39(9): 1446–1454.
64. Ahola R & Halonen K (2003) A 4 GHz Δ-Σ fractional-N frequency synthesizer. Analog Integrated Circuits and Signal Processing 34(2): 77–87.
65. Rogers J, Dai F, Cavin M & Rahn D (2005) A multiband Δ-Σ fractional-N frequency synthesizer for a MIMO WLAN transceiver RFIC. IEEE J. Solid-State Circuits 40(3): 678–689.
66. Martin FL (1989). Frequency synthesizer with spur compensation. U. S. Patent 4,816,774.
67. Kozak MA & Kale I (2004) Rigorous analysis of delta-sigma modulators for fractional-N PLL frequency synthesis. IEEE Trans. Circuits Syst. I 51: 1148–1162.
68. Kozak MA & Kale I (2003) Oversampled Delta-Sigma Modulators Analysis, Applications and Novel Topologies. Kluwer Academic Publishers.
69. Kozak MA & Kale I (2001) A pipelined noise shaping coder for fractional-N frequency synthesis. IEEE Trans. Instrum. Meas. 50: 1154–1161.
70. Sun L, Lepley T, Nozahic F, Bellissant A, Kwasniewski T & Heim B (1999) Reduced complexity, high performance digital delta-sigma modulator for fractional-N frequency synthesis. In: ISCAS '99. Proceedings of the 1999 IEEE International Symposium on Circuits and Systems, pp. 152–155. Orlando, FL, USA.
71. Hosseini K & Kennedy MP (2007) Mathematical analysis of a prime modulus quantizer MASH digital delta-sigma modulator. IEEE Trans. Circuits Syst. II 54(12): 1105–1109.
72. Hosseini K & Kennedy MP (2007) Maximum sequence length MASH digital delta-sigma modulators. IEEE Trans. Circuits Syst. I 54(12): 2628–2638.
73. Aziz PM, Sorensen HV & VanderSpiegel J (1996) An overview of sigma-delta converters. IEEE Signal Processing Mag. 13(1): 61–84.
74. Bennett WR (1948) Spectra of quantized signals. Bell System Technical Journal 27: 446–471.
75. Gray RM & Neuhoff DL (1998) Quantization. IEEE Trans. Inform. Theory 44: 2325–2383.
76. Wannamaker RA, Lipshitz SP, Vanderkooy J & Wright JN (2000) A theory of nonsubtractive dither. IEEE Trans. Signal Processing 48: 499–516.
77. Widrow B, Kollar I & Liu MC (1996) Statistical theory of quantization. IEEE Trans. Instrum. Meas. 45(2): 353–361.

78. Kollar I (1986) The noise model of quantization. In: 1st ZMEKO TC4 Symp. Noise in Electrical Measurements, pp. 125–129. Budapest, OMIKK-Technoinform, Como, Italy. URL http://home.mit.bme.hu/~kollar/.
79. Widrow B & Kollár I (2008) Quantization Noise; Roundoff Error in Digital Computation, Signal Processing, Control and Communications. Cambridge University Press, New York. URL http://www.mit.bme.hu/books/quantization/index.html.
80. Farrell R & Feely O (1997) Application of non-linear analysis techniques to sigma-delta modulators. In: Signals Systems and Chaos (Ref. No. 1997/393), IEE Colloquium on, pp. 4/1–4/6.
81. Feely O (1997) A tutorial introduction to non-linear dynamics and chaos and their application to sigma-delta modulators. International Journal of Circuit Theory and Applications 25(5): 347–367.
82. Hyun D & Fischer G (2002) Limit cycles and pattern noise in single-stage single-bit delta-sigma modulators. IEEE Trans. Circuits Syst. I 49: 646–656.
83. Reefman D, Reiss J, Janssen E & Sandler M (2005) Description of limit cycles in sigma-delta modulators. IEEE Trans. Circuits Syst. I 52(6): 1211–1223.
84. Friedman V (1988) The structure of the limit cycles in sigma delta modulation. IEEE Trans. Commun. 36: 972–979.
85. Gray RM (1990) Quantization noise spectra. IEEE Trans. Inform. Theory 36: 1220–1244.
86. Gray RM (1989) Spectral analysis of quantization noise in a single-loop sigma-delta modulator with DC input. IEEE Trans. Commun. 37: 588–599.
87. Gray RM, Chou W & Wong PW (1989) Quantization noise in single-loop sigma-delta modulation with sinusoidal inputs. IEEE Trans. Commun. 37(9): 956–968.
88. Candy J & Benjamin O (1981) The structure of quantization noise from sigma-delta modulation. IEEE Trans. Commun. 29(9): 1316–1323.
89. Risbo L (1995) On the design of tone-free sigma delta modulators. IEEE Trans. Circuits Syst. II 42(1): 52–55.
90. Hein S & Zakhor A (1993) On the stability of sigma delta modulators. IEEE Trans. Signal Processing 41: 2322–2348.
91. Schreier R (1994) On the use of chaos to reduce idle-channel tones in delta-sigma modulators. IEEE Trans. Circuits Syst. I 41: 539–547.
92. Wong N & Ng TS (2003) DC stability analysis of high-order, lowpass ΣΔ modulators with distinct unit circle NTF zeros. IEEE Trans. Circuits Syst. II 50(1): 12–30.
93. Schreier R (1993) An empirical study of high-order single-bit delta-sigma modulators. IEEE Trans. Circuits Syst. II 40: 461–466.
94. Schreier R, Goodson MV & Zhang B (1997) An algorithm for computing convex positively invariant sets for delta-sigma modulators. IEEE Trans. Circuits Syst. I 44: 38–44.
95. Ouzounov S, Roza E, Hegt J, van der Weide G & van Roermund A (2006) Analysis and design of high-performance asynchronous sigma-delta modulators with a binary quantizer. IEEE J. Solid-State Circuits 41(3): 588–596.
96. Chou W & Gray RM (1991) Dithering and its effects on sigma-delta and multistage sigma-delta modulation. IEEE Trans. Inform. Theory 37: 500–513.
97. Pamarti S, Welz J & Galton I (2007) Statistics of the quantization noise in 1-bit dithered single-quantizer digital delta-sigma modulators. IEEE Trans. Circuits Syst. I 54(3): 492–503.
98. Dunn C & Sandler M (1995) Linearising sigma-delta modulators using dither and chaos. In: IEEE International Symposium on Circuits and Systems ISCAS, pp. 625–628.

99. Hein S (1993) Exploiting chaos to suppress spurious tones in general double-loop Σ-Δ modulators. IEEE Trans. Circuits Syst. II 40: 651–659.
100. Ho CYF, Ling BWK & Reiss JD (2006) Estimation of an initial condition of sigma-delta modulators via projection onto convex sets. IEEE Trans. Circuits Syst. I 53(12): 2729–2738.
101. Pamarti S & Galton I (2007) LSB dithering in MASH Δ-Σ D/A converters. IEEE Trans. Circuits Syst. I 54(4): 779–790.
102. Galton I (1993) One-bit dithering in delta-sigma modulator-based D/A conversion. In: ISCAS '93, IEEE International Symposium on Circuits and Systems, pp. 1310–1313 vol.2.
103. Candy J & Temes G (1991) Oversampling methods for data conversion. In: IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, pp. 498–502 vol.2.
104. Hosseini K & Kennedy M (2006) Mathematical analysis of digital MASH delta-sigma modulators for fractional-N frequency synthesizers. In: PRIME, Ph. D. Research in Microelectronics and Electronics 2006, pp. 309–312.
105. Hosseini K, Kennedy M & McCarthy C (2007) Calculation of sequence lengths in MASH 1-1-1 digital delta sigma modulators with a constant input. In: PRIME, Ph. D. Research in Microelectronics and Electronics 2007, pp. 13–16.
106. Ye Z & Kennedy MP (2007) Reduced complexity MASH delta sigma modulator. IEEE Trans. Circuits Syst. II 54(8): 725–729.
107. Oppenheim AV & Schaffer RW (1975) Digital Signal Processing. Prentice-Hall, Englewood Cliffs, New Jersey.
108. IEEE (1999). IEEE standard definitions of physical quantities for fundamental frequency and time metrology - random instabilities.
109. Keaveney MF & Hunt WP (2005). Variable modulus interpolator, and a variable frequency synthesizer incorporating the variable modulus interpolator. U. S. Patent 6,927,716.
110. Riley T (2007). Multiple stage delta sigma modulators. U. S. Patent Application 20070018871.
111. Perrott MH, Trott MD & Sodini CG (2002) A modelling approach for sigma-delta fractional-N frequency synthesizers allowing straightforward noise analysis. IEEE J. Solid-State Circuits 37: 1028–1038.
112. ETSI (1999). Radio transmission and reception. GSM standard. URL `http://www.etsi.org/`. 3GPP TS 05.05 version 8.9.0.
113. ETSI (1998). Radio subsystem synchronization. GSM standard. URL `http://www.etsi.org/`. GSM 05.10 version 7.3.1.
114. Analog Devices (1999). TxDAC AD9750 data sheet, rev. 0. URL `http://www.analog.com`.

Appendix 1 Spectrum representation for discrete-time periodic signals

The spectra of discrete-time periodic signals studied here represent a special case of spectral analysis, as the spectrum is not estimated but can be computed exactly. In this appendix we prove Parseval's relation for the power spectrum representation used in this work and study the relationship between the average power, mean and variance of the real-valued discrete time periodic signal. The appendix has a formal character, as it expresses some very well known relationships with the notation used in this work. It is nevertheless included on account of the specific character of the spectral analysis performed here.

Discrete power spectrum

Let us consider a discrete-time signal x with fundamental period[9] L_s:

$$x(n) = x(n + k \cdot L_s), \text{ where } k \in \mathbb{Z}. \tag{56}$$

The discrete Fourier transform (DFT) of $x(n)$ and the inverse DFT (IDFT) can be expressed as [20, 107]:

DFT

$$X(k) = \sum_{n=0}^{L_s-1} x(n) e^{-j2\pi kn/L_s} \qquad k = 0, 1, 2, \ldots, L_s - 1, \tag{57}$$

IDFT

$$x(n) = \frac{1}{L_s} \sum_{k=0}^{L_s-1} X(k) e^{j2\pi kn/L_s} \qquad n = 0, 1, 2, \ldots, L_s - 1. \tag{58}$$

The definition of the DFT presented above is implemented in the same form in Matlab® and therefore the expressions derived here can be used directly to construct algorithms for use in that environment. Formally, the spectra of the discrete-time periodic signals should be expressed by the discrete Fourier series (DFS), but in practise the DFT and DFS have the same form [20, 107]. Some authors introduce a scaling

[9] We deliberately avoid using the common symbol N for denoting a period, in order to not to confuse it with the fractional N in fractional-N frequency synthesis.

coefficient L_s^{-1} into the definition of the DFS [20], but others do not [107]. More importantly, the difference between the DFT and the DFS lies in the interpretation of the spectra.

We will represent the power spectrum of a discrete-time periodic signal x as:

$$P_x(k) = \frac{1}{L_s^2} |X(k)|^2 . \tag{59}$$

Since we will use the DFT to compute a power spectrum for one period of a periodic signal with period L_s, the spectrum will be computed exactly and will be composed of L_s discrete tones with no power between them.

Parseval's relation

The average power P_x of a discrete-time periodic signal with period L_s is defined as:

$$P_x = \frac{1}{L_s} \sum_{n=0}^{L_s-1} |x(n)|^2 . \tag{60}$$

The power spectrum for a discrete-time periodic signal (59) is defined in such a way that the average power in the signal is the sum of the powers of the individual frequency components [20, p. 251]. Parseval's relation for a discrete-time periodic signal can be stated as:

$$P_x = \frac{1}{L_s} \sum_{n=0}^{L_s-1} |x(n)|^2 = \sum_{k=0}^{L_s-1} P_x(k). \tag{61}$$

For convenience, we repeat here the standard proof of Parseval's relation for discrete-time periodic signals [20, p. 250], adapted to the notation used in this thesis:

$$\begin{aligned}
P_x &= \frac{1}{L_s} \sum_{n=0}^{L_s-1} x(n) x^*(n) \\
&= \frac{1}{L_s} \sum_{n=0}^{L_s-1} x(n) \left(\frac{1}{L_s} \sum_{k=0}^{L_s-1} X^*(k) e^{-j2\pi kn/L_s} \right) \\
&= \frac{1}{L_s^2} \sum_{k=0}^{L_s-1} X^*(k) \left(\sum_{n=0}^{L_s-1} x(n) e^{-j2\pi kn/L_s} \right) \\
&= \frac{1}{L_s^2} \sum_{k=0}^{L_s-1} X^*(k) X(k) = \frac{1}{L_s^2} \sum_{k=0}^{L_s-1} |X(k)|^2 .
\end{aligned}$$

Relationship between the average power, mean and variance

Here we derive a relationship between the average power P_x, mean $\bar{x}$ and variance σ_x^2 for real-valued discrete time periodic signal x. Taking (56) in to acount the mean value can be expressed as:

$$\bar{x} = \frac{1}{L_s} \sum_{n=0}^{L_s-1} x(n). \tag{62}$$

Variance is defined as:

$$\sigma_x^2 = \frac{1}{L_s} \sum_{n=0}^{L_s-1} (x(n) - \bar{x})^2. \tag{63}$$

Expanding the definition of the variance we obtain the following:

$$\begin{aligned} \sigma_x^2 &= \frac{1}{L_s} \sum_{n=0}^{L_s-1} (x(n) - \bar{x})^2 \\ &= \frac{1}{L_s} \sum_{n=0}^{L_s-1} x^2(n) - 2\bar{x} \left(\frac{1}{L_s} \sum_{n=0}^{L_s-1} x(n) \right) + \bar{x}^2 \\ &= \frac{1}{L_s} \sum_{n=0}^{L_s-1} x^2(n) - \bar{x}^2. \end{aligned} \tag{64}$$

Considering that for a real-valued signal x the following is true:

$$x^2(n) = |x(n)|^2, \tag{65}$$

then the average power defined by (60) is related to the mean value and the variance:

$$P_x = \sigma_x^2 + \bar{x}^2. \tag{66}$$

Relationship between the power spectrum and the variance

It automatically follows from the definition of the DFT (57) that:

$$X(0) = \sum_{n=0}^{L_s-1} x(n), \tag{67}$$

and therefore, considering (65) for a real-valued signal x, the zero component of the discrete power spectrum is:

$$P_x(0) = \frac{1}{L_s^2} |X(0)|^2 = \left| \frac{1}{L_s} \sum_{n=0}^{L_s-1} x(n) \right|^2 = |\bar{x}|^2 = \bar{x}^2. \tag{68}$$

According to Parseval's relation, (61) we can write:

$$P_x = \sum_{k=0}^{L_s-1} P_x(k) = \bar{x}^2 + \sum_{k=1}^{L_s-1} P_x(k). \tag{69}$$

Finally, taking into account (66) we can see that the variance σ_x^2 is distributed over $L_s - 1$ spectral tones, excluding the DC component at $k = 0$:

$$\sigma_x^2 = \sum_{k=1}^{L_s-1} P_x(k). \tag{70}$$

Appendix 2 The arbitrary modulus quantizer

This appendix presents a more detailed analysis of the arbitrary modulus quantizer introduced in Sec. 4.1.2. This quantizer represents a generalized version of the truncation quantizer commonly used in digital DSM implementations. Assuming a uniformly distributed quantization error e, we will derive expressions for the quantization error mean $\bar{e}$ and variance σ_e^2 and the average quantization noise power P_e. The input-output characteristic and the quantization error of the arbitrary modulus quantizer are shown in Fig. 35.

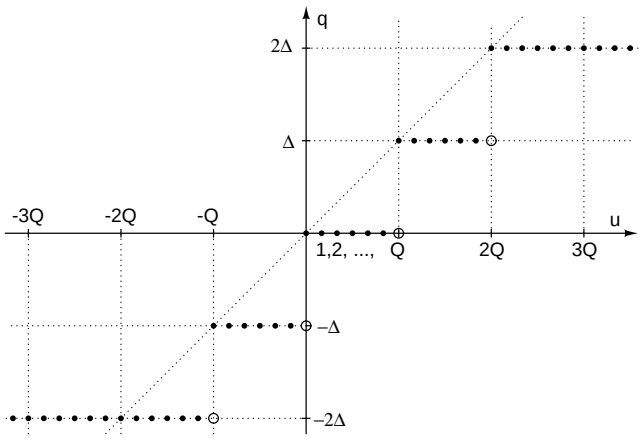

(a) Input-output characteristic.

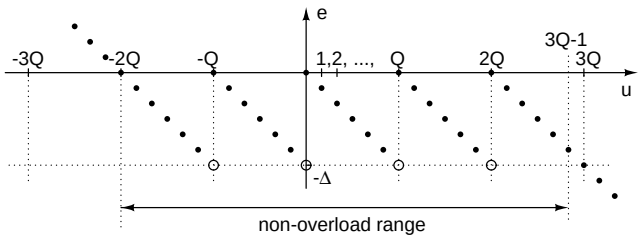

(b) Quantization error.

Fig 35. A five-step arbitrary modulus quantizer.

The operation of the quantizer in the non-overload range can be defined as:

$$\begin{cases} q &= \frac{\Delta}{Q}u + e \\ e &= -\frac{\Delta}{Q}(u \bmod Q). \end{cases} \tag{71}$$

The quantization error differs from that of the CMQ and the mid-rise and mid-tread quantizers discussed in Sec. 2.1 in three ways. First of all, it is discrete, it is asymmetrical (has a negative offset) and it is distributed in a range of magnitude smaller than Δ:

$$e = -\Delta\left(0, \frac{1}{Q}, \frac{2}{Q}, \ldots, \frac{Q-1}{Q}\right). \tag{72}$$

The distribution range approaches Δ as the modulus Q increases. The discrete and uniformly distributed quantization error in the given range is shown in Fig. 36. The mean value $\bar{e}$ of the quantization error can be expressed as:

$$\begin{aligned} \bar{e} &= -\frac{1}{Q}\sum_{i=0}^{Q-1}\Delta\cdot\frac{i}{Q} \\ &= -\frac{\Delta}{2}\cdot\frac{Q-1}{Q} \\ \lim_{Q\to\infty}\bar{e} &= -\frac{\Delta}{2}. \end{aligned} \tag{73}$$

The quantization error variance can be expressed as:

$$\begin{aligned} \sigma_e^2 &= \sum_{i=0}^{Q-1}(e(i)-\bar{e})^2\cdot p(e(i)) \\ &= \frac{1}{Q}\sum_{i=0}^{Q-1}\left(-\Delta\frac{i}{Q}+\frac{\Delta}{2}\cdot\frac{Q-1}{Q}\right)^2 \\ &= \frac{\Delta^2}{12}\cdot\frac{(Q-1)(Q+1)}{Q^2} \\ \lim_{Q\to\infty}\sigma_e^2 &= \frac{\Delta^2}{12}. \end{aligned} \tag{74}$$

As expected, as the quantizer resolution approaches infinity the quantization error variance approaches the well-known value $\Delta^2/12$ that is characteristic of a continuous system. As the bus width increases, the real value of the variance quickly approaches this limit value, and will have already reached 99 % of it with $b = 4$ bits. We can therefore assume for most practical systems that the variance for the arbitrary modulus quantizer, and an equivalent truncation quantizer, is $\Delta^2/12$.

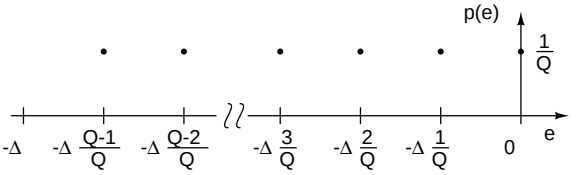

Fig 36. Probability function of the quantization error.

Due to the asymmetry of the quantizer, the quantization error power is not simply equal to the variance but is affected by the DC offset:

$$\begin{aligned} P_e &= \frac{1}{Q}\sum_{i=0}^{Q-1}\left(-\Delta\frac{i}{Q}\right)^2 \\ &= \frac{\Delta^2}{6}\cdot\frac{(Q-1)(2Q-1)}{Q^2} \\ \lim_{Q\to\infty} P_e &= \frac{\Delta^2}{3}. \end{aligned} \tag{75}$$

This makes the arbitrary modulus quantizer significantly different from a symmetrical quantizer. Due to the DC offset the power is 4 times higher than the variance. This result may matter in modulator architectures which do not attenuate the quantization noise at DC. In such cases mid-rise or mid-tread quantizers should be used. According to (66), the average quantization noise power is related to the mean and variance as follows:

$$P_e = \bar{e}^2 + \sigma_e^2 \approx \frac{\Delta^2}{4} + \frac{\Delta^2}{12} = \frac{\Delta^2}{3}. \tag{76}$$

In this work all modulator architectures have a maximum attenuation at DC, which allows us to neglect the DC component of the quantization error.

Appendix 3 Selected GSM transceiver specifications

This appendix contains a selection of GSM specifications that are relevant to a DSM fractional-N frequency synthesizer working as a local oscillator. The DSM generates pseudo-random sequences with two essential properties: a mean value corresponding to the selected radio channel and a noise-shaping property. The resolution required for the mean can be derived from the transceiver channel spacing specification, frequency source accuracy and reference frequency, while the DSM scaling related to the noise shaping property can be derived from the minimum measurement bandwidth used to qualify the transceiver.

GSM specifications

The required specifications can be found in two documents: ETSI 3GPP TS 05.05, which defines the requirements for the transceiver, see [112], and ETSI GSM 05.10, which defines the requirements for synchronization on the GSM radio sub-system [113]. The parameters in question fall into the following categories:

Frequency bands and channel arrangement. The technical specification [112] defines all the available GSMS bands. The channel spacing Δ_{ch} for all bands is:

$$\Delta_{\text{ch}} = 200\,\text{kHz}. \tag{77}$$

Radio frequency tolerance. The technical specification [113] defines the requirements for the frequency source separately for a base station (BS) and a mobile station (MS). In this work we will use the BS requirement, as it is the stricter of the two and assumes a relative accuracy in the frequency source:

$$\delta_f = 0.05\,\text{ppm}. \tag{78}$$

Transmitter characteristics. The technical specification [112] defines the output RF spectrum and permitted spurious emissions in the form of spectral masks obtained by measuring the power in a given bandwidth at various frequencies. The bandwidth is usually increased as the frequency offset between the measurement

frequency and either the carrier or the edge of the MS or BS transmit band increases. The minimum bandwidth for the spectral mask due to modulation, wide band noise, and switching transients, is 30 kHz for the BS. An example mask is shown in Fig. 37. The minimum measurement bandwidth for determining the spurious emissions in the MS is 10 kHz. For the purpose of this work we will use a commonly used specification for the BS measurement bandwidth [28, 39, 43] and denote:

$$\mathrm{BW}_{\min} = 30\,\mathrm{kHz}. \tag{79}$$

Absolute frequency accuracy

The radio frequency tolerance of the GSM system is given as a relative accuracy (78). This means that, as the generated frequency increases, the permitted absolute frequency error increases as well. Thus the strictest requirement for the accuracy of the synthesizer arises at the lowest frequency it is supposed to generate. Here we derive the maximum permitted absolute frequency error for the lowest GSM band, GSM 450. For a complete specification of the GSM bands, refer to [112]. The maximum permitted frequency errors for GSM 450, GSM 900 and GSM 1800 are shown in Table 6.

Since the lowest frequency in the GSM 450 system is $f_{\min} = 450.4\,\mathrm{MHz}$, the absolute frequency accuracy for this band can be calculated as:

$$\begin{aligned} \Delta f &= \delta_f \cdot f_{\min} \\ &= 5 \cdot 10^{-8} \cdot 450.4\,\mathrm{MHz} = 22.5\,\mathrm{Hz}. \end{aligned} \tag{80}$$

Similarly $\Delta f = 43.8\,\mathrm{Hz}$ and $\Delta f = 85.5\,\mathrm{Hz}$ in the GSM 900 and GSM 1800 bands, respectively.

Table 6. Maximum permitted absolute frequency errors Δf for various GSM bands.

Band	Δf
GSM 450	22.5 Hz
GSM 900	43.8 Hz
GSM 1800	85.5 Hz

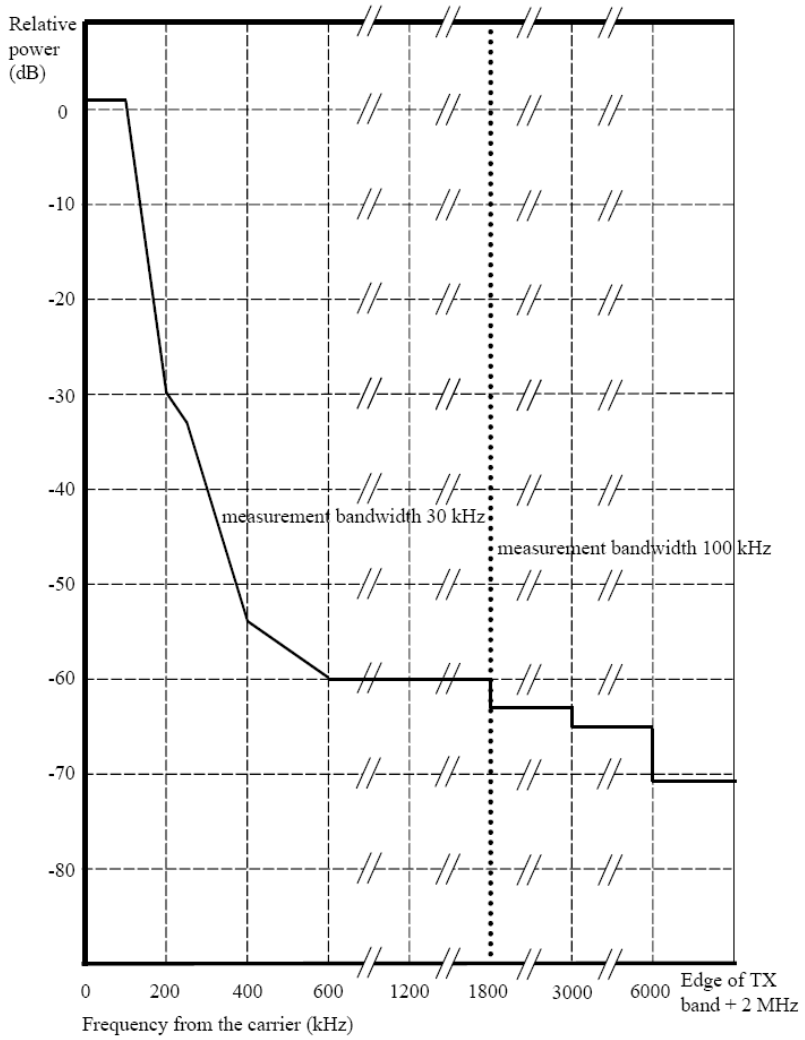

Figure A.1b: GSM 400, GSM 900 and GSM 850 MS spectrum due to 8-PSK modulation

Fig 37. Example spectral mask. Figure reprinted from [112]. © European Telecommunications Standards Institute 2008. Further use, modification, redistribution is strictly prohibited. ETSI standards are available from http://pda.etsi.org/pda/

Appendix 4 Measurement setup

A simple test system has been built to enable DSM output spectra to be measured directly using a spectrum analyzer. The test system comprises a Xilinx® ML401 evaluation board and a custom-built DAC extension board. The block diagram of the test system is shown in Fig. 38. The DAC board is based on a 10-bit, 125 MSPS, DAC AD9750. The differential DAC current output is converted to a single-ended voltage output by means of an operational amplifier AD8056, according to the standard configuration suggested in the DAC data sheet [114]. The spurious-free dynamic range of the DAC can vary between 70 and 87 dBc depending on the clock frequency and the differential to single-ended coupling method [114]. In the selected arrangement, based on the operational amplifier and a 20 MHz clock frequency, the noise floor remained at 70 dBc, as shown in Fig. 20.

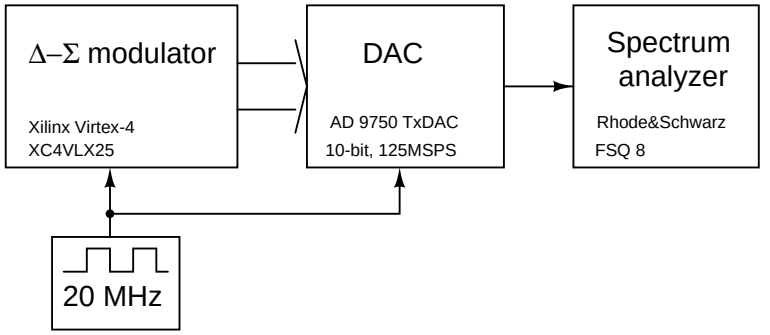

Fig 38. Block diagram of the measurement system.

This test system is not equipped with any image rejection filter and allows direct measurement of the DSM output spectrum over the full frequency range. This measurement system can be used to measure the out-of band spectrum, as the in-band noise is covered by the noise floor of the off-the-shelf multi-bit DAC. When comparing the measured spectra with simulation results based on the DFT, the following parameters have to be taken to account:

Sample and hold. Since the DAC performs first-order sample and hold operations, the spectrum is multiplied by the frequency response of the sinc function (see

((25) in Sec. 3.2.3).

DAC gain. The quantized output from the DSM passes though the multi-bit DAC, in which the gain is selected to maximize SNR but not to overdrive the spectrum analyzer input. With the actual selected DAC gain taken into account, the value of the DSM quantization step is $\Delta = 75\,\mathrm{mV}$.

Load impedance. The spectrum analyzer provides a standard load impedance of $R_L = 50\,\Omega$.

Clock frequency. The DSM and DAC are driven at the same clock frequency $f_s = 20\,\mathrm{MHz}$. In the context of the DSM implementation the clock frequency is referred to as the sampling frequency f_s.

Original articles

I Häkkinen J, Borkowski MJ & Kostamovaara J (2003) A PLL-based RF synthesizer test system. In IFIP VLSI-SOC 2003, International Conference on Very Large Scale Integration. Darmstadt, Germany, December 2003: 211–214.

II Borkowski MJ, Häkkinen J & Kostamovaara J (2003) A sigma-delta modulator development environment for fractional-N frequency synthesis. In IFIP VLSI-SOC 2003, International Conference on Very Large Scale Integration. Darmstadt, Germany, December 2003:50–54.

III Borkowski MJ & Kostamovaara J (2004) Post modulator filtering in Δ-Σ fractional-N frequency synthesis. In MWSCAS-04, The 2004 IEEE International Midwest Symposium on Circuits and Systems, volume I. Hiroshima, Japan, July 2004: 325–328.

IV Borkowski MJ & Kostamovaara J (2005) Spurious tone free digital delta-sigma modulator design for DC inputs. In ISCAS 2005, IEEE International Symposium on Circuits and Systems. Kobe, Japan, May 2005: 5601–5604.

V Borkowski MJ & Kostamovaara J (2006) On randomization of digital delta-sigma modulators with DC inputs. In ISCAS 2006, IEEE International Symposium on Circuits and Systems. Kos, Greece, May 2006: 3770–3773.

VI Borkowski MJ, Riley TAD, Häkkinen J & Kostamovaara J (2005) A practical Δ-Σ modulator design method based on periodical behavior analysis. IEEE Trans Circuits Syst II 52: 626–630.

VII Borkowski MJ & Kostamovaara J (2007) Variable modulus delta-sigma modulation in fractional-N frequency synthesis. Electronics Letters 25: 1399–1400.

Reprinted with permission from Department of Electrical Engineering and Information Technology, Darmstadt University of Technology [I], [II]; IEEE, The Institute of Electrical and Electronics Engineers, Incorporated [III], [IV], [V], [VI]; IET, The Institution of Engineering and Technology [VII].

Original publications are not included in the electronic version of the dissertation.

289. Nissilä, Mauri (2008) Iterative receivers for digital communications via variational inference and estimation

290. Toivonen, Tuukka (2007) Efficient methods for video coding and processing

291. Lyöri, Veijo (2007) Structural monitoring with fibre-optic sensors using the pulsed time-of-flight method and other measurement techniques

292. Stoica, Lucian (2008) Non-coherent energy detection transceivers for Ultra Wideband Impulse radio systems

293. Koski, Anna (2008) Applicability of crude tall oil for wood protection

294. Gore, Amol (2008) Exploring the competitive advantage through ERP systems. From implementation to applications in agile networks

295. Kirillin, Mikhail (2008) Optical coherence tomography of strongly scattering media

296. Tölli, Antti (2008) Resource management in cooperative MIMO-OFDM cellular systems

297. Karkkila, Harri (2008) Consumer pre-purchase decision taxonomy

298. Rabbachin, Alberto (2008) Low complexity UWB receivers with ranging capabilities

299. Kunnari, Esa (2008) Multirate MC-CDMA. Performance analysis in stochastically modeled correlated fading channels, with an application to OFDM-UWB

300. Särkkä, Jussi (2008) A novel method for hazard rate estimates of the second level interconnections in infrastructure electronics

301. Mäkelä, Juha-Pekka (2008) Effects of handoff algorithms on the performance of multimedia wireless networks

302. Teräs, Jukka (2008) Regional science-based clusters. A case study of three European concentrations

303. Lahti, Markku (2008) Gravure offset printing for fabrication of electronic devices and integrated components in LTCC modules

304. Popov, Alexey (2008) TiO2 nanoparticels as UV protectors in skin

305. Illikainen, Mirja (2008) Mechanisms of thermomechanical pulp refining

Book orders:
OULU UNIVERSITY PRESS
P.O. Box 8200, FI-90014
University of Oulu, Finland

Distributed by
OULU UNIVERSITY LIBRARY
P.O. Box 7500, FI-90014
University of Oulu, Finland

UNIVERSITY OF OULU P.O.B. 7500 FI-90014 UNIVERSITY OF OULU FINLAND

ACTA UNIVERSITATIS OULUENSIS

SERIES EDITORS

A **SCIENTIAE RERUM NATURALIUM**

Professor Mikko Siponen

B **HUMANIORA**

University Lecturer Elise Kärkkäinen

C **TECHNICA**

Professor Hannu Heusala

D **MEDICA**

Professor Olli Vuolteenaho

E **SCIENTIAE RERUM SOCIALIUM**

Senior Researcher Eila Estola

F **SCRIPTA ACADEMICA**

Information officer Tiina Pistokoski

G **OECONOMICA**

University Lecturer Seppo Eriksson

EDITOR IN CHIEF

Professor Olli Vuolteenaho

PUBLICATIONS EDITOR

Publications Editor Kirsti Nurkkala

ISBN 978-951-42-8909-5 (Paperback)
ISBN 978-951-42-8910-5 (PDF)
ISSN 0355-3213 (Print)
ISSN 1796-2226 (Online)

UNIVERSITY of OULU
OULUN YLIOPISTO