
On Time, Time Synchronization and Noise in Time Measurement Systems

A dissertation submitted towards the degree Doctor of
Natural Sciences of the Faculty of Mathematics and
Computer Sciences of Saarland University

by Attila Kinali-Dogan

Saarbrücken / 2022

Day of Colloquium: 12 October 2022
Dean of the Faculty: Prof. Jürgen Steimle

Chair: Prof. Karl Bringmann
Academic Assistant: Dr. Roohani Sharma
Reporters:
Dr. Christoph Lenzen
Prof. Kurt Mehlhorn
Prof. Matthias Függer
Prof. François Vernotte

Abstract

Time plays an important role in our modern lives. Especially having accurate time, which in turn depends on having clocks being synchronized to each other.

This thesis is split into three distinct parts.

The first part deals with the mathematical description of noise that is required to model clocks and electronics accurately. In particular we will address the problem that the generally used tools from signal theory fail for noise signals which are neither of finite energy nor periodic in nature. For this we will introduce a new function space based on the P^p -seminorm that is an extension of the L^p -norm for functions of potentially infinite energy but limited power. Using this new semi-norm we will modify the Fourier transform to work on signals from this P^p -space. And last but not least, we will introduce, based on the above, a new mathematical model of noise that captures all the properties associated with $1/f^\alpha$ -noise.

In the second part, we will look at how noise propagates in a few classes of electronics, especially how the non-linear behavior of electronics leads to an amplification of noise and how it could be mitigated.

Lastly, in the third part we will look at one approach of fault-tolerant clock synchronization. After explaining its working principle and showing an implementation in an FPGA we will focus on metastability, the problems it can cause and how to handle them on two different circuit levels.

Zusammenfassung

Zeit spielt eine wichtige Rolle in unserem Leben. Insbesondere die Verfügbarkeit einer genauen Zeit. Welches wiederum davon abhängt, dass man Uhren hat die auf einander synchronisiert laufen.

Diese Arbeit ist in drei Teile aufgeteilt:

Im ersten Teil betrachten wir die mathematische Beschreibung von Rauschen um elektronische Systeme und Uhren korrekt beschreiben zu können. Im Besonderen betrachten wir die Probleme die die generell benutzten Methoden der Signalverarbeitung beim Umgang mit Rauschsignalen haben, die weder energiebegrenzt noch periodisch sind. Dafür erweitern wir den Funktionenraum der L^p -Norm auf leistunglimitierte Funktionen und führen die P^p -Halbnorm ein und modifizieren die Fouriertransformation zur Verwendung auf diesen Raum. Und letztlich führen wir ein neues mathematisches Model zur Beschreibung von Rauschen ein, welches alle üblicherweise angenommenen Eigenschaften gleichzeitig erfüllt.

Im zweiten Teil analysieren wir wie sich einige Klassen von elektronischen Schaltungen im Bezug auf Rauschen verhalten. Insbesondere im Bezug auf das nicht-lineare Verhalten der elektronischen Elemente, welches zu einer Verstärkung des Rauschens führt.

Im dritten Teil betrachten wir eine Möglichkeit um fehlertolerante Synchronization von Uhren zu erreichen. Nach einem Überblick über den verwendeten Algorithmus und wie dieser einem FPGA implementiert werden kann, schauen wir uns den Einfluss von Metastabilität an und wie dieser eingedämmt werden kann.

Acknowledgments

It is kind of tradition to start the acknowledgments of a thesis with the advisor. I am not the one to break with this tradition, but I would like to emphasize that *Christoph Lenzen* went in many ways beyond what would be expected from an advisor and often applied non-traditional methods whenever required to make problems go away. Even how I got hired as a PhD student was not quite the usual way, the whole story of which might be shared over some drinks, and involved the gentle prodding of *Mimi*. Christoph not only gave me the opportunity to pursue a PhD, but also allowed me to find my own path through the lands of science, even as I strayed more and more from his own research into territories uncharted and full of dragons. Some of these dragons have been successfully tamed and can be found in this thesis.

Of the people I have met at MPI *Stephan Friedrichs* and *Matthias Függer* were by far the most influential. Not only did we share many hours brooding over various scientific problems, we also became good friends over many shared hot and cold beverages. And even after both left MPI, we still stayed in contact and engaged in scientific and non-scientific discussions whenever the opportunity arose.

Others, like (in no particular order) *Daniel Vaz*, *Bhaskar Ray Chaudhury*, *Philip Wellnitz*, *Pavel Kolev*, and many more made the often long evenings enjoyable, even when deadlines were looming. And not to forget *Christina Fries*, our secretary at D1, who made many impossible things possible.

Given that I was an electrical engineer doing electrical engineering things at a theoretical computer science department I often had questions nobody in our group could answer. Therefore I had to go and look elsewhere for answers. *Magnus Danielson*, *Enrico Rubiola*, and *Claudio Calosso* were among those who had to answer most of my questions and did so quite admirably and with lots of patience. I am very thankful for their help, without which I would not have been able to gain as much understanding in many topics, more than could ever fit into a thesis like this. Especially considering that my questions were quite often not the easy kind and involved a lot of discussion. Another name in this illustrious group is *François Vernotte*, whose SigmaTheta program I (partially) hijacked for some time, to generate the noise for my simulations. It was also François who invited Magnus and me to a short workshop on noise in atomic clocks back in 2019, that ultimately led to the mathematical analysis in first part of this thesis.

Another large group that should not be forgotten are the members of the time-nuts mailinglist lead by *Tom Van Baak* and administrated by *John Ackermann*. The time-nuts mailinglist is a loose group of people who are interested, in some way or other, in precision time and frequency measurement, many of which collect atomic clocks as a hobby. Many of these people are experts in very specific parts of this interesting field with knowledge and experience that is hard to come by anywhere else. Since the time I joined this mailinglist way back in February 2008, as a total novice, I learned not only the craft of precision measurements, but also a lot about low-noise electronics, atomic clocks and other frequency sources, time and frequency distribution, and everything else related to time in some form or other. I asked many, quite often naive questions over the years that people answered with lots of patience, often adding historical details that would be impossible to find elsewhere. The members of the time-nuts community were the people who instilled in me the passion for time and frequency and are the source of much of my knowledge therein. I would like to thank all members of the time-nuts community for all their help and kindness over all these years. I would like to highlight (again, in no particular order) *Bruce Griffiths*, *Charles Steinmetz*, *Rick Karlquist*, and *Bob Camp* who answered many of my most difficult questions and gave deep insights into the how and why with their incredible expertise.

Contents

Contents	ix
Symbols	xi
1 Notation and Preliminaries	1
1.1 Analysis and Measure Theory	1
1.2 Hyperreal Numbers and Nonstandard Analysis	8
1.3 Fourier Analysis	10
1.4 Probability Theory	11
1.5 Fractional Calculus	14
1.6 Special Functions	14
I Noise and its Mathematical Description	17
2 White and $1/f^\alpha$-Noise in Nature and Technology	19
3 The P^p-seminorm and its Integral Transforms	21
3.1 The Problem with Power-Limited Functions	21
3.2 The P^p -seminorm and its Associated Function Space	22
3.3 P^p -seminorm Integral Transforms	26
4 White Noise	37
4.1 Mathematical Description	38
5 $1/f^\alpha$-Noise	41
5.1 Mathematical Description	42
II Noise Propagation in Electronics	49
6 The Importance of Time in Modern Metrology	51
6.1 The Current State of Art in Low-Noise Amplifier Design	52
7 Noise Propagation in Electronics	55

7.1	Circuit Model	55
7.2	Noise Sources	57
7.3	Noise Translation and Scaling	59
7.4	Conclusion	65
8	Noise Propagation in Hard Limiters	67
8.1	Circuit Model	67
8.2	Circuit Analysis	68
8.3	Conclusion	71
III	Faults and Fault-tolerance in Clocking Applications	73
9	Clock Synchronization and Reliability	75
10	Fault-tolerant Clock Synchronization	77
10.1	Introduction	77
10.2	Using Fault-Tolerant Algorithms for Clock Synchronization	78
10.3	Implementation and Experiments	81
11	Metastability in Time to Digital Converters	89
11.1	Introduction	89
11.2	Binary Cycle Counter Designs	94
11.3	Metastability-Aware time to digital converters	96
11.4	Conclusion	104
12	Efficient Metastability-Containing Multiplexers	107
12.1	Introduction	107
12.2	Problem Statement	111
12.3	Circuits	112
12.4	Simulation	114
12.5	Impact on Metastability-Containing Circuits	117
12.6	Conclusion	118
IV	Backmatter	119
	Bibliography	129

Symbols

\mathbb{N}	Natural numbers starting from 1
\mathbb{N}_0	Natural numbers starting from 0
\mathbb{Z}	Integer numbers
\mathbb{R}	Real numbers
$\mathbb{R}_{\geq 0}$	Non-negative real numbers
$\mathbb{R}_{> 0}$	Strictly positive real numbers
${}^*\mathbb{R}$	Hyperreal numbers
\mathbb{C}	Complex numbers
\emptyset	Empty set
${}^*r, {}^*f()$	Hyperreal number / function
e	Euler's number $e \approx 2.718281828$
j	Imaginary unit $j^2 = -1$
γ	Euler-Ascheroni constant $\gamma \approx 0.5772156649$
$\mathbb{1}_A$	Indicator function on a set A
C^0	Class of continuous functions
C^∞	Class of infinitely differentiable / smooth functions
\bar{x}	Complex conjugate of x
$\{a_k\}_{k \in I}$	Sequence over an index set I
$ x $	Lebesgue measure or absolute value of x
$\lceil x \rceil$	Ceil of x
$\lfloor x \rfloor$	Floor of x
$\ x\ _{L^p}$	L^p -norm of x
$\ x\ _{P^p}$	P^p -seminorm of x
$x^{\bar{n}}$	Rising factorial (Pochhammer's function) $x^{\bar{n}} = \prod_{k=0}^{n-1} (x - k)$
$B(r, x)$	Ball with radius r around x
$B(a, b; x)$	Incomplete beta function
$F(a, b; c; x)$	Gauss hypergeometric function
$\text{supp}(f)$	Support of a function f
$\mathcal{F}\{f\}$	Fourier transform of f
$\mathcal{F}^{-1}\{f\}$	Inverse Fourier transform of f
$ \mathcal{F}\{f\} ^2$	energy spectral density of f
$\mathcal{F}_{B(r,y)}\{f\}, \mathcal{F}_r\{f\}, \mathcal{F}_\infty\{f\}$	P^p -Fourier transform of f
$\mathcal{F}_{B(r,y)}^{-1}\{f\}, \mathcal{F}_r^{-1}\{f\}, \mathcal{F}_\infty^{-1}\{f\}$	Inverse P^p -Fourier transform of f
$ \mathcal{F}_\infty\{f\} ^2$	power spectral density of f

ν_0	Signal frequency
ξ	Fourier / offset frequency
*	Convolution $(f * g)(x) = \int_{\mathbb{R}^n} f(x - y)g(y) dy$
\approx	Approximately equal, i.e., equal when ignoring higher order terms
\lesssim	Approximately smaller, i.e., smaller while ignoring higher order terms
\sim	Similar, i.e., if $X \sim Y$ then X and Y have the same probability distribution
$\mathcal{N}(\mu, \sigma^2)$	Normal distribution with mean μ and variance σ^2
$\mathcal{U}(A)$	Uniform distribution over a set A
$\mathbb{E}[X]$	Expectation of a random variable X
$\mathbb{V}[X]$	Variance of a random variable X
$\text{erf}(z)$	Error function $\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$
$\Gamma(z)$	Gamma function $\Gamma(z) = \int_0^\infty x^{z-1} e^{-x} dx$
${}_a^{\tau} I_t^p f(\tau)$	Fractional integral of $f(\tau)$
$\text{sinc}(x)$	Sinc function: $\text{sinc}(x) = \sin(x)/x$
$\Pi(x)$	Rectangular unit pulse of height 1 and width 1, centered around 0

CHAPTER 1

Notation and Preliminaries

This chapter contains the basic definitions of some of the symbols used and introduces some of the used theorems. These are either all common definitions. Most of the definitions and theorems in this chapter are taken verbatim from textbooks, without change, with the exception of using a consistent notation and some slight simplification of language, here and there. As it is not feasible to reproduce all the definitions that are required to cover the topics at hand, we refer to the referenced books for further details and explanations. As proving the theorems mentioned here goes beyond the scope of this dissertation, we refer to the sources these theorems are taken from, for their respective proofs.

This chapter starts with more basic definitions and theorems than it would be strictly necessary for the following discussion. The aim here is to provide a proper basis in order to avoid confusion or inaccuracies due to simplifications often found in engineering texts.

1.1 Analysis and Measure Theory

In this thesis, unless noted otherwise, the triple (A, \mathcal{A}, μ) denotes a space A with a sigma algebra \mathcal{A} and a σ -finite, non-negative measure μ on \mathcal{A} . We call this triple a *measure space*. Due to space limitations we will only focus on the key building blocks of σ -algebra and measure theory. For a proper and complete introduction we would like to refer to [Bog07] and [Hal50].

On \mathbb{R}^n and its subspaces, we will use the Lebesgue measure if not noted otherwise.

Integration and Convergence

Definition 1.1 (Metric Space). [Sim15, p. 3] A metric d on a space E is a map $d : E \times E \rightarrow \mathbb{R}_{\geq 0}$ that obeys

$$(1) \quad \forall x, y \in E, \quad d(x, y) = d(y, x)$$

This chapter does not constitute original work and the author does not lay any claim to anything discussed here.

$$(2) \quad \forall x, y, z \in E, \quad d(x, z) \leq d(x, y) + d(y, z)$$

$$(3) \quad \forall x, y \in E, \quad d(x, y) = 0 \Leftrightarrow x = y$$

A metric space is a set E , equipped with a metric.

Definition 1.2 (Cauchy Sequence). [Sim15, p. 5] Given a metric space E , a sequence $\{x_k\}_{k \in \mathbb{N}}$ of points $x_k \in E$ is called Cauchy if and only if

$$\forall \varepsilon, \exists N : n, m > N \Rightarrow d(x_n, x_m) < \varepsilon \quad (1.1)$$

Definition 1.3 (Complete Metric Space). [Sim15, p. 5] A metric space E is called complete if and only if every Cauchy sequence in E converges to some $x_\infty \in E$.

Definition 1.4 (Banach Space). [Sim15, p. 113] A norm on a vector space E is a function $\|\cdot\| : E \rightarrow \mathbb{R}_{\geq 0}$ with $x, y \in E$ and $\lambda \in \mathbb{R}$ that obeys

$$(1) \quad \|x\| = 0 \Leftrightarrow x = 0$$

$$(2) \quad \|\lambda x\| = |\lambda| \|x\|$$

$$(3) \quad \|x + y\| \leq \|x\| + \|y\|$$

If Item (1) is dropped, $\|\cdot\|$ is called a seminorm.

A space E equipped with a norm is called a normed linear space. In any normed linear space, the function $d(x, y) = \|x - y\|$ defines a metric.

A Banach space is a normed linear space which is complete in the metric induced by the norm.

Theorem 1.5 (Limit Rules). [Rud76, Theorem 4.4., p. 85] Suppose E is a subset of a metric space, p is a limit point of E and f, g are functions on E with values in a Banach space, Given

$$\begin{aligned} F &:= \lim_{x \rightarrow p} f(x) \\ G &:= \lim_{x \rightarrow p} g(x) \end{aligned} \quad (1.2)$$

Then

$$\begin{aligned} \lim_{x \rightarrow p} (f + g)(x) &= F + G \\ \lim_{x \rightarrow p} (fg)(x) &= FG \\ \lim_{x \rightarrow p} \left(\frac{f}{g} \right)(x) &= \frac{F}{G} \quad G \neq 0 \end{aligned} \quad (1.3)$$

Definition 1.6 (Algebra). [Bog07, Vol. 1, Definition 1.2.1., p. 3] An algebra of sets \mathcal{A} is a class of subsets of some fixed set A (called the space) such that

- (1) A and the empty set \emptyset belong to \mathcal{A}
 (2) if $E, F \in \mathcal{A}$ then $E \cap F \in \mathcal{A}$, $E \cup F \in \mathcal{A}$, $E \setminus F \in \mathcal{A}$

Definition 1.7 (σ -Algebra). [Bog07, Vol. 1, Definition 1.2.2., p. 4]
 An algebra of sets \mathcal{A} is called a σ -algebra if for any sequence of sets $\{A_k\}_{k \in \mathbb{N}}$, $A_k \in \mathcal{A}$ one has $\bigcup_{k=1}^{\infty} A_k \in \mathcal{A}$

Definition 1.8 (σ -Finite). [Bog07, Vol. 1, Definition 1.6.1., p. 24] Let \mathcal{A} be a σ -algebra in a space A and let μ be a set function on \mathcal{A} with values in $[0, +\infty]$. If μ satisfies the condition $\mu(\emptyset) = 0$ and is countably additive in the sense that $\mu\left(\bigcup_{k=1}^{\infty} A_k\right) = \sum_{k=1}^{\infty} \mu(A_k)$ for all pairwise disjoint sets $A_k \in \mathcal{A}$, where infinite values are admissible as well, then μ is called a measure with values in $[0, +\infty]$. We call μ a σ -finite measure if $A = \bigcup_{k=1}^{\infty} A_k$ where $A_k \in \mathcal{A}$, $\mu(A_k) < \infty$.

Definition 1.9 (Lebesgue Measure). On \mathbb{R}^n we (informally) define the Lebesgue measure $|A|$ of a subset $A \subseteq \mathbb{R}^n$ as the volume of A .

See [Bog07, Vol. 1, Definition 1.7.11., p. 26] and [Hal50, §15, p. 62] for a rigorous definition.

Definition 1.10 (Measurable Function). [Bog07, Vol. 1, Definition 2.1.1., p. 105] A function $f : A \rightarrow \mathbb{R}$ is called measurable with respect to \mathcal{A} (or \mathcal{A} -measurable) if $\{x : f(x) < c\} \in \mathcal{A}$ for every $c \in \mathbb{R}$.

Definition 1.11 (μ -Measurable Function). [Bog07, Vol. 1, Definition 2.1.10., p. 108] A function $f : A \rightarrow \mathbb{R}$ is called μ -measurable if it is measurable with respect to the σ -algebra \mathcal{A}_μ of all μ -measurable sets¹ of A .

Definition 1.12 (Simple Functions). [Bog07, Vol. 1, §2.1., p. 106] A simple function is a \mathcal{A} -measurable function $f : A \rightarrow \mathbb{R}$ that assumes only finitely many values, i.e., has the form

$$f(x) = \sum_{k=1}^n c_k \mathbb{1}_{A_k}(x)$$

with $c_k \in \mathbb{R}$ and $A_k \in \mathcal{A}$.

Definition 1.13 (Integral of Simple Functions). [Bog07, Vol. 1, §2.3., p. 115] For any simple function $f : A \rightarrow \mathbb{R}$ that assumes finitely many values $c_k \in \mathbb{R}$ on disjoint sets $A_k \subseteq A$, $k \in \mathbb{N}$, the Lebesgue integral of f with respect to μ is defined by the equality

$$\int_A f(x) \mu(dx) := \sum_{k=1}^n c_k \mu(A_k).$$

1. Informally, a μ -measurable set A_k is a set for which the measure $\mu(A_k)$ is defined. Or equivalently, if A_k is an element of \mathcal{A} . See [Bog07, Vol. 1, Definition 1.5.1., p. 17] for a rigorous definition.

Definition 1.14 (Fundamental in Mean). [Bog07, Vol. 1, Definition 2.3.1., p. 116] A sequence $\{f_k\}$ of simple functions is called fundamental in the mean or mean fundamental if for every $\varepsilon > 0$ there exists a number k such that

$$\int_A |f_i(x) - f_j(x)| \mu(dx) < \varepsilon \quad \text{for all } i, j \geq k.$$

Definition 1.15 (Lebesgue Integral). [Bog07, Vol. 1, Definition 2.4.1., p. 118] Let a function f be defined and finite μ -almost everywhere (a.e.) (i.e., f may be undefined or infinite on a set of measure zero). The function f is called Lebesgue integrable with respect to the measure μ (or μ -integrable) if there exists a sequence of simple functions $\{f_k\}$ such that $f_k(x) \rightarrow f(x)$ a.e. and the sequence $\{f_k\}$ is fundamental in the mean. The finite value

$$\lim_{k \rightarrow \infty} \int_A f_k(x) \mu(dx),$$

is called the Lebesgue integral of the function f and is denoted by

$$\int_A f(x) \mu(dx) \quad \text{or by} \quad \int_A f \, d\mu.$$

We note here that it is possible to extend the Lebesgue integral, using the same construction, known as the Bochner integral [Boc33], to functions $f : A \rightarrow B$, with B being a Banach space (see also [Mik78, §3, p. 15]). We will use the Lebesgue integral for functions with their codomain in \mathbb{R} and the Bochner integral for all others.

Definition 1.16 (Locally Integrable Functions). [DK10, Definition 20.37., p. 343] A function f is said to be locally integrable if the product $(f \mathbb{1}_{A_k})$ is integrable for all $A_k \in \mathcal{A}$ whose closure is compact.

Theorem 1.17 (Translation Invariance). [Mik78, Theorem 3.3., p. 7] Given a measure space (A, \mathcal{A}, μ) , with A closed under addition, a measure μ such that $\mu(B) = \mu(B + a)$, where $B + a := \{x + a : x \in B\}$ and $B + a \in \mathcal{A}$ for all $B \in \mathcal{A}$ and $a \in A$. If a function f is Lebesgue integrable and $f_a(x) = f(x - a)$ then f_a is also Lebesgue integrable and

$$\int_A f_a(x) \mu(dx) = \int_A f(x) \mu(dx) \quad (1.4)$$

In other words, the integral is translation invariant.

Theorem 1.18 (Absolute Continuity of the Lebesgue Integral). [Bog07, Vol. 1, Theorem 2.5.7., p. 124] Let f be a μ -integrable function. Then for every $\varepsilon > 0$ there exists a $\delta > 0$ such that for every $D \in \mathcal{A}$, $\mu(D) < \delta$

$$\int_D |f(x)| \mu(dx) < \varepsilon \quad (1.5)$$

Proposition 1.19. [Bog07, Vol. 1, Proposition 2.6.2., p. 125] and [Bog07, Vol. 1, Remark 2.6.3., p. 127]

- (1) If a function f is integrable with respect to a countably additive measure μ (c.f. Definition 1.10) with values in $[0, +\infty]$, then the measure μ is σ -finite on the set $\{x : f(x) \neq 0\}$.
- (2) Let μ be a σ -finite measure on a space A that is the union of an increasing sequence of μ -measurable subsets A_k of finite measures. Then the function f is integrable with respect to μ precisely when the restrictions of f on to the sets A_k are integrable and

$$\sup_k \int_{A_k} |f| \mu(dx) < \infty \quad (1.6)$$

In this case one has, with $A_0 := \emptyset$

$$\int_A f \mu(dx) = \lim_{k \rightarrow \infty} \int_{A_k} f \mu(dx) = \sum_{k=1}^{\infty} \int_{A_k \setminus A_{k-1}} f \mu(dx) \quad (1.7)$$

Theorem 1.20 (Egoroff's Theorem). [Bog07, Vol. 1, Theorem 2.2.1., p. 110] A sequence of μ -measurable functions $\{f_k\}$, such that μ -almost everywhere there is a finite limit $f(x) := \lim_{k \rightarrow \infty} f_k(x)$. Then for every $\varepsilon > 0$ there exists a set $A_\varepsilon \in \mathcal{A}$ such that $\mu(A \setminus A_\varepsilon) < \varepsilon$ and functions $\{f_k\}$ converge to f uniformly on A_ε .

Theorem 1.21 (Dominated Convergence). [Bog07, Vol. 1, Theorem 2.8.1., p. 130] and [Bog07, Vol. 1, Corollary 2.8.6., p. 132] Suppose that a sequence of μ -integrable functions $\{f_k\}$ converge almost everywhere to a function f . If there exists a μ -integrable function Φ such that

$$|f_k(x)| \leq \Phi(x) \quad \text{a.e. for every } k,$$

then the function f is integrable and

$$\int_A f(x) \mu(dx) = \lim_{k \rightarrow \infty} \int_A f_k(x) \mu(dx) \quad (1.8)$$

Dominated convergence also holds for Bochner integrals [Mik78, Theorem 9.6, p. 35].

Theorem 1.22 (Tannery's Theorem). [Hof02, p. 199] If $s(l) = \sum_{k=1}^{\infty} f_k(l)$, $l \in \mathbb{N}$ is a finite sum (or a convergent series) for each l , $g_k := \lim_{l \rightarrow \infty} f_k(l)$, $|f_k(l)| \leq M_k$, and $\sum_{k=1}^{\infty} M_k < \infty$ then

$$\lim_{l \rightarrow \infty} s(l) = \sum_{k=1}^{\infty} g_k \quad (1.9)$$

Theorem 1.23 (Moore-Osgood Theorem). [Rud76, Theorem 7.11., p. 149] Suppose a sequence of functions $\{f_k\}$ on a set A in a metric space, converging uniformly to a function f . Let x be a limit point of A and suppose that

$$F_k := \lim_{t \rightarrow x} f_k(t) \quad (1.10)$$

Then the sequence $\{F_k\}$ converges and

$$\lim_{t \rightarrow x} f(t) = \lim_{k \rightarrow \infty} F_k \quad (1.11)$$

Product Spaces and Fubini's Theorem

Definition 1.24 (Products of Measure Spaces). [Bog07, Vol. 1, §3.3., p. 180] Let (A, \mathcal{A}, μ) and (B, \mathcal{B}, ν) be two spaces with finite non-negative measures. On the space $A \times B$ we consider sets of the form $A_k \times B_k$ where $A_k \in \mathcal{A}$ and $B_k \in \mathcal{B}$ are called measurable rectangles. Let

$$\mu \times \nu(A_k \times B_k) := \mu(A_k)\nu(B_k)$$

be the product of the measures. And let $\mathcal{A} \otimes \mathcal{B}$ denote the σ -algebra generated by all measurable rectangles; this σ -algebra is called the product of the σ -algebra \mathcal{A} and \mathcal{B} .

Theorem 1.25. [Bog07, Vol. 1, §3.3.1., p. 180] The set function $\mu \times \nu$ is countably additive on the algebra generated by all measurable rectangles and uniquely extends to a countably additive measure, denoted by $\mu \otimes \nu$, on the Lebesgue completion of this algebra denoted by $\mathcal{A} \otimes \mathcal{B}$.

Theorem 1.26 (Fubini's Theorem). [Bog07, Vol. 1, Theorem 3.4.1., p. 183] Let a set $X \subset A \times B$ be measurable with respect to the measure $\mu \otimes \nu$, i.e., belong to $(\mathcal{A} \otimes \mathcal{B})_{\mu \otimes \nu}$. Then, for μ -a.e. $a \in A$, the set X_a is ν -measurable and the function $a \mapsto \nu(X_a)$ is μ -measurable; similarly for ν -a.e. $b \in B$, the set X_b is μ -measurable and the function $b \mapsto \mu(X_b)$ is ν -measurable. In addition, one has

$$\mu \otimes \nu(X) = \int_A \nu(X_a)\mu(da) = \int_B \mu(X_b)\nu(db) \quad (1.12)$$

We would like to emphasize here that Eq. (1.12) of Theorem 1.26 allows the exchange of order of integration in iterated integrals i.e., $\int_X \int_Y f(x, y) dy dx = \int_Y \int_X f(x, y) dx dy$ under the condition that the function f is measurable on $X \times Y$ with respect to the measure $dx \otimes dy$. In particular it allows the exchange of order of expectation and integration $\int \mathbb{E}[X(t, \omega)] dt = \mathbb{E}[\int X(t, \omega) dt]$ which is often put under the linearity of expectation umbrella. Fubini's theorem also holds for Bochner integrals [Mik78, Theorem 2, p. 93]).

Definition 1.27 (Approximate Identity). [Str03, §3.4., p. 40] Given a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, locally integrable and a sequence of functions $\{\phi_\varepsilon\}_{\varepsilon \in \mathbb{R}_{>0}}$, $\phi_\varepsilon : \mathbb{R}^n \rightarrow \mathbb{R}$, infinitely differentiable and the following properties

$$(1) \int_{\mathbb{R}^n} \phi_\varepsilon(x) dx = 1, \forall \varepsilon > 0$$

$$(2) \phi_\varepsilon(x) \geq 0, \forall x \in \mathbb{R}^n, \varepsilon > 0$$

$$(3) \lim_{\varepsilon \rightarrow 0} \phi_\varepsilon(x) = 0, \forall x \neq 0, \text{ uniformly in } |x| \geq y, \text{ for any } y > 0$$

then the function $f \mapsto f * \phi_\varepsilon$ is called an approximate identity.

Lemma 1.28. [DK10, Lemma 1.6., p. 12] and [DK10, Lemma 11.6., p. 118] For any function f and ϕ_ε as in Definition 1.27 and $|(f * \phi_\varepsilon)(x)| < \infty$ a.e. for all $\varepsilon > 0$, then

$$\lim_{\varepsilon \rightarrow 0} f * \phi_\varepsilon = f \quad \text{a.e.} \quad (1.13)$$

L^p -norm, L^p -spaces and Hölder's Inequality

Definition 1.29 (L^p -norm). [Bog07, Vol. 1, §2.11., p. 140] On a measure space (A, \mathcal{A}, μ) , a set $B \in \mathcal{A}$, a μ -measurable function, and $1 \leq p < \infty$, we define the L^p -norm as

$$\|f\|_{L^p(B)} := \left(\int_B |f(x)|^p \mu(dx) \right)^{\frac{1}{p}} \quad (1.14)$$

If the L^p -norm is over the whole space A we write for short

$$\|f\|_{L^p} := \|f\|_{L^p(A)} \quad (1.15)$$

Definition 1.30 (L^p -space). [Bog07, Vol. 1, §2.11., p. 139] We define L^p as the set of all equivalence classes of functions for which the L^p -norm is finite. In particular, we define functions f and g equivalent, if the difference norm of $\|f - g\|_{L^p} = 0$.

Theorem 1.31 (Hölder's Inequality). [Bog07, Vol. 1, Theorem 2.11.1., p. 140] Suppose that $1 < p < \infty$, $q = p/(p - 1)$, $f \in L^p$, and $g \in L^q$. Then $fg \in L^1$ and

$$\|fg\|_{L^1} \leq \|f\|_{L^p} \|g\|_{L^q}$$

1.2 Hyperreal Numbers and Nonstandard Analysis

We will present an abbreviated construction of the hyperreal numbers, which is mostly taken from [Gol98]. More rigorous and detailed introductions can be found in [ACH97] and [DR89]. A concise, yet quite complete exposition can be found in [LW15].

Definition 1.32 (Filters). [Gol98, §2.3., p. 18] *Let I be a non-empty set. The power set of I is the set*

$$\mathcal{P}(I) := \{A : A \subseteq I\} \quad (1.16)$$

of all subsets of I . A filter on I is the non-empty collection $\mathcal{F} \subseteq \mathcal{P}(I)$ of subsets of I satisfying the following axioms:

- $A, B \in \mathcal{F} \Rightarrow A \cap B \in \mathcal{F}$
- $A \in \mathcal{F}, A \subseteq B \subseteq I \Rightarrow B \in \mathcal{F}$

A filter \mathcal{F} contains \emptyset if and only if $\mathcal{F} = \mathcal{P}(I)$. We say that \mathcal{F} is a proper filter if $\emptyset \notin \mathcal{F}$.

An ultrafilter is a proper filter that satisfies that for any $A \subseteq I$, either $A \in \mathcal{F}$ or $I \setminus A \in \mathcal{F}$.

$\mathcal{F}^i := \{A \subseteq I : i \in A\}$ is an ultrafilter, called the principal ultrafilter generated by i .

Definition 1.33 (Set of All Sequences of Real Numbers). [Gol98, §3.1., p. 24] *Let $\mathbb{R}^{\mathbb{N}}$ be the set of all sequences of real numbers of the form $r = \{r_k\}_{k \in \mathbb{N}}$. Further, we denote constant sequences as:*

$$\underline{r} := \{r, r, r, \dots\}_{k \in \mathbb{N}} \quad r \in \mathbb{R} \quad (1.17)$$

Definition 1.34 (Equivalence Modulo an Ultrafilter). [Gol98, §3.2., p. 24] *Let \mathcal{F} be a fixed non-principal ultrafilter on the set \mathbb{N} . Given the two sequences $r = \{r_k\}_{k \in \mathbb{N}}$ and $s = \{s_k\}_{k \in \mathbb{N}}$, we define the relation \equiv on $\mathbb{R}^{\mathbb{N}}$ as*

$$r \equiv s \Leftrightarrow \{k \in \mathbb{N} : r_k = s_k\} \in \mathcal{F} \quad (1.18)$$

When this relation holds, we say the two sequences agree on a large set, or agree almost everywhere modulo \mathcal{F} .

We denote the set $\{k \in \mathbb{N} : r_k = s_k\}$ by $\llbracket r = s \rrbracket$.

Similarly, we define

$$\begin{aligned} \llbracket r < s \rrbracket &:= \{k \in \mathbb{N} : r_k < s_k\} \\ \llbracket r > s \rrbracket &:= \{k \in \mathbb{N} : r_k > s_k\} \\ \llbracket r \leq s \rrbracket &:= \{k \in \mathbb{N} : r_k \leq s_k\} \\ \llbracket r \geq s \rrbracket &:= \{k \in \mathbb{N} : r_k \geq s_k\} \end{aligned} \quad (1.19)$$

We denote the equivalence class of a sequence $r \in \mathbb{R}^{\mathbb{N}}$ under \equiv by

$$[r] := \{s \in \mathbb{R}^{\mathbb{N}} : r \equiv s\} \quad (1.20)$$

Definition 1.35 (Set of Hyperreal Numbers). [Gol98, §3.6., p. 25] Given a fixed non-principal ultrafilter \mathcal{F} on the set \mathbb{N} , we define the set of hyperreal numbers as the quotient set of $\mathbb{R}^{\mathbb{N}}$ by \equiv :

$${}^*\mathbb{R} := \{[r] : r \in \mathbb{R}^{\mathbb{N}}\} \quad (1.21)$$

Given $r, s \in \mathbb{R}^{\mathbb{N}}$, we define:

$$\begin{aligned} [r] + [s] &:= [\{r_k + s_k\}_{k \in \mathbb{N}}] \\ [r] \cdot [s] &:= [\{r_k \cdot s_k\}_{k \in \mathbb{N}}] \end{aligned} \quad (1.22)$$

and

$$[r] < [s] \Leftrightarrow \llbracket r < s \rrbracket \in \mathcal{F} \quad (1.23)$$

We denote hyperreal numbers and hyperreal valued functions with a prescript-*, e.g., ${}^*r \in {}^*\mathbb{R}$, ${}^*f : A \rightarrow {}^*\mathbb{R}$.

Theorem 1.36 (Structure of Hyperreal Numbers). [Gol98, Theorem 3.6.1., p. 25] The structure $({}^*\mathbb{R}, +, \cdot, <)$ is an ordered field with zero element $[0]$, unity element $[1]$, and the corresponding additive and multiplicative inverses.

Theorem 1.37. [Gol98, Theorem 3.7.1., p. 27] Given $r \in \mathbb{R}$ the map $r \mapsto [r]$ is an order preserving field isomorphism from \mathbb{R} into a subset of ${}^*\mathbb{R}$.

Definition 1.38. For any operation \otimes between a real number $r \in \mathbb{R}$ and a hyperreal number ${}^*b \in {}^*\mathbb{R}$ we use the above map, to first map the real number into the space of hyperreal numbers:

$$r \otimes {}^*b := [r] \otimes {}^*b \quad (1.24)$$

Definition 1.39 (Infinitesimals and Infinite Numbers). [Gol98, §3.8., p. 27] We define the sequence

$$h_\varepsilon := \left\{ \frac{1}{k} \right\}_{k \in \mathbb{N}} \quad (1.25)$$

Then $[0] < [h_\varepsilon]$, but for any $r \in \mathbb{R}_{>0}$ we have $[h_\varepsilon] < [r]$. We call $[h_\varepsilon]$ an infinitesimal.

Similarly, we define the sequence

$$h_\omega := \{k\}_{k \in \mathbb{N}} \quad (1.26)$$

Then for any $r \in \mathbb{R}_{>0}$ we have $[0] < [h_\omega]$. We call $[h_\omega]$ an infinite number.

Theorem 1.40 (Standard Part). [Gol98, Theorem 5.6.1., p. 53] For every non-infinite hyperreal number *b there exists exactly one real number r for which the difference $([x] - {}^*b)$ is infinitesimal. We call this number r the standard part or shadow of *b . We denote this as $r = \text{st}({}^*b)$.

Theorem 1.41 (Arithmetic Operations and the Standard Part). [Gol98, Theorem 5.6.2., p. 53] Given two hyperreal ${}^*r, {}^*s \in {}^*\mathbb{R}$, both non-infinite, then

$$\begin{aligned} \text{st}({}^*r \pm {}^*s) &= \text{st}({}^*r) \pm \text{st}({}^*s) \\ \text{st}({}^*r \cdot {}^*s) &= \text{st}({}^*r) \text{st}({}^*s) \\ \text{st}({}^*r / {}^*s) &= \text{st}({}^*r) / \text{st}({}^*s) \\ \text{st}(|{}^*r|) &= |\text{st}({}^*r)| \\ {}^*r \leq {}^*s &\Rightarrow \text{st}({}^*r) \leq \text{st}({}^*s) \end{aligned} \tag{1.27}$$

Note that the above theorem cannot be generalized to infinitely repeated operations. E.g., given a sequence of infinitesimal hyperreal numbers *r_k , the infinite sum $\sum_k {}^*r_k$ could become non-infinitesimal. Or in other words $\sum_k \text{st}({}^*r_k) = 0 \not\Rightarrow \text{st}(\sum_k {}^*r_k) = 0$.

The Lebesgue integral can be extended to hyperreal valued functions and measures using the Loeb measure. For details see [Gol98, §16., p. 203], [LW15, §6., p. 179], and [ACH97, p. 91].

In this work we will restrict ourselves to the simpler case of hyperreal valued functions on measure spaces with real valued measures. Thus we can treat integrals over hyperreal valued functions as Bochner integrals.

1.3 Fourier Analysis

Definition 1.42 (Fourier Transform). [Gra14, Definition 2.2.8., p. 108 and §2.2.4., p. 113] Given a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $f \in L^1 \cup L^2$. We define the Fourier transform as

$$\mathcal{F}\{f\}(\xi) := \int_{\mathbb{R}^n} f(x) e^{-2\pi i x \cdot \xi} dx$$

For this discussion we restrict ourselves to f being in $L^1 \cup L^2$. The Fourier transform is more generally defined on tempered distributions, whose introduction would require quite a bit more discussion, while offering only marginal benefit for the later chapters. We refer the reader who is so inclined to [Gra14, §2.3., p. 119 and §2.4., p. 133] for an introduction to the topic.

Definition 1.43 (Inverse Fourier Transform). [Gra14, Definition 2.2.13., p. 111] Given a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $f \in L^1 \cup L^2$ we define the inverse Fourier transform as

$$\mathcal{F}^{-1}\{f\}(x) := \int_{\mathbb{R}^n} f(\xi) e^{2\pi j x \cdot \xi} d\xi$$

In a slight abuse of notation, in case f is a function of multiple variables or when we want to be explicit about which variable the Fourier transform is performed over, we will use $\mathcal{F}\{f(x)\}$ to denote that the Fourier of a function $f(x, y, \dots)$ is performed over the variable x with all other variables held constant.

Definition 1.44 (Self-Similarity Function). [Lap09, Definition 11.2.1., p. 186] We define the self-similarity function $R_{ff} : A \rightarrow \mathbb{C}$ of a function $f : A \rightarrow \mathbb{C}$ as

$$R_{ff}(\tau) = \int_A f(t + \tau) \overline{f(t)} dt$$

This function is more commonly referred as *auto-correlation function* although the auto-correlation function of a continuous-time stochastic process is only then equivalent to the self-similarity function, if the stochastic process is real valued and ergodic². Following [Lap09] and in order to avoid confusion with the auto-correlation function of stochastic processes, we use the term self-similarity function instead. It follows immediately from Theorem 1.31 that the self-similarity function converges for all functions $f \in L^2$. It can be also shown that the self-similarity function is finite a.e. for functions in L^1 (c.f. [Lap09, §11.4., p. 198]).

Theorem 1.45 (Fourier Transform of the Self-Similarity Function). [Lap09, §11.4., p. 198] The Fourier transform of the self-similarity function R_{ff} is equal to the energy spectral density (ESD) of f :

$$\mathcal{F}\{R_{ff}\}(\xi) = |\mathcal{F}\{f\}(\xi)|^2$$

1.4 Probability Theory

Definition 1.46 (Probability Space and Probability Measure). [Øks03, Definition 2.1.1., p. 7] We call the measure space $(\Omega, \mathcal{F}, \mathbb{P})$ a probability space if its measure \mathbb{P} fulfils the following conditions:

- (1) \mathbb{P} is a non-negative, countably additive measure
- (2) $\mathbb{P}(\emptyset) = 0$ and $\mathbb{P}(\Omega) = 1$

We call \mathbb{P} a probability measure.

2. Simplified, ergodic means that, for a random process, the moments of time samples converge to the same values as the moments of the ensemble samples. For the equality mentioned here, only ergodicity in mean is required, i.e., that the time average converges to the same value as the ensemble average.

Definition 1.47 (Random Variable and Random Function). [*Çm11*, §2.1., p. 51] Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a measure space (A, \mathcal{A}, μ) , a mapping $X : \Omega \rightarrow A$ is called a random variable provided it is measurable with respect to \mathcal{F} and \mathcal{A} , that is $\{\omega : X(\omega) \in A_k\} \in \mathcal{F}$ for all $A_k \in \mathcal{A}$.

If X is a mapping from a product space $X : T \times \Omega \rightarrow A$ for some set (or space) T , then X is called a random function. This is also known as a stochastic process denoted as $\{X_t\}_{t \in T}$ or $X(t, \omega)$ (c.f. [*Øks03*, Definition 2.1.4., p. 10]).

Definition 1.48 (Probability of an Event). [*Çm11*, §2.1., p. 51] Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a random variable $X : \Omega \rightarrow A$ and an event $E \subseteq A$, we denote the probability of the event E occurring as

$$\mathbb{P}[X \in E] := \mathbb{P}(\{\omega : \omega \in \Omega, X(\omega) \in E\})$$

We denote by

$$\mathbb{P}[X \in E | X_0 \in E_0] := \frac{\mathbb{P}(\{\omega : \omega \in \Omega, X(\omega) \in E, X_0(\omega) \in E_0\})}{\mathbb{P}(\{\omega : \omega \in \Omega, X_0(\omega) \in E_0\})}$$

the conditional probability of event E under the condition of the random variable $X : \Omega \rightarrow A$ being equal to $X_0 \in E_0$ with $\mathbb{P}[X_0 \in E_0] > 0$ for this discussion. In a slight abuse of notation, we denote with $\mathbb{P}(\cdot | X_0 \in E_0)$ the probability measure of a set $W \subseteq \Omega$ and the set $W_0 = \{\omega : \omega \in \Omega, X_0(\omega) \in E_0\}$ as $\mathbb{P}(W | X_0 \in E_0) = \mathbb{P}(W | W_0) := \mathbb{P}(W \cap W_0) / \mathbb{P}(W_0)$.

Definition 1.49 (Expectation). [*Øks03*, §2.1., p. 9] and [*Øks03*, Appendix B, p. 319] For a random variable $X : \Omega \rightarrow \mathbb{R}$, we call the integral

$$\mathbb{E}[X] := \int_{\Omega} X(\omega) \mathbb{P}(d\omega)$$

the expectation or the mean of X . We denote by

$$\mathbb{E}[X | X_0 \in E_0] := \int_{\Omega} X(\omega) \mathbb{P}(d\omega | X_0 \in E_0)$$

the conditional expectation of X subject to an event $X_0 \in E_0$.

Definition 1.50 (Variance). [*Çm11*, §2.2., p. 60] For a random variable $X : \Omega \rightarrow \mathbb{R}$, we call

$$\mathbb{V}[X] := \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$$

the variance of X .

Definition 1.51 (Normal Distribution). We denote by $\mathcal{N}(\mu, \sigma^2)$ a random variable $X : \Omega \rightarrow \mathbb{R}$ with the probability distribution $\mathbb{P}[X = x] = 1/(\sigma\sqrt{2\pi})e^{(-1/2)((x-\mu)/\sigma)^2}$. We call X normal distributed or Gauss distributed with mean μ and variance σ^2 .

Definition 1.52 (Uniform Distribution). We denote by $\mathcal{U}(A)$ a random variable $X : \Omega \rightarrow \mathbb{R}^n$ with the probability distribution $\mathbb{P}[X = x] = \{1/|A|$ if $x \in A, 0$ otherwise} over a measurable set $A \subset \mathbb{R}^n$. We call X uniformly distributed over A .

Theorem 1.53 (Kolmogorov's Extension Theorem). [Øks03, Theorem 2.1.5., p. 11] Given a subset $T \subseteq \mathbb{R}$, a finite-dimensional stochastic process

$$\begin{aligned} X : T \times \Omega &\rightarrow \mathbb{R}^n \\ (t, \omega) &\mapsto X(t, \omega) \end{aligned}$$

the distributions of the process X are the measures μ_{t_1, \dots, t_k} defined on \mathbb{R}^{nk} , $k \in \mathbb{N}$, with measurable sets $F_k \subseteq \mathbb{R}^n$ at times $t_1, \dots, t_k \in T$ by

$$\mu_{t_1, \dots, t_k}(F_1 \times \dots \times F_k) := \mathbb{P}[X(t_1, \cdot) \in F_1, \dots, X(t_k, \cdot) \in F_k]$$

For all $t_1, \dots, t_k \in T$, $k \in \mathbb{N}$ let ν_{t_1, \dots, t_k} be probability measures on \mathbb{R}^{nk} , and measurable sets $F_k \subseteq \mathbb{R}^n$, such that

$$\nu_{t_{\sigma(1)}, \dots, t_{\sigma(k)}}(F_1 \times \dots \times F_k) = \nu_{t_1, \dots, t_k}(F_{\sigma^{-1}(1)} \times \dots \times F_{\sigma^{-1}(k)})$$

for all permutations σ on $\{1, 2, \dots, k\}$ and

$$\nu_{t_1, \dots, t_k}(F_1 \times \dots \times F_k) = \nu_{t_1, \dots, t_k, t_{k+1}, \dots, t_{k+m}}(F_1 \times \dots \times F_k \times \mathbb{R}^n \times \dots \times \mathbb{R}^n)$$

for all $m \in \mathbb{N}$, where the set on the right hand side has a total of $k + m$ factors.

Then there exists a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a stochastic process X as defined above such that

$$\nu_{t_1, \dots, t_k}(F_1 \times \dots \times F_k) = \mathbb{P}[X_{t_1} \in F_1, \dots, X_{t_k} \in F_k]$$

for all $t_i \in T$, $k \in \mathbb{N}$ and all measurable sets F_k .

A more general version and its proof can be found in [Bog07, Vol. 2, §7.7., p. 95]. For the discussion here, this simpler version, which is restricted to stochastic processes with values in \mathbb{R}^n is sufficient.

1.5 Fractional Calculus

Definition 1.54 (Holmgren-Riemann-Liouville Integral). [Pod99, §2.3.2., p. 65]. We define the left-sided Holmgren-Riemann-Liouville fractional order integral of a function $f : \mathbb{R} \rightarrow \mathbb{R}^n$ from a to t of order $p \in \mathbb{R}_{>0}$ as

$${}_a^{\tau}I_t^p f(\tau) = \frac{1}{\Gamma(p)} \int_a^t (t - \tau)^{p-1} f(\tau) d\tau$$

We differ here slightly from the common definition by introducing the integration variable τ as a pre-superscript, which is usually the same as the end of the integration limit t , i.e., ${}_a I_t^p f(t)$, in order to be able to distinguish the integration variable in repeated fractional integration, e.g., ${}_a^{\tau}I_t^p {}_a^{\tau'}I_t^p f(\tau)g(\tau')$. We also note that we do not define it as differ-integral as we limit ourselves to fractional integration only. The introduction of fractional differentiation leads to various complications. A discussion of this can be found in [Cam18].

Theorem 1.55 (Fourier Transform of Fractional Integrals). [Pod99, §2.9.3., p. 111]. Given a fractional integral $g(t) = {}_{-\infty}^{\tau}I_t^p f(\tau)$, the Fourier transform of the fractional integral is

$$\mathcal{F}\{g\}(\xi) = \frac{1}{(2\pi j\xi)^p} \mathcal{F}\{f\}(\xi)$$

1.6 Special Functions

Definition 1.56 (Incomplete Beta Function). [WJ +10, Formula 8.17.1, p. 183]. For $0 < x < 1$ and $a, b > 0$ we define the incomplete beta function as

$$B(a, b; x) = \int_0^x t^{a-1} (1-t)^{b-1} dt \quad (1.28)$$

There exists an analytic continuation of $B(a, b; x)$ for $x \in \mathbb{C}$.

Definition 1.57 (Hypergeometric Function). [WJ +10, Formula 15.2.1, p. 384]. For $x \in \mathbb{C}$, and $a, b, c \in \mathbb{R}$ we define the (Gauss) hypergeometric function as

$$F(a, b; c; x) = \sum_{n=0}^{\infty} \frac{a^{\bar{n}} b^{\bar{n}}}{c^{\bar{n}} n!} x^n \quad (1.29)$$

For real x , $F(a, b; c; x)$ is only defined for $-\infty < x < 1$, unless one of the four quantities a , b , $(c - a)$, or $(c - b)$ is a strictly negative integer, then $F(a, b; c; x)$ may be reduced to a polynomial.

Lemma 1.58. [WJ +10, Formula 8.17.7, p. 183]. The hypergeometric function and the incomplete beta function are related by the following formula

$$F(a, b; b + 1; x) = \frac{b}{x^b} B(b, 1 - a; x) \quad (1.30)$$

For more details on the incomplete beta function and the hypergeometric function see [WJ +10, §8.17., p. 183], [OMS09, §58, p. 603], [WJ +10, §15., p. 383], [OMS09, §60, p. 627], and [Bat53, Vol. 1, §2, p. 56].

PART I

Noise and its Mathematical Description

CHAPTER 2

White and $1/f^\alpha$ -Noise in Nature and Technology

Noise has been and still is the single most important limit in everything involving measurements. Thus it is not surprising that noise has been very well studied over the past one and a half centuries.

Based on its power spectral density we categorize noise into two broad groups: white noise and $1/f^\alpha$ -noise. White noise processes are usually processes with events that happen independently in large numbers and whose effect becomes only apparent as an average, which, by the central limit theorem, lets the distribution of those apparent averages approach a Gaussian probability distribution. The independence of each single event leads to the overall visible average having no memory and thus no time correlation. Which in turn means that its power spectral density is independent of the frequency, which we call white.

White noise has been historically the more important of the two, as many measurements are done fast¹. Thus it comes with no surprise that white noise is well covered in literature and every engineering discipline has its ways how to deal with it. But as technology evolved, the white noise level decreased and measurement times became longer and longer, thus $1/f^\alpha$ -noise became more and more important. So much so that certain types of measurements today are limited by $1/f^\alpha$ -noise and not by white noise anymore.

$1/f^\alpha$ -Noise goes by a few names, such as power law noise, for its power spectral density (PSD) decaying with a power of the frequency. Or flicker noise, by its most common member $1/f$ -noise. There is a certain mystique that surrounds $1/f^\alpha$ -noise, partially because it still defies our understanding and partially because it is so ubiquitous. This widespread occurrence seems to suggest that a generic mathematical and physical explanation might exist, however, the only generally accepted mathematical description is fractional Brownian motion (i.e., the half integral of white noise). But even for fractional Brownian motion, no physical process is known to explain its occurrence or gives an adequate explanation for $1/f^\alpha$ -noise.

1. Fast here means in relation to the corner frequency, below which $1/f^\alpha$ -noise has a higher power spectral density than white noise.

The first description of $1/f$ -noise is attributed to Johnson [Joh25] in his analysis of data of an experiment to test Schottky's theory of shot noise ("Schrot-Effekt") a few years earlier [Sch18]. Schottky gave a first attempt to describe this noise mathematically in [Sch26] in terms of a train of pulses from exponential relaxation processes. A train of such pulses would produce a power spectrum proportional to $\frac{1}{\lambda^2 + f^2}$ with some constant λ , which is constant near $f = 0$ and goes toward $1/f^2$ for large frequencies f , with a narrow transition region where the power spectrum resembles the $1/f$ -noise found by Johnson. Bernamont later pointed out in [Ber37] that to extend Schottky's description to fully describe $1/f$ -noise was a superposition of such relaxation processes with a uniform distribution of λ . If λ was uniformly distributed between λ_1 and λ_2 then the power spectrum would be proportional to $1/f$ in a range $\lambda_1 \ll f \ll \lambda_2$.

To the best of our knowledge Lévy was the first to introduce fractional integration using the Riemann-Liouville integral in his work in the 1930s and 1940s (c.f. [Lév65]). This approach was picked up by several others. Of these we want to pick out Barnes and Allan [BA66], who were the first to use it as a tool for simulating $1/f$ -noise, and Mandelbrot and van Ness [MN68] whose work is the basis for a majority of the research done on fractional Brownian motion today.

Unfortunately, there is a slight disconnect here. While the traditional $1/f^\alpha$ -noise definition uses the power spectrum, the fractional integration definition creates a function in the time domain that is of infinite energy, i.e., not in $L^1 \cup L^2$ and neither a tempered distribution. Thus neither its Fourier transform nor its self-similarity function² are defined. While normalization over the integration time has been used in the past to define the power spectral density, either by direct Fourier transform or by using the self-similarity function, to the best of our knowledge, the function space on which it converges has never been shown. Neither is there any inverse transform been reported that would give a time domain representation of a function defined by its power spectrum.

In the following chapters we will attempt to remedy this situation by introducing a new semi-norm that normalizes over the integration period (or more generally area). Using this, we will introduce a variant of the Fourier transform with normalization and its inverse transform and prove that the equivalent of the Fourier inverse theorem holds for this modified Fourier transform as well. We will then continue and introduce a mathematical description of white and $1/f^\alpha$ -noise for which the modified Fourier transform is defined.

2. Usually called auto-correlation function in this context. But we would like to point out that the auto-correlation function is only then equal to the self-similarity function, if the stochastic function is ergodic, which is not the case for $1/f^\alpha$ -noise.

CHAPTER 3

The P^p -seminorm and its Integral Transforms

3.1 The Problem with Power-Limited Functions

In signal systems it is common to have functions that do not decay for $t \rightarrow \pm\infty$, i.e., have infinite energy, but have bounded power. Or more formally speaking, functions $f(t) : \mathbb{R} \rightarrow \mathbb{C}$ for which the limit integral

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} |f(t)|^2 dt \quad (3.1)$$

is finite.

While these functions of the class of power limited functions are quite common, they are quite often treated as if they were energy limited functions (i.e., functions in L^2), leading to errors in the mathematical discussion. It is quite important to note that for power limited functions, which are not in L^2 , two of the most commonly used tools in signal analysis, the Fourier transform and the self-similarity function, do not converge and are not defined. Some text get around this by explicitly stating that the functions are time limited (i.e., have compact support), thus forcing them to have limited energy (by virtue of only existing for a limited time) and thus making them L^2 . Other text limit themselves to periodic functions only, for which the Fourier transform leads to a series of weighted Dirac distributions¹, which are still infinite, but allow a way of integration that gives a sensible inverse transform. But one should be careful here, as many textbooks then tacitly extend this analysis to general functions, which is mathematically incorrect.

Unfortunately, these approaches make the proper treatment of the important class of power-limited functions of infinite length impossible. But by slight modification of the self-similarity function and Fourier transform it is possible to handle power-limited functions and extend

1. Dirac distributions behave similar to approximate identities (c.f. Definition 1.27 (Approximate Identity)).

This chapter is work done solely by the author and has not been published otherwise yet.

our well known tools to this class as well. To this end, we will start by an extension of the L^p -norm and then show how to modify well known integral transforms such that they converge for functions from the space induced by this extended norm.

3.2 The P^p -seminorm and its Associated Function Space

We generalize Eq. (3.1) to a semi-norm borrowing from the L^p -norm as follows:

Definition 3.1 (P^p -seminorm). *Given a measure space (A, \mathcal{A}, μ) with a σ -finite measure μ , an increasing sequence of sets $\{A_k\}_{k \in \mathbb{N}}$, $A_k \in \mathcal{A}$ such that $A_k \subseteq A_{k+1}$ and $\mu(A_k) > 0$ for all k , $\lim_{k \rightarrow \infty} A_k = A$, and a measurable function $f : A \rightarrow \mathbb{R}^n$, $p \in]0, \infty[$, we define the P^p -seminorm as*

$$\|f\|_{P^p} := \lim_{k \rightarrow \infty} \left(\frac{1}{\mu(A_k)} \int_{A_k} |f(x)|^p \mu(dx) \right)^{\frac{1}{p}} \quad (3.2)$$

We say that the P^p -seminorm of a function f exists, if Eq. (3.2) converges to a finite value.

Similar to the L^p -norm the P^p -seminorm induces a function space, with the set of functions in P^p being a superset of the set of functions in L^p . It is interesting to note that for a function $f \in P^p$ the expression $f(x)/(\mu(A_k))^{1/p}$ mimics a function in L^p in the limit $k \rightarrow \infty^2$. Unfortunately this expression only makes sense when f is evaluated within A_k , hence a general mapping from P^p to L^p is not possible.

2. Please note that $f(x)/(\mu(A_k))^{1/p}$ does not necessarily converge to zero for all x

Lemma 3.2. *Given a function $f \in P^p$ and a sequence of sets $\{A_k\}_{k \in \mathbb{N}}$ as in Definition 3.1, there exists a subsequence $\{A_j\}_{j \in \mathbb{N}}$ of $\{A_k\}$ such that the sequence of functions*

$$f_j(x) := \left(\frac{1}{\mu(A_j)} \right)^{\frac{1}{p}} f(x) \quad (3.3)$$

are in $L^p(A_j)$ for all j .

Proof. It is possible to choose the sequence $\{A_j\}$ such that

$$\frac{1}{\mu(A_j)} \int_{A_j} |f(x)|^p \mu(dx) \quad (3.4)$$

exists and is finite for all j . As $\mu(A_j)$ is in \mathbb{R} and strictly positive, we can pull it into the integral and into the absolute value:

$$\int_{A_j} \left| \left(\frac{1}{\mu(A_j)} \right)^{\frac{1}{p}} f(x) \right|^p \mu(dx) = \int_{A_j} |f_j(x)|^p \mu(dx) = \left(\|f_j\|_{L^p(A_j)} \right)^p < \infty \quad (3.5)$$

□

We note that for all j the sequence of functions $\{f_k\}_{k \geq j}$ converges uniformly on A_j almost everywhere (a.e.).

Next we would like to define a type of Fourier transform on functions in P^p . But before we can define the P^p -Fourier transform and show that the basic properties of the regular Fourier transform hold for it, we need two additional theorems. The first one is the P^p equivalent of Theorem 1.21 (Dominated Convergence) and a second one that allows the exchange of integration and limit when multiplying a function with a convergent sequence of functions.

Theorem 3.3 (P^p -Convergence). *Given a function $f \in P^p$ and a sequence of sets $\{A_k\}_{k \in \mathbb{N}}$ as in Definition 3.1. Let $\{f_k\}$ be the sequence of functions*

$$f_k(x) := \begin{cases} \left(\frac{1}{\mu(A_k)}\right)^{\frac{1}{p}} f(x) & \text{for } x \in A_k \\ 0 & \text{otherwise} \end{cases} \quad (3.6)$$

and ³

$$\begin{aligned} {}^*f_\infty : A &\rightarrow {}^*\mathbb{R}^n \\ x &\mapsto [\{f_k(x)\}_{k \in \mathbb{N}}] \end{aligned} \quad (3.7)$$

3. ${}^*f_\infty$ is the hyper-real equivalent of $f(x) = \lim_{k \rightarrow \infty} f_k(x)$ in Theorem 1.21 (Dominated Convergence).

Then ${}^*f_\infty$ is in L^p (i.e., $\|{}^*f_\infty\|_{L^p}$ is defined and finite) and

$$\int_A ({}^*f_\infty(x))^p \mu(dx) = \lim_{k \rightarrow \infty} \int_A (f_k(x))^p \mu(dx) \quad (3.8)$$

Proof. Without loss of generality, we can assume that sequence $\{A_k\}$ has been chosen such that the functions f_k are in L^p for all k .

${}^*f_\infty \in L^p$ follows directly from Lemma 3.2 and L^p being a complete metric space. Thus $({}^*f_\infty)^p$ is measurable and integrable. We also note that for all $k \in \mathbb{N}$

$$\begin{aligned} |{}^*f_\infty(x)| &\leq |f_k(x)| && \text{for all } x \in A_k \\ |{}^*f_\infty(x)| &\geq |f_k(x)| && \text{for all } x \notin A_k \end{aligned} \quad (3.9)$$

Further, following the proof of Theorem 1.21 (Dominated Convergence) in [Bog07, Vol. 1, Theorem 2.8.1., p. 130]: By Theorem 1.18 (Absolute Continuity of the Lebesgue Integral), for every $\varepsilon > 0$ and every $k \in \mathbb{N}$ there exists a $\delta > 0$ such that for every $D \in \mathcal{A}$, $\mu(D) < \delta$

$$\int_D |f_k(x)|^p \mu(dx) < \frac{\varepsilon}{8} \quad \text{and} \quad \int_D |{}^*f_\infty(x)|^p \mu(dx) < \frac{\varepsilon}{8} \quad (3.10)$$

Let B be any element of \mathcal{A} such that $0 < \mu(B) < \infty$, then by Theorem 1.20 (Egoroff's Theorem) there is a set $A_\delta \in \mathcal{A}$ such that $A_\delta \subseteq B$, $\mu(B \setminus A_\delta) < \delta$, and the functions $\{f_k\}$ converge uniformly to ${}^*f_\infty$ on A_δ . Hence there exists a $K \in \mathbb{N}$ such that for all $k > K$

$$|(f_k(x))^p - ({}^*f_\infty(x))^p| \leq \frac{\varepsilon}{2\mu(B) + 1} \quad x \in A_\delta \quad (3.11)$$

Therefore, for all $k > K$ we have

$$\begin{aligned} & \left| \int_B (f_k(x))^p \mu(dx) - \int_B ({}^*f_\infty(x))^p \mu(dx) \right| \\ & \leq \int_B |(f_k(x))^p - ({}^*f_\infty(x))^p| \mu(dx) \\ & \leq \int_{B \setminus A_\delta} |(f_k(x))^p - ({}^*f_\infty(x))^p| \mu(dx) \\ & \quad + \int_{A_\delta} |(f_k(x))^p - ({}^*f_\infty(x))^p| \mu(dx) \end{aligned} \quad (3.12)$$

Using Eq. (3.9) on the first integral and Eq. (3.11) on the second integral this simplifies to

$$\begin{aligned} & \leq 2 \int_{(B \cap A_k) \setminus A_\delta} |f_k(x)|^p \mu(dx) \\ & \quad + 2 \int_{B \setminus A_k \setminus A_\delta} |{}^*f_\infty(x)|^p \mu(dx) + \frac{\varepsilon}{2\mu(B) + 1} \mu(A_\delta) \\ & \leq \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{2} \\ & = \varepsilon \end{aligned} \quad (3.13)$$

I.e., for any subset $B \subset A$ with $0 < \mu(B) < \infty$ it holds that

$$\int_B ({}^*f_\infty(x))^p \mu(dx) = \lim_{k \rightarrow \infty} \int_B (f_k(x))^p \mu(dx) \quad (3.14)$$

Because the measure μ is σ -finite, there exists a set of sets $B_l, l \in \mathbb{N}$, $B_l \in \mathcal{A}$ such that

$$\begin{aligned} & 0 < \mu(B_l) < \infty \quad \text{for all } l \in \mathbb{N} \\ & B_l \cap B_m = \emptyset \quad \text{for all } l \neq m \\ & \sum_{l=1}^{\infty} \mu(B_l) = \mu(A) \end{aligned} \quad (3.15)$$

Using Proposition 1.19 and Theorem 1.22 (Tannery's Theorem) it follows

$$\begin{aligned}
 & \int_A (*f_\infty(x))^p \mu(dx) \\
 &= \sum_{l=1}^{\infty} \int_{B_l} (*f_\infty(x))^p \mu(dx) \\
 &= \sum_{l=1}^{\infty} \lim_{k \rightarrow \infty} \int_{B_l} (f_k(x))^p \mu(dx) \quad (3.16) \\
 &= \lim_{k \rightarrow \infty} \sum_{l=1}^{\infty} \int_{B_l} (f_k(x))^p \mu(dx) \\
 &= \lim_{k \rightarrow \infty} \int_A (f_k(x))^p \mu(dx)
 \end{aligned}$$

□

Theorem 3.4. *Given a function $f : A \rightarrow \mathbb{R}^n$, an increasing sequence of sets $\{A_k\}_{k \in \mathbb{N}}$, $A_k \subseteq A_{k+1}$, $\lim_{k \rightarrow \infty} A_k = A$, and a sequence of functions $\{g_k(x)\}_{k \in \mathbb{N}}$, $g_k : A \rightarrow \mathbb{R}$, converging uniformly to $g_\infty(x) := \lim_{k \rightarrow \infty} g_k(x)$, and $\lim_{k \rightarrow \infty} |g_k(x) - g_\infty(x)| \mu(A_k) = 0$ (i.e., g_k converging faster than A_k grows), if $\lim_{k \rightarrow \infty} 1/\mu(A_k) \int_{A_k} f(x)g_k(x) dx$ exists and is finite, then*

$$\lim_{k \rightarrow \infty} \frac{1}{\mu(A_k)} \int_{A_k} f(x)g_k(x) dx = \lim_{k \rightarrow \infty} \frac{1}{\mu(A_k)} \int_{A_k} f(x)g_\infty(x) dx \quad (3.17)$$

Proof. Extending the functions f , $g_\infty(x)$ and the sequence of functions $\{g_k(x)\}_{k \in \mathbb{N}}$ to the hyperreal numbers:

$$\begin{aligned}
 *f(x) &:= \left[\left\{ \frac{\mathbb{1}_{A_k}}{\mu(A_k)(x)} f(x) \right\}_{k \in \mathbb{N}} \right] \\
 *g(x) &:= [\{g_k(x)\}_{k \in \mathbb{N}}] \\
 *g_\infty(x) &:= [\underline{g_\infty(x)}]
 \end{aligned} \quad (3.18)$$

Using Theorem 3.3 it follows

$$\begin{aligned}
 \lim_{k \rightarrow \infty} \frac{1}{\mu(A_k)} \int_A f(x)g_k(x) dx &= \text{st} \left(\int_A *f(x) *g(x) dx \right) \\
 \lim_{k \rightarrow \infty} \frac{1}{\mu(A_k)} \int_A f(x)g_\infty(x) dx &= \text{st} \left(\int_A *f(x) *g_\infty(x) dx \right)
 \end{aligned} \quad (3.19)$$

Using that the difference $*g(x) - *g_\infty(x)$ is infinitesimal because g_k converges uniformly to g_∞ and that the integral $1/\mu(A_k) \int_{A_k} f(x)g_k(x) dx$ can only be finite if $\mu(\{x : f(x) \notin \mathbb{R}^n\} \cap \{x : g_k(x) \neq 0\}) = 0$.

$$\begin{aligned} & \text{st} (*f(x) *g(x) - *f(x) *g_\infty(x)) \\ & \text{st} (*f(x) (*g(x) - *g_\infty(x))) \\ & = 0 \quad \text{a.e.} \end{aligned} \tag{3.20}$$

Using $\lim_{k \rightarrow \infty} |g_k(x) - g_\infty(x)| \mu(A_k) = 0$ we can show

$$\begin{aligned} & \text{st} \left(\int_A *f(x) *g(x) dx \right) - \text{st} \left(\int_A *f(x) *g_\infty(x) dx \right) \\ & = \text{st} \left(\int_A *f(x) *g(x) dx - \int_A *f(x) *g_\infty(x) dx \right) \\ & = \text{st} \left(\int_A *f(x) *g(x) - *f(x) *g_\infty(x) dx \right) \\ & = \text{st} \left(\int_A *f(x) (*g(x) - *g_\infty(x)) dx \right) \\ & = 0 \end{aligned} \tag{3.21}$$

Thus it follows

$$\begin{aligned} & \lim_{k \rightarrow \infty} \frac{1}{\mu(A_k)} \int_A f(x)g_k(x) dx \\ & = \text{st} \left(\int_A *f(x) *g(x) dx \right) \\ & = \text{st} \left(\int_A *f(x) *g_\infty(x) dx \right) \\ & = \lim_{k \rightarrow \infty} \frac{1}{\mu(A_k)} \int_A f(x)g_\infty(x) dx \end{aligned} \tag{3.22}$$

□

3.3 P^p -seminorm Integral Transforms

Even though it is not generally possible to map functions in P^p to functions in L^p , it is possible to modify integral transforms for functions in L^p such that they work for functions in P^p . The two transforms we will focus on are the self-similarity function and the Fourier transform:

Definition 3.5 (P^p -Self-Similarity Function). *Given a metric space A closed under addition, a complete measure space (A, \mathcal{A}, μ) , a sequence of sets $\{A_k\}_{k \in \mathbb{N}}$, with $A_k \in \mathcal{A}$ such that $A_k \subseteq A_{k+1}$ and $\mu(A_k) > 0$ for all k , $\lim_{k \rightarrow \infty} A_k = A$, a function $f : A \rightarrow \mathbb{C}$, $f \in P^p$, and using a point $\tau \in A_1$ we define the P^p -self-similarity function as the sequence of functions*

$$R_{P^p, ff, k}(\tau) := \left(\frac{1}{\mu(A_k)} \int_{A_k} \left(f(x + \tau) \overline{f(x)} \right)^{\frac{p}{2}} \mu(dx) \right)^{\frac{2}{p}} \quad (3.23)$$

and we call

$$R_{P^p, ff, \infty}(\tau) := \lim_{k \rightarrow \infty} R_{P^p, ff, k}(\tau) \quad (3.24)$$

the limit of the P^p -self-similarity function of f .

Please note, we square the result of the integration in order to make the P^p -self-similarity function match the common self-similarity function (c.f. Definition 1.44).

We would also like to point out that the P^2 -self-similarity function is equal to the auto-covariance function (c.f. [Lap09, Definition 25.4.4., p. 517]) $K_{XX}(\tau) = \mathbb{E} [X(t, \omega) \overline{X(t + \tau, \omega)}]$ for wide-sense stationary signals⁴.

From here on, we will use $B(r, x) \subseteq \mathbb{C}^n$, the n -dimensional ball of radius r around a point x instead of A_k , as it leads to slightly simpler notation. It is possible to replace $B(r, x)$ by any increasing sequence $\{A_k\}$, $A_k \subseteq A_{k+1}$, $A_k \subseteq \mathbb{C}^n$ with $k \in \mathbb{N}$ and $\lim_{k \rightarrow \infty} A_k = \mathbb{C}^n$.

We will also dispense with the introduction of the P^p -Fourier transform over functions in the P^p -equivalent of Schwartz space and directly head for functions in P^p . See [Gra14, §2.2.4., p. 113] for a justification of this.

Definition 3.6 (P^p -Fourier Transform). *Given a function $f : \mathbb{C}^n \rightarrow \mathbb{C}^m$, $f \in P^p$, $p \in [1, 2]$, we define the P^p -Fourier transform as the sequence of functions with an increasing parameter r and a fixed parameter y*

$$\mathcal{F}_{B(r, y)} \{f\}(\xi) := \frac{1}{|B(r, y)|} \int_{B(r, y)} f(x) e^{-2\pi j x \cdot \xi} dx \quad (3.25)$$

We use the shorthand

$$\mathcal{F}_r \{f\}(\xi) := \mathcal{F}_{B(r, 0)} \{f\}(\xi) \quad (3.26)$$

when using an integration region of radius r centered around the origin, and we call

$$\mathcal{F}_\infty \{f\}(\xi) := \lim_{r \rightarrow \infty} \mathcal{F}_r \{f\}(\xi) \quad (3.27)$$

the limit of the P^p -Fourier transform of f .

4. Wide-sense stationary signals require that only the first moment or mean is stationary and that the second moment or variance is finite at all times. See [Lap09, Definition 25.4.2., p. 517] for details.

Definition 3.7 (Inverse P^p -Fourier Transform). *Given a function $f : \mathbb{C}^n \rightarrow \mathbb{C}^m$, $f \in P^p$, $p \in [1, 2]$, and its P^p -Fourier transform $F_{B(r,y)}(\xi) = \mathcal{F}_{B(r,y)} \{f\}(\xi)$ we define the inverse P^p -Fourier transform as the sequence of functions with an increasing parameter r and a fixed parameter y*

$$\mathcal{F}_{B(r,y)}^{-1} \{F_{B(r,y)}\}(x) := |B(r,y)| \int_{B(r,y)} F_{B(r,y)}(\xi) e^{2\pi i x \cdot \xi} d\xi \quad (3.28)$$

Similarly, we use the shorthand

$$\mathcal{F}_r^{-1} \{F_r\}(\xi) := \mathcal{F}_{B(r,0)} \{F_{B(r,0)}\}(\xi) \quad (3.29)$$

and we call

$$\mathcal{F}_\infty^{-1} \{F_\infty\}(\xi) := \lim_{r \rightarrow \infty} \mathcal{F}_r^{-1} \{F_r\}(\xi) \quad (3.30)$$

the limit of the inverse P^p -Fourier transform of $F_{B(r,y)}$.

It is worth noting, that similar to the regular Fourier transform of functions in L^p , $p \in [1, 2]$, it is possible to extend the P^p -Fourier transform to distributions. We will not pursue this here to not complicate the discussion unnecessarily.

Similarly to the regular Fourier transform, we will, in a slight abuse of notation, write $\mathcal{F}_r \{f(x)\}$ whenever we want to be explicit over which variable of f the P^p -Fourier transform is taken.

Lemma 3.8. *The P^p -Fourier transform converges for all functions $f \in P^p$, $p \in [1, 2]$*

Proof. Using Lemma 3.2 it can be seen that for a choice of a sequence of $\{r_k\}$ the P^p -Fourier transform of a function f leads to the sequence of (regular) Fourier transforms of a sequence of functions

$$f_r(x) := \begin{cases} f(x)/|B(r_k, 0)| & \text{for } x \in B(r_k, 0) \\ 0 & \text{otherwise} \end{cases} \quad (3.31)$$

with $f_r \in L^p$. As the L^p -norm of the sequence $\{f_r\}$ converges to the P^p -seminorm of f the sequence of Fourier transforms of the sequence $\{f_r\}$ must converge. \square

Next, in order to be able to use the P^p -Fourier transform in the same way as the regular Fourier transform is being used, we need to show that the most commonly used features of the regular Fourier transform have P^p -Fourier transforms equivalents.

We start with some basic properties (c.f. [Gra14, Proposition 2.2.11., p. 109]) and establish that they do hold true:

Proposition 3.9. *Given $f, g \in P^p$, $p \in [1, 2]$, $b \in \mathbb{C}$, and $t > 0$. If the following transforms exist, then*

$$\|\mathcal{F}_\infty \{f\}\|_{L^\infty} \leq \|f\|_{P^1} \quad (3.32)$$

$$\mathcal{F}_r \{f + g\} = \mathcal{F}_r \{f\} + \mathcal{F}_r \{g\} \quad (3.33)$$

$$\mathcal{F}_r \{bf\} = b\mathcal{F}_r \{f\} \quad (3.34)$$

$$\mathcal{F}_r \{f(-x)\}(\xi) = \mathcal{F}_r \{f(x)\}(-\xi) \quad (3.35)$$

$$\mathcal{F}_r \{\overline{f}\}(\xi) = \overline{\mathcal{F}_r \{f\}(-\xi)} \quad (3.36)$$

$$\mathcal{F}_{B(r,0)} \{f(x-y)\}(\xi) = e^{-2\pi jy \cdot \xi} \mathcal{F}_{B(r,-y)} \{f(x)\}(\xi) \quad (3.37)$$

$$\mathcal{F}_r \{e^{2\pi jx \cdot y} f(x)\}(\xi) = \mathcal{F}_r \{f(x)\}(\xi - y) \quad (3.38)$$

$$\mathcal{F}_r \{f(tx)\}(\xi) = \mathcal{F}_{tr} \{f(x)\}(t^{-1}\xi) \quad (3.39)$$

Proof. Equation (3.32) follows directly from the definition of the P^p -Fourier transform, the proof of Lemma 3.8, and the equivalent property of the regular Fourier transform $\|\mathcal{F}\{f\}\|_{L^\infty} \leq \|f\|_{L^1}$.

Equation (3.33) through Eq. (3.39) follow from the definition with a few simple transformations:

$$\begin{aligned} \mathcal{F}_r \{f + g\} &= \frac{1}{|B(r, 0)|} \int_{B(r,0)} (f(x) + g(x)) e^{-2\pi jx \cdot \xi} dx \\ &= \frac{1}{|B(r, 0)|} \int_{B(r,0)} f(x) e^{-2\pi jx \cdot \xi} dx \\ &\quad + \frac{1}{|B(r, 0)|} \int_{B(r,0)} g(x) e^{-2\pi jx \cdot \xi} dx \\ &= \mathcal{F}_r \{f\} + \mathcal{F}_r \{g\} \end{aligned} \quad (3.40)$$

$$\begin{aligned} \mathcal{F}_r \{bf\} &= \frac{1}{|B(r, 0)|} \int_{B(r,0)} bf(x) e^{-2\pi jx \cdot \xi} dx \\ &= \frac{b}{|B(r, 0)|} \int_{B(r,0)} f(x) e^{-2\pi jx \cdot \xi} dx \\ &= b\mathcal{F}_r \{f\} \end{aligned} \quad (3.41)$$

$$\begin{aligned} \mathcal{F}_r \{f(-x)\}(\xi) &= \frac{1}{|B(r, 0)|} \int_{B(r,0)} f(-x) e^{-2\pi jx \cdot \xi} dx \\ &= \frac{1}{|B(r, 0)|} \int_{B(r,0)} f(x) e^{-2\pi j(-x) \cdot \xi} dx \\ &= \frac{1}{|B(r, 0)|} \int_{B(r,0)} f(x) e^{-2\pi jx \cdot (-\xi)} dx \\ &= \mathcal{F}_r \{f(x)\}(-\xi) \end{aligned} \quad (3.42)$$

$$\begin{aligned}
\mathcal{F}_r \{\bar{f}\}(\xi) &= \frac{1}{|B(r,0)|} \int_{B(r,0)} \overline{f(x)} e^{-2\pi j x \cdot \xi} dx \\
&= \frac{1}{|B(r,0)|} \int_{B(r,0)} \overline{f(x) e^{-2\pi j x \cdot \xi}} dx \\
&= \frac{1}{|B(r,0)|} \int_{B(r,0)} \overline{f(x) e^{-2\pi j x \cdot (-\xi)}} dx \\
&= \overline{\mathcal{F}_r \{f\}(-\xi)}
\end{aligned} \tag{3.43}$$

$$\begin{aligned}
\mathcal{F}_r \{f(x-y)\}(\xi) &= \frac{1}{|B(r,0)|} \int_{B(r,0)} f(x-y) e^{-2\pi j x \cdot \xi} dx \\
&= \frac{1}{|B(r,-y)|} \int_{B(r,-y)} f(x) e^{-2\pi j(x+y) \cdot \xi} dx \\
&= \frac{e^{-2\pi j y \cdot \xi}}{|B(r,-y)|} \int_{B(r,-y)} f(x) e^{-2\pi j x \cdot \xi} dx \\
&= e^{-2\pi j y \cdot \xi} \mathcal{F}_{B(r,-y)} \{f(x)\}(\xi)
\end{aligned} \tag{3.44}$$

$$\begin{aligned}
\mathcal{F}_r \{e^{2\pi j x \cdot y} f(x)\}(\xi) &= \frac{1}{|B(r,0)|} \int_{B(r,0)} e^{2\pi j x \cdot y} f(x) e^{-2\pi j x \cdot \xi} dx \\
&= \frac{1}{|B(r,0)|} \int_{B(r,0)} f(x) e^{-2\pi j x \cdot (\xi - y)} dx \\
&= \mathcal{F}_r \{f(x)\}(\xi - y)
\end{aligned} \tag{3.45}$$

$$\begin{aligned}
\mathcal{F}_r \{f(tx)\}(\xi) &= \frac{1}{|B(r,0)|} \int_{B(r,0)} f(tx) e^{-2\pi j x \cdot \xi} dx \\
&= \frac{t^{-n}}{t^{-n} |B(tr,0)|} \int_{B(tr,0)} f(x') e^{-2\pi j x' \cdot (t^{-1} \xi)} dx' \\
&= \mathcal{F}_{tr} \{f(x)\}(t^{-1} \xi)
\end{aligned} \tag{3.46}$$

□

Theorem 3.10 (P^p -Fourier Transform Convolution Theorem). *Given $f, g \in P^p$, $p \in [1, 2]$. If the following transform exist, then*

$$\left\{ \lim_{r \rightarrow \infty} \frac{1}{|B(r,0)|} \int_{B(r,0)} f(x-y)g(y) dy \right\} = \mathcal{F}_\infty \{f(x)\} \mathcal{F}_\infty \{g(x)\} \tag{3.47}$$

Proof.

$$\begin{aligned}
 & \lim_{r_1 \rightarrow \infty} \mathcal{F}_{r_1} \left\{ \lim_{r_2 \rightarrow \infty} \frac{1}{|B(r_2, 0)|} \int_{B(r_2, 0)} f(x-y)g(y) \, dy \right\} \\
 &= \lim_{r_1 \rightarrow \infty} \frac{1}{|B(r_1, 0)|} \int_{B(r_1, 0)} \\
 & \quad \lim_{r_2 \rightarrow \infty} \frac{1}{|B(r_2, 0)|} \int_{B(r_2, 0)} f(x-y)g(y) \, dy e^{-2\pi j x \cdot \xi} \, dx \\
 &= \lim_{r_1 \rightarrow \infty} \frac{1}{|B(r_1, 0)|} \int_{B(r_1, 0)} \\
 & \quad \lim_{r_2 \rightarrow \infty} \frac{1}{|B(r_2, 0)|} \int_{B(r_2, 0)} f(x-y)e^{-2\pi j(x-y) \cdot \xi} g(y)e^{-2\pi j y \cdot \xi} \, dy \, dx
 \end{aligned} \tag{3.48}$$

Using Theorem 3.3, Theorem 1.23 (Moore-Osgood Theorem) and Theorem 1.26 (Fubini's Theorem) we can rearrange the above:

$$\begin{aligned}
 &= \lim_{r_2 \rightarrow \infty} \frac{1}{|B(r_2, 0)|} \int_{B(r_2, 0)} g(y)e^{-2\pi j y \cdot \xi} \\
 & \quad \lim_{r_1 \rightarrow \infty} \frac{1}{|B(r_1, 0)|} \int_{B(r_1, 0)} f(x-y)e^{-2\pi j(x-y) \cdot \xi} \, dx \, dy \\
 &= \lim_{r_2 \rightarrow \infty} \frac{1}{|B(r_2, 0)|} \int_{B(r_2, 0)} g(y)e^{-2\pi j y \cdot \xi} \\
 & \quad \lim_{r_1 \rightarrow \infty} \frac{1}{|B(r_1, 0)|} \int_{B(r_1, -y)} f(x)e^{-2\pi j x \cdot \xi} \, dx \, dy
 \end{aligned} \tag{3.49}$$

By translation invariance of the Lebesgue integral (Theorem 1.17) the inner integral is, in the limit, constant with respect to y , and thus can be factored out:

$$\begin{aligned}
 &= \lim_{r_2 \rightarrow \infty} \frac{1}{|B(r_2, 0)|} \int_{B(r_2, 0)} g(y)e^{-2\pi j y \cdot \xi} \, dy \\
 & \quad \lim_{r_1 \rightarrow \infty} \frac{1}{|B(r_1, 0)|} \int_{B(r_1, 0)} f(x)e^{-2\pi j x \cdot \xi} \, dx \\
 &= \mathcal{F}_\infty \{f(x)\} \mathcal{F}_\infty \{g(x)\}
 \end{aligned} \tag{3.50}$$

□

Lemma 3.11. *Given f, g in P^p , $p \in [1, 2]$*

$$\frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \mathcal{F}_r \{g\}(x) dx = \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \mathcal{F}_r \{f\}(x) g(x) dx \quad (3.51)$$

Proof.

$$\begin{aligned} & \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \mathcal{F}_r \{g\}(x) dx \\ &= \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \frac{1}{|B(r, 0)|} \int_{B(r, 0)} g(y) e^{-2\pi j y \cdot x} dy dx \quad (3.52) \\ &= \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) g(y) e^{-2\pi j y \cdot x} dy dx \end{aligned}$$

Using Theorem 1.26 (Fubini's Theorem) we can rearrange to

$$\begin{aligned} & \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) g(y) e^{-2\pi j y \cdot x} dx dy \\ &= \frac{1}{|B(r, 0)|} \int_{B(r, 0)} g(y) \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) e^{-2\pi j y \cdot x} dx dy \quad (3.53) \\ &= \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \mathcal{F}_r \{f\}(y) g(y) dy \end{aligned}$$

□

Theorem 3.12 (P^p -Fourier Transform Inversion Theorem). *Given a function $f : \mathbb{C}^n \rightarrow \mathbb{C}^m$, $f \in P^p$, $p \in [1, 2]$*

$$\lim_{r \rightarrow \infty} \mathcal{F}_{B(r, 0)}^{-1} \{ \mathcal{F}_{B(r, 0)} \{f\} \} = f \quad a.e. \quad (3.54)$$

Proof. This proof is based on [Gra14, Theorem 2.2.14., p. 112].

Using the same ansatz we define:

$$g_r(\xi) = |B(r, 0)| e^{2\pi j \xi \cdot t} e^{-\pi |\varepsilon \xi|^2} \quad (3.55)$$

The P^p -Fourier transform of $g_r(\xi)$ is thus:

$$\begin{aligned} \mathcal{F}_r \{g_r(\xi)\}(x) &= \frac{1}{|B(r, 0)|} \int_{B(r, 0)} |B(r, 0)| e^{2\pi j \xi \cdot t} e^{-\pi |\varepsilon \xi|^2} e^{-2\pi j \xi \cdot x} d\xi \\ &= \int_{B(r, 0)} e^{2\pi j \xi \cdot t} e^{-\pi |\varepsilon \xi|^2} e^{-2\pi j \xi \cdot x} d\xi \end{aligned} \quad (3.56)$$

From which follows that (c.f. [Gra14, Example 2.2.9., p. 108])

$$\mathcal{F}_\infty \{g_\infty(\xi)\} (x) = \frac{1}{\varepsilon^n} e^{-\pi|(x-t)/\varepsilon|^2} \quad (3.57)$$

We now use g_r in Lemma 3.11, which gives for the left side

$$\begin{aligned} & \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \mathcal{F}_r \{g_r(\xi)\} (x) \, dx \\ &= \lim_{r \rightarrow \infty} \int_{\mathbb{C}^n} \left(\frac{\mathbb{1}_{B(r, 0)}(x)}{|B(r, 0)|} f(x) \mathcal{F}_r \{g_r(\xi)\} (x) \right) \, dx \end{aligned} \quad (3.58)$$

Using the hyperreal valued helper functions

$$\begin{aligned} *f(x) &= \left[\left\{ \frac{\mathbb{1}_{B(r, 0)}(x)}{|B(r, 0)|} f(x) \right\}_{r \in \mathbb{N}} \right] \\ *G(x) &= [\{\mathcal{F}_r \{g_r(\xi)\} (x)\}_{r \in \mathbb{N}}] \\ *G_\infty(x) &= \left[\left\{ \mathcal{F}_\infty \{g_\infty(\xi)\} (x) \right\} \right] \end{aligned} \quad (3.59)$$

We note that $\mathcal{F}_r \{g_r(\xi)\} (x)$ converges exponentially fast in r to $\mathcal{F}_\infty \{g_\infty(\xi)\} (x)$, while $|B(r, 0)|$ only grows as r^n , thus we can apply Theorem 3.4

$$\begin{aligned} & \lim_{r \rightarrow \infty} \int_{\mathbb{C}^n} \left(\frac{\mathbb{1}_{B(r, 0)}(x)}{|B(r, 0)|} f(x) \mathcal{F}_r \{g_r(\xi)\} (x) \right) \, dx \\ &= \text{st} \left(\int_{\mathbb{C}^n} *f(x) *G(x) \, dx \right) \\ &= \text{st} \left(\int_{\mathbb{C}^n} *f(x) *G_\infty(x) \, dx \right) \\ &= \lim_{r \rightarrow \infty} \int_{\mathbb{C}^n} \left(\frac{\mathbb{1}_{B(r, 0)}(x)}{|B(r, 0)|} f(x) \mathcal{F}_\infty \{g_\infty(\xi)\} (x) \right) \, dx \\ &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \mathcal{F}_\infty \{g_\infty(\xi)\} (x) \, dx \\ &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \frac{1}{\varepsilon^n} e^{-\pi|(x-t)/\varepsilon|^2} \, dx \end{aligned} \quad (3.60)$$

And for the right side

$$\begin{aligned} & \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \mathcal{F}_r \{f(x)\} (\xi) g_r(\xi) \, d\xi \\ &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \mathcal{F}_r \{f(x)\} (\xi) |B(r, 0)| e^{2\pi j \xi \cdot t} e^{-\pi|\varepsilon \xi|^2} \, d\xi \\ &= \lim_{r \rightarrow \infty} \int_{B(r, 0)} \mathcal{F}_r \{f(x)\} (\xi) e^{2\pi j \xi \cdot t} e^{-\pi|\varepsilon \xi|^2} \, d\xi \end{aligned} \quad (3.61)$$

Putting both sides together and letting $\varepsilon \rightarrow 0$ for both sides

$$\begin{aligned} & \lim_{\varepsilon \rightarrow 0} \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x) \frac{1}{\varepsilon^n} e^{-\pi|(x-t)/\varepsilon|^2} dx \\ &= \lim_{\varepsilon \rightarrow 0} \lim_{r \rightarrow \infty} \int_{B(r, 0)} \mathcal{F}_r \{f(x)\}(\xi) e^{2\pi j \xi \cdot t} e^{-\pi|\varepsilon \xi|^2} d\xi \end{aligned} \quad (3.62)$$

Using Theorem 1.23 (Moore-Osgood Theorem) and Theorem 1.21 (Dominated Convergence)

$$\begin{aligned} & \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \lim_{\varepsilon \rightarrow 0} \int_{B(r, 0)} f(x) \frac{1}{\varepsilon^n} e^{-\pi|(x-t)/\varepsilon|^2} dx \\ &= \lim_{r \rightarrow \infty} \int_{B(r, 0)} \mathcal{F}_r \{f(x)\}(\xi) e^{2\pi j \xi \cdot t} \left(\lim_{\varepsilon \rightarrow 0} e^{-\pi|\varepsilon \xi|^2} \right) d\xi \end{aligned} \quad (3.63)$$

On the left side, the integral is an approximate identity (c.f. Definition 1.27 and Lemma 1.28). Thus the following equation holds a.e.

$$\lim_{\varepsilon \rightarrow 0} \int_{B(r, 0)} f(x) \frac{1}{\varepsilon^n} e^{-\pi|(x-t)/\varepsilon|^2} dx = \begin{cases} f(t) & t \in B(r, 0) \\ 0 & \text{otherwise} \end{cases} \quad (3.64)$$

On the right side we have $\lim_{\varepsilon \rightarrow 0} e^{-\pi|\varepsilon \xi|^2} = 1$. Now multiplying both sides with $|B(r, 0)|$ we get

$$\begin{aligned} \lim_{r \rightarrow \infty} \mathbb{1}_{B(r, 0)}(t) f(t) &= \lim_{r \rightarrow \infty} |B(r, 0)| \int_{B(r, 0)} \mathcal{F}_r \{f(x)\}(\xi) e^{2\pi j \xi \cdot t} d\xi \\ f(t) &= \lim_{r \rightarrow \infty} \mathcal{F}_{B(r, 0)}^{-1} \{ \mathcal{F}_{B(r, 0)} \{f\} \}(t) \end{aligned} \quad (3.65)$$

□

With this we have proven that the most commonly used properties of the Fourier transform apply to the P^p -Fourier transform as well, and often are almost the same.

Last but not least, we will provide a definition of the PSD that extends the commonly used definition to functions in P^p . The common definition of the PSD is based on the auto-covariance function⁵, but only works for wide-sense stationary functions (c.f. [Lap09, §13.6., p. 213 and §25.7., p. 522]). For wide-sense stationary functions our definition is equivalent to the commonly used one, through the equality of the auto-correlation function and the P^2 -self-similarity function.

Definition 3.13 (Power Spectral Density). *Given a function $f : \mathbb{C}^n \rightarrow \mathbb{C}^n$, $f \in P^p$, $p \in [1, 2]$ we define the P^p -Fourier transform $|\mathcal{F}_\infty \{f\}|^2(\xi)$ to be the PSD of the function f .*

5. There is a second, commonly used definition based on the self-similarity function (c.f. Definition 1.44), but this is then either an energy density function in disguise or there is an implicit re-normalization step.

Theorem 3.14 (Fourier Transform of the P^2 -Self-Similarity Function).
The Fourier transform of the P^2 -self-similarity function of a function $f : \mathbb{C} \rightarrow \mathbb{C}$, $f \in P^2$ is equal to the PSD of f :

$$\lim_{r \rightarrow \infty} \mathcal{F}_r \{R_{P^2, ff, r}(\tau)\}(\xi) = |\mathcal{F}_\infty \{f(x)\}(\xi)|^2 \quad a.e. \quad (3.66)$$

Proof.

$$\begin{aligned} & \lim_{r \rightarrow \infty} \mathcal{F}_r \{R_{P^2, ff, r}(\tau)\}(\xi) \\ &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} f(x + \tau) \overline{f(x)} dx e^{-2\pi j \tau \cdot \xi} d\tau \end{aligned} \quad (3.67)$$

We replace $f(x + \tau)$ by $\mathcal{F}_r^{-1} \{\mathcal{F}_r \{f(x + \tau)\}\}$ and use Theorem 1.26 (Fubini's Theorem) to rearrange

$$\begin{aligned} &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} |B(r, 0)| \\ & \quad \int_{B(r, 0)} \mathcal{F}_r \{f\}(\zeta) e^{2\pi j \tau \cdot \zeta} e^{2\pi j x \cdot \zeta} d\zeta \overline{f(x)} dx e^{-2\pi j \tau \cdot \xi} d\tau \\ &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} |B(r, 0)| \int_{B(r, 0)} \mathcal{F}_r \{f\}(\zeta) \\ & \quad \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \overline{f(x)} e^{2\pi j x \cdot \zeta} dx e^{2\pi j \tau \cdot \zeta} d\zeta e^{-2\pi j \tau \cdot \xi} d\tau \\ &= \lim_{r \rightarrow \infty} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} \\ & \quad |B(r, 0)| \int_{B(r, 0)} \mathcal{F}_r \{f\}(\zeta) \overline{\mathcal{F}_r \{f\}(\zeta)} e^{2\pi j \tau \cdot \zeta} d\zeta e^{-2\pi j \tau \cdot \xi} d\tau \\ &= \lim_{r \rightarrow \infty} |B(r, 0)| \int_{B(r, 0)} \frac{1}{|B(r, 0)|} \int_{B(r, 0)} |\mathcal{F}_r \{f\}(\zeta)|^2 e^{-2\pi j \tau \cdot \zeta} d\zeta e^{2\pi j \tau \cdot \xi} d\tau \\ &= \lim_{r \rightarrow \infty} \mathcal{F}_r^{-1} \left\{ \mathcal{F}_r \left\{ |\mathcal{F}_r \{f\}(\zeta)|^2 \right\}(\tau) \right\}(\xi) \\ &= |\mathcal{F}_\infty \{f\}(\xi)|^2 \end{aligned} \quad (3.68)$$

□

CHAPTER 4

White Noise

White noise is the most commonly used noise type, due its simplicity in description and being a good fit for most application. Yet, the two most common definitions, often used implicitly, are not without problems.

Definition 4.1. *White noise is noise in which the frequency and power spectrum is constant and independent of frequency [Man02, pp 10-11]. (For completeness, phase of the Fourier transform of the signal is independent and identically distributed (iid) $\mathcal{U}([0, 2\pi[)$)*

or equivalently

Definition 4.2. *White noise is a stationary, continuous-time process, for which every point is iid $\mathcal{N}(0, \sigma^2)$.*

The equivalence of these two definitions can be seen using Definition 1.44 (Self-Similarity Function) and its Fourier transform (Theorem 1.45): Starting from Definition 4.2, if the self-similarity function of a continuous-time process approaches zero everywhere but at $\tau = 0$, i.e., approaches a Dirac distribution, then its Fourier transform will be constant. Conversely, if the Fourier transform of the self-similarity function is constant, then the self-similarity function must be the Dirac distribution.

While infinite bandwidth does not immediately cause any problems (besides being unphysical due to infinite power), its equivalence to the function being discontinuous in \mathbb{R} is quite troublesome. By this, the function belongs into the class of functions that are not Lebesgue integrable, which prevents the use of many mathematical tools that are generally assumed to work. Foremost is Fubini's theorem, which is the basis behind linearity of expectation and behind the equivalence of energy spectral density (ESD) and the Fourier transform of the self-similarity function.

This chapter is work done solely by the author and has not been published otherwise yet.

In the following, a slightly more restrictive definition of white noise will be proposed, which circumvents the problems associated with discontinuity, while still being almost as general.

4.1 Mathematical Description

We will henceforth operate on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with Ω being a set, \mathcal{F} being a σ -algebra over Ω and \mathbb{P} being a probability measure over (Ω, \mathcal{F}) . For the sake of brevity, this will not be explicitly mentioned from here on. We further will use $\omega \in \Omega$ to denote a specific outcome.

Definition 4.3 (White Noise). *We call a function X_0 a continuous-time white noise process if it has the following properties*

- (1) $X_0 : \mathbb{R} \times \Omega \rightarrow \mathbb{R}$
 $(t, \omega) \mapsto X_0(t, \omega)$
- (2) $\forall t > 0 \ X_0(t, \cdot) \sim \mathcal{N}(\mu, \sigma^2)$
- (3) $\forall \omega \ X_0(\cdot, \omega)$ is continuous a.e.
- (4) $\exists \varepsilon_{\text{BW}}$ s.t. $\forall t_1, t_2 \in \mathbb{R}_{\geq 0}$ with $|t_2 - t_1| > \varepsilon_{\text{BW}}$, $X_0(t_1, \cdot)$ and $X_0(t_2, \cdot)$ are independent.
- (5) $\forall \omega \ X_0(\cdot, \omega) \in \{-\infty, +\infty\}$ at most on a set of Lebesgue measure zero.
- (6) $\forall t < 0, \forall \omega \ X_0(t, \omega) = 0$

In the following X_0 will always denote a white noise process as defined above.

We show that the following propositions hold for continuous-time white noise processes:

Proposition 4.4. *A random process X_0 as defined above exists.*

Proof. Both conditions of Theorem 1.53 (Kolmogorov's Extension Theorem) are fulfilled by Definition 4.3. \square

Such a function can be easily constructed. Given a set $A = \bigcup_{k=0}^{\infty} A_k$ of non-overlapping ranges $A_k \subset \mathbb{R}_{\geq 0}$ such that $A = \mathbb{R}_{\geq 0}$ and $0 < |A_k| \leq \varepsilon_{\text{BW}}$, a simple function $f(t, \omega) = \sum_{k=0}^{\infty} a_k(\omega) \mathbb{1}_{A_k}(t)$ with $a_k(\omega) \in \mathbb{R}$ being a random variable with a probability distribution such that $f(t, \cdot) \sim \mathcal{N}(\mu, \sigma^2) \forall t$ and independent for each k , then $f(t, \omega)$ fulfills the above properties for X_0 .

If a smooth function is required, this can be achieved using a suitable mollifier function $m(t) \in C^\infty$ with compact support $|\text{supp}(m)| \leq \varepsilon_{\text{BW}}$ and $\int_{\mathbb{R}} m(t) dt = 1$, then the convolution

$$f_m(t, \omega) = \int_{\mathbb{R}} m(\tau - t) f(t, \omega) d\tau \quad (4.1)$$

is a smooth function fulfilling the properties of X_0 .

Proposition 4.5. $X_0(\cdot, \omega)$ is measurable.

Proof. $X_0(\cdot, \omega)$ can be $\pm\infty$ only on a set of Lebesgue measure zero. Likewise Item (3) in Definition 4.3 demands that the set of points where $X_0(\cdot, \omega)$ is discontinuous has a Lebesgue measure zero. \square

Proposition 4.6. $X_0(\cdot, \omega) \in P^p$ for $p \in [1, \infty[$

Proof. Reminding ourselves that the P^p -seminorm is

$$\|X_0(\cdot, \omega)\|_{P^p} = \lim_{T \rightarrow \infty} \left(\frac{1}{T} \int_{-T/2}^{T/2} |X_0(t, \omega)|^p dt \right)^{\frac{1}{p}} \quad (4.2)$$

From Item (5) of Definition 4.3 follows immediately that $|X_0(\cdot, \omega)| < \infty$ a.e., thus $|X_0(\cdot, \omega)|^p < \infty$ a.e.. Similarly, due to Item (2) of Definition 4.3, the Lebesgue measure $|\cdot|$ of the set of values of $X_0(t, \omega)$, $t \in [-T/2, T/2]$ being larger than a value $c \in \mathbb{R}_{\geq 0}$ is can be calculated, namely:

$$\lim_{T \rightarrow \infty} |\{t : |t| < T/2, |X_0(t, \omega)| > c\}| = \lim_{T \rightarrow \infty} T \left(1 + \text{erf} \left(\frac{c - \mu}{\sigma\sqrt{2}} \right) \right) \quad (4.3)$$

Or equivalently, using the probability distribution

$$\lim_{T \rightarrow \infty} |\{t : |t| < T/2, |X_0(t, \omega)| = c\}| = \lim_{T \rightarrow \infty} T \frac{2}{\sigma\sqrt{\pi}} e^{-\frac{1}{2} \left(\frac{c - \mu}{\sigma} \right)^2} \quad (4.4)$$

Using this, we can calculate the P^p -seminorm:

$$\begin{aligned} \|X_0(\cdot, \omega)\|_{P^p} &= \lim_{T \rightarrow \infty} \left(\frac{1}{T} \int_{\mathbb{R}_{\geq 0}} c^p d|\{t : |t| < T/2, |X_0(t, \omega)| = c\}| \right)^{\frac{1}{p}} \\ &= \lim_{T \rightarrow \infty} \left(\frac{2}{\sigma\sqrt{\pi}} \int_{\mathbb{R}_{\geq 0}} c^p e^{-\frac{1}{2} \left(\frac{c - \mu}{\sigma} \right)^2} dc \right)^{\frac{1}{p}} \\ &< \infty \end{aligned} \quad (4.5)$$

\square

Lemma 4.7. The PSD of $X_0(\cdot, \omega)$ exists and is almost white (flat) for $\xi \ll 1/\varepsilon_{\text{BW}}$.

Proof. It follows directly from Proposition 4.6 and Lemma 3.8 that the P^p -Fourier transform of $X_0(\cdot, \omega)$ exists.

Given the P^2 -self-similarity function $R_{P^2, X_0, X_0, r}(\tau)$ of X_0 , its PSD is (Theorem 3.14)

$$|\mathcal{F}_\infty \{f\}|^2(\xi) = \lim_{r \rightarrow \infty} \mathcal{F}_r \{R_{P^2, X_0, X_0, r}\}(\xi) \quad (4.6)$$

The P^2 -self-similarity function $R_{P^2, X_0, X_0, \infty}(\tau)$ will be zero for $|\tau| > \varepsilon_{\text{BW}}$ as $X_0(t_1)$ and $X_0(t_2)$ are independent for $|t_1 - t_2| > \varepsilon_{\text{BW}}$ (Item (4) of Definition 4.3). For $|\tau| < \varepsilon_{\text{BW}}$ the value of $R_{P^2, X_0, X_0, r}$ is not restricted by the definition, but its absolute value is bounded from above by $R_{P^2, X_0, X_0, \infty}(0) = \|X_0\|_{P^2}^2$. I.e., $R_{P^2, X_0, X_0, \infty}$ can be approximated by a pulse of width of ε_{BW} .

Given a rectangular pulse

$$r_{\varepsilon_{\text{BW}}}(\tau) := \begin{cases} 1 & \text{if } \tau < \varepsilon_{\text{BW}} \\ 0 & \text{otherwise} \end{cases} \quad (4.7)$$

its Fourier transform is

$$\mathcal{F}\{r_{\varepsilon_{\text{BW}}}\} = \frac{\varepsilon_{\text{BW}}}{2} \text{sinc}\left(\frac{\varepsilon_{\text{BW}}\xi}{2}\right) \quad (4.8)$$

Hence, if $R_{P^2, X_0, X_0, \infty}$ is approximated by a rectangular pulse of width ε_{BW} , then for every $\varepsilon > 0$ there exists an δ such that for all $\xi < \delta$

$$\left| \frac{d}{d\xi} \lim_{r \rightarrow \infty} \mathcal{F}_r \{R_{P^2, X_0, X_0, r}\}(\xi) \right| < \varepsilon \quad (4.9)$$

□

CHAPTER 5

$1/f^\alpha$ -Noise

$1/f^\alpha$ -Noise noise or, as it is often more commonly known, flicker noise¹, power-law noise, or fractional Brownian motion (fBm), is characterized by its PSD being proportional to $1/f^\alpha$. By giving its frequency characteristic, through Theorem 1.45 (Fourier Transform of the Self-Similarity Function) and Theorem 3.14 (Fourier Transform of the P^2 -Self-Similarity Function) the (P^2 -)self-similarity function and thus its time correlation is defined. One important property to note here is, because $1/f^\alpha$ -noise has some time correlation it is not stationary (not even wide-sense stationary) and thus not ergodic.

Although it is often not further specified, it is usually implied that the ensemble distribution is normally distributed, owing to the fact that the underlying source of $1/f^\alpha$ -noise is assumed to be the average of a large number of stochastic processes. Although we do not know with certainty what the source of $1/f^\alpha$ -noise is, this assumption seems to hold up well enough for practical purposes².

Looking at $1/f^\alpha$ -noise from the fBm perspective, the most common mathematical description is by Van Ness and Mandelbrot [MN68], which is based on the Holmgren-Riemann-Liouville integral [Hol65; Rie47; Lio32]

$$B_H^0(t, \omega) = \frac{1}{\Gamma(H + \frac{1}{2})} \int_0^t (t-s)^{H-\frac{1}{2}} dB(s, \omega) \quad (5.1)$$

with $B(s, \omega)$ being the ordinary Brownian motion and H the Hurst number with $\alpha = 2H + 1$. One slight problem with the analysis in [MN68] is that it is restricted to $0 < H < 1$ or equivalently $1 < \alpha < 3$. This leaves out the, for physics and engineering, very important case of $\alpha = 1$.

1. Flicker noise is usually only used to denote $1/f$ -noise, i.e., for α equals 1, but depending on the text in question it can encompass all of $1/f^\alpha$ -noise.

2. In case this assumption does not hold, a modification to use a more general Levy-process instead of a Wiener-process can be made

This chapter is work done solely by the author and has not been published otherwise yet.

5.1 Mathematical Description

Going back to Eq. (5.1) we follow Barnes and Allan [BA66] but with a slightly different mathematical notation. As noted in the previous chapter, white noise, as it is usually defined, is not Lebesgue integrable. Hence the derivation in [BA66] does not work as intended. Instead we will be using the white noise X_0 as given by 4.3.

Definition 5.1 ($1/f^\alpha$ -Noise). *Given a continuous-time white noise process $X_0(t, \omega)$ as defined in Definition 4.3, in particular with specific μ , σ , and ε_{BW} , we call a function $X_\alpha(t, \omega)$, $t \in \mathbb{R}_{\geq 0}$ defined as the fractional integral*

$$X_\alpha(t, \omega) := {}_0^{\tau}I_t^{\alpha/2} X_0(\tau, \omega) \quad (5.2)$$

a continuous-time $1/f^\alpha$ -noise process.

We note that this function is similar to the definition in [BA66] but circumvents the problem of integrating over a function that is not integrable. It is also similar to the definition in [MN68] without the contribution of noise for $t < 0$. See [SL95] for a discussion of the consequences of this difference in starting point.

Lemma 5.2. *The expected value of $X_\alpha(t, \omega)$ and the conditional expected value, given $X_\alpha(t_0, \omega)$ at time t_0 are*

$$\begin{aligned} \mathbb{E}[X_\alpha(t, \omega)] &= {}_0^{\tau}I_t^{\alpha/2} \mu \\ \mathbb{E}[X_\alpha(t, \omega) | X_\alpha(t_0, \omega)] &= {}_{t_0}^{\tau}I_t^{\alpha/2} \mu + X_\alpha(t_0, \omega) \end{aligned} \quad (5.3)$$

with

$${}_{t_0}^{\tau}I_t^{\alpha/2} \mu = \frac{2\mu(t-t_0)^{\alpha/2}}{\alpha\Gamma(\alpha/2)} \quad (5.4)$$

Proof.

$$\begin{aligned} \mathbb{E}[X_\alpha(t, \omega) | X_\alpha(t_0, \omega)] &= \mathbb{E}\left[{}_{t_0}^{\tau}I_t^{\alpha/2} X_0(\tau, \omega) + X_\alpha(t_0, \omega)\right] \\ &= {}_{t_0}^{\tau}I_t^{\alpha/2} \mathbb{E}[X_0(\tau, \omega)] + X_\alpha(t_0, \omega) \\ &= {}_{t_0}^{\tau}I_t^{\alpha/2} \mu + X_\alpha(t_0, \omega) \end{aligned} \quad (5.5)$$

Using the definition of the fractional integral (c.f. Definition 1.54 (Holmgren-Riemann-Liouville Integral)):

$$\begin{aligned} {}_{t_0}^{\tau}I_t^{\alpha/2} \mu &= \frac{1}{\Gamma(\alpha/2)} \int_{t_0}^t (t-\tau)^{\alpha/2-1} \mu \, d\tau \\ &= \frac{1}{\Gamma(\alpha/2)} \frac{2}{\alpha} (t-t_0)^{\alpha/2} \mu \end{aligned} \quad (5.6)$$

The unconditional case follows from setting $t_0 = 0$ and $X_\alpha(t_0, \omega) = 0$. \square

We note that this means that X_α is not stationary, not even in the wide sense, and thus not ergodic.

Lemma 5.3. *Given $a, x, y, z \in \mathbb{R}$*

$$\int (y-x)^a (z-x)^{a+1} dx = (y-x)^{a+1} (y-z)^{a+1} \left(\frac{z-x}{x-z} \right)^a \frac{1}{a+1} \text{F} \left(-a-1, a+1; a+2; \frac{y-x}{y-z} \right) \quad (5.7)$$

Proof. Using Lemma 1.58 we get

$$\begin{aligned} & (y-x)^{a+1} (y-z)^{a+1} \left(\frac{z-x}{x-z} \right)^a \\ & \quad \frac{1}{a+1} \text{F} \left(-a-1, a+1; a+2; \frac{y-x}{y-z} \right) \\ & = (y-x)^{a+1} (y-z)^{a+1} \left(\frac{z-x}{x-z} \right)^a \\ & \quad \frac{1}{a+1} (a+1) \left(\frac{y-z}{y-x} \right)^{a+1} \int_0^{\frac{y-x}{y-z}} t^a (1-t)^{a+1} dt \end{aligned}$$

Substituting $u = y - t(y-z)$ and further simplifying

$$\begin{aligned} & = (y-z)^{a+2} \left(\frac{z-x}{x-z} \right)^a \left(\frac{1}{z-y} \right)^{a+2} \int_0^x (y-u)^a (z-u)^{a+1} du \\ & = \int_0^x (y-u)^a (z-u)^{a+1} du \end{aligned} \quad (5.8)$$

□

Lemma 5.4. *The conditional variance of $X_\alpha(t, \omega)$ given its value $X_\alpha(t_0, \omega)$ at a time t_0 is bounded from above by*

$$\begin{aligned} & \mathbb{V} [X_\alpha(t, \omega) | X_\alpha(t_0, \omega)] \\ & \leq \frac{4(\sigma^2 + \mu^2)}{\alpha^2 \Gamma^2(\alpha/2)} \left((t-t_0)^{\alpha/2} \left((t-t_0)^{\alpha/2} - (t-t_0 - \varepsilon_{\text{BW}})^{\alpha/2} \right) \right. \\ & \quad + (-1)^{\alpha/2-1} \left((t-t_0)^{\alpha/2} \varepsilon_{\text{BW}}^{\alpha/2} \text{F} \left(-\frac{\alpha}{2} - 2, \frac{\alpha}{2}; \frac{\alpha}{2} + 1; \frac{t-t_0}{\varepsilon_{\text{BW}}} \right) \right. \\ & \quad \left. - (t-t_0 - \varepsilon_{\text{BW}})^{\alpha/2} \varepsilon_{\text{BW}}^{\alpha/2} \text{F} \left(-\frac{\alpha}{2} - 2, \frac{\alpha}{2}; \frac{\alpha}{2} + 1; \frac{t-t_0 - \varepsilon_{\text{BW}}}{-\varepsilon_{\text{BW}}} \right) \right. \\ & \quad \left. - \varepsilon_{\text{BW}}^\alpha \text{F} \left(-\frac{\alpha}{2} - 2, \frac{\alpha}{2}; \frac{\alpha}{2} + 1; 1 \right) \right) \\ & \quad \left. - \left({}^\tau I_{t_0}^{\alpha/2} \mu \right)^2 \right) \end{aligned} \quad (5.9)$$

Proof.

$$\begin{aligned} & \mathbb{V} [X_\alpha(t, \omega) | X_\alpha(t_0, \omega)] \\ &= \mathbb{E} \left[(X_\alpha(t, \omega))^2 | X_\alpha(t_0, \omega) \right] - (\mathbb{E} [X_\alpha(t, \omega) | X_\alpha(t_0, \omega)])^2 \end{aligned} \quad (5.10)$$

with

$$(\mathbb{E} [X_\alpha(t, \omega) | X_\alpha(t_0, \omega)])^2 = \left({}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} \mu + X_\alpha(t_0, \omega) \right)^2 \quad (5.11)$$

and, using Theorem 1.26 (Fubini's Theorem)

$$\begin{aligned} & \mathbb{E} \left[(X_\alpha(t, \omega))^2 | X_\alpha(t_0, \omega) \right] \\ &= \mathbb{E} \left[{}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} X_0(\tau, \omega) {}_{t_0}^{\tau'} \mathbf{I}_t^{\alpha/2} X_0(\tau', \omega) \right] \\ & \quad + 2X_\alpha(t_0, \omega) \mathbb{E} \left[{}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} X_0(\tau, \omega) \right] + (X_\alpha(t_0, \omega))^2 \quad (5.12) \\ &= {}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} {}_{t_0}^{\tau'} \mathbf{I}_t^{\alpha/2} \mathbb{E} [X_0(\tau, \omega) X_0(\tau', \omega)] \\ & \quad + 2X_\alpha(t_0, \omega) {}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} \mu + (X_\alpha(t_0, \omega))^2 \end{aligned}$$

In general, ${}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} {}_{t_0}^{\tau'} \mathbf{I}_t^{\alpha/2} \mathbb{E} [X_0(\tau, \omega) X_0(\tau', \omega)]$ cannot be further simplified, as the expectation $\mathbb{E} [X_0(\tau, \omega) X_0(\tau', \omega)]$ depends on the auto-correlation of $X_0(\tau, \omega)$. But we do know $\mathbb{E} [X_0(\tau, \omega) X_0(\tau', \omega)] = 0$ for $|\tau - \tau'| > \varepsilon_{\text{BW}}$ due to Item (4) in Definition 4.3. We also know $\mathbb{E} [X_0(\tau, \omega) X_0(\tau', \omega)] \leq \mathbb{E} [(X_0(\tau, \omega))^2]$ for all τ, τ' . Thus we can get a bound from above:

$$\begin{aligned} & {}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} {}_{t_0}^{\tau'} \mathbf{I}_t^{\alpha/2} \mathbb{E} [X_0(\tau, \omega) X_0(\tau', \omega)] \\ & \leq {}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} {}_{t_0}^{\tau'} \mathbf{I}_t^{\alpha/2} \mathbb{1}_{\{|\tau - \tau'| < \varepsilon_{\text{BW}}\}} \mathbb{E} [(X_0(\tau, \omega))^2] \quad (5.13) \\ & = {}_{t_0}^{\tau} \mathbf{I}_t^{\alpha/2} {}_{t_0}^{\tau'} \mathbf{I}_t^{\alpha/2} \mathbb{1}_{\{|\tau - \tau'| < \varepsilon_{\text{BW}}\}} (\sigma^2 + \mu^2) \end{aligned}$$

Where in the last step, we used the well known equation for the variance of any random variable Y : $\mathbb{V} [Y] = \mathbb{E} [Y^2] - (\mathbb{E} [Y])^2 \iff \mathbb{E} [Y^2] = \mathbb{V} [Y] + (\mathbb{E} [Y])^2$.

This leaves $\int_{t_0}^{\tau} \Gamma_t^{\alpha/2} \int_{t_0}^{\tau'} \Gamma_t^{\alpha/2} \mathbb{1}_{\{|\tau-\tau'| < \varepsilon_{\text{BW}}\}}$ as the last bit to solve:

$$\begin{aligned}
 & \int_{t_0}^{\tau} \Gamma_t^{\alpha/2} \int_{t_0}^{\tau'} \Gamma_t^{\alpha/2} \mathbb{1}_{\{|\tau-\tau'| < \varepsilon_{\text{BW}}\}} \\
 &= \frac{1}{\Gamma^2(\alpha/2)} \int_{t_0}^t (t-\tau)^{\alpha/2-1} \int_{t_0}^t (t-\tau')^{\alpha/2-1} \mathbb{1}_{\{|\tau-\tau'| < \varepsilon_{\text{BW}}\}} d\tau' d\tau \\
 &= \frac{1}{\Gamma^2(\alpha/2)} \int_{t_0}^t (t-\tau)^{\alpha/2-1} \int_{\max\{t_0, \tau-\varepsilon_{\text{BW}}\}}^{\min\{t, \tau+\varepsilon_{\text{BW}}\}} (t-\tau')^{\alpha/2-1} d\tau' d\tau \\
 &= \frac{1}{\Gamma^2(\alpha/2)} \left(\int_{t_0}^{t_0+\varepsilon_{\text{BW}}} (t-\tau)^{\alpha/2-1} \int_{t_0}^{\tau+\varepsilon_{\text{BW}}} (t-\tau')^{\alpha/2-1} d\tau' d\tau \right. \\
 &\quad + \int_{t_0+\varepsilon_{\text{BW}}}^{t-\varepsilon_{\text{BW}}} (t-\tau)^{\alpha/2-1} \int_{\tau-\varepsilon_{\text{BW}}}^{\tau+\varepsilon_{\text{BW}}} (t-\tau')^{\alpha/2-1} d\tau' d\tau \\
 &\quad \left. + \int_{t-\varepsilon_{\text{BW}}}^t (t-\tau)^{\alpha/2-1} \int_{\tau-\varepsilon_{\text{BW}}}^t (t-\tau')^{\alpha/2-1} d\tau' d\tau \right)
 \end{aligned} \tag{5.14}$$

Using the well known relation $\int_a^b (t-\tau)^\alpha d\tau = 2/\alpha ((t-a)^{\alpha+1} - (t-b)^{\alpha+1})$ and using Lemma 5.3 we can further simplify

$$\begin{aligned}
 &= \frac{2}{\alpha \Gamma^2(\alpha/2)} \left(\int_{t_0}^{t_0+\varepsilon_{\text{BW}}} (t-\tau)^{\alpha/2-1} \left((t-t_0)^{\alpha/2} - (t-\tau-\varepsilon_{\text{BW}})^{\alpha/2} \right) d\tau \right. \\
 &\quad + \int_{t_0+\varepsilon_{\text{BW}}}^{t-\varepsilon_{\text{BW}}} (t-\tau)^{\alpha/2-1} \left((t-\tau+\varepsilon_{\text{BW}})^{\alpha/2} - (t-\tau-\varepsilon_{\text{BW}})^{\alpha/2} \right) d\tau \\
 &\quad \left. + \int_{t-\varepsilon_{\text{BW}}}^t (t-\tau)^{\alpha/2-1} (t-\tau+\varepsilon_{\text{BW}})^{\alpha/2} d\tau \right) \\
 &= \frac{2}{\alpha \Gamma^2(\alpha/2)} \left((t-t_0)^{\alpha/2} \int_{t_0}^{t_0+\varepsilon_{\text{BW}}} (t-\tau)^{\alpha/2-1} d\tau \right. \\
 &\quad + \int_{t_0+\varepsilon_{\text{BW}}}^t (t-\tau)^{\alpha/2-1} (t-\tau+\varepsilon_{\text{BW}})^{\alpha/2} d\tau \\
 &\quad \left. - \int_{t_0}^{t-\varepsilon_{\text{BW}}} (t-\tau)^{\alpha/2-1} (t-\tau-\varepsilon_{\text{BW}})^{\alpha/2} d\tau \right) \\
 &= \frac{2}{\alpha \Gamma^2(\alpha/2)} \left((t-t_0)^{\alpha/2} \frac{2}{\alpha} \left((t-t_0)^{\alpha/2} - (t-t_0+\varepsilon_{\text{BW}})^{\alpha/2} \right) \right. \\
 &\quad - (t-t_0+\varepsilon_{\text{BW}})^{\alpha/2} (-\varepsilon_{\text{BW}})^{\alpha/2} (-1)^{\alpha/2-1} \frac{2}{\alpha} \text{F} \left(-\frac{\alpha}{2} - 2, \frac{\alpha}{2}; \frac{\alpha}{2} + 1; \frac{t-t_0-\varepsilon_{\text{BW}}}{-\varepsilon_{\text{BW}}} \right) \\
 &\quad + \varepsilon_{\text{BW}}^{\alpha} (-1)^{\alpha/2-1} \frac{2}{\alpha} \text{F} \left(-\frac{\alpha}{2} - 2, \frac{\alpha}{2}; \frac{\alpha}{2} + 1; 1 \right) \\
 &\quad \left. - (t-t_0)^{\alpha/2} \varepsilon_{\text{BW}}^{\alpha/2} (-1)^{\alpha/2-1} \frac{2}{\alpha} \text{F} \left(-\frac{\alpha}{2} - 2, \frac{\alpha}{2}; \frac{\alpha}{2} + 1; \frac{t-t_0}{\varepsilon_{\text{BW}}} \right) \right)
 \end{aligned} \tag{5.15}$$

□

Lemma 5.5. For $(t - t_0) \gg \varepsilon_{\text{BW}}$ the variance $\mathbb{V}[X_\alpha(t, \omega) | X_\alpha(t_0, \omega)]$ can be approximately bounded from above by

$$\begin{aligned} & \mathbb{V}[X_\alpha(t, \omega) | X_\alpha(t_0, \omega)] \\ & \lesssim \frac{2(\sigma^2 + \mu^2)}{\alpha \Gamma^2(\alpha/2)} \left(\left(\frac{\alpha \varepsilon_{\text{BW}}}{\alpha/2 - 1} - \frac{2\alpha \varepsilon_{\text{BW}}}{\alpha - 1} \right) (t - t_0)^{\alpha-1} \right. \\ & \quad \left. + \left(\frac{\alpha(\alpha - 1) \varepsilon_{\text{BW}}^2}{\alpha - 2} + \alpha \varepsilon_{\text{BW}}^2 \right) (t - t_0)^{\alpha-2} \right) \\ & \quad - \left({}_{t_0}^{\tau} \mathbb{I}_t^{\alpha/2} \mu \right)^2 \end{aligned} \tag{5.16}$$

Proof. Starting from Eq. (5.15) and remembering that $(x + \varepsilon)^a = x^a + ax^{\alpha-1}\varepsilon + \mathcal{O}(x^{\alpha-2}\varepsilon^2)$ and $(x + \varepsilon)^a - (x - \varepsilon)^a = 2ax^{\alpha-1}\varepsilon + \mathcal{O}(x^{\alpha-3}\varepsilon^3)$ we can approximate:

$$\begin{aligned} & \int_{t_0}^{t_0 + \varepsilon_{\text{BW}}} (t - \tau)^{\alpha/2-1} \left((t - t_0)^{\alpha/2} - (t - \tau - \varepsilon_{\text{BW}})^{\alpha/2} \right) d\tau \\ & + \int_{t_0 + \varepsilon_{\text{BW}}}^{t - \varepsilon_{\text{BW}}} (t - \tau)^{\alpha/2-1} \left((t - \tau + \varepsilon_{\text{BW}})^{\alpha/2} - (t - \tau - \varepsilon_{\text{BW}})^{\alpha/2} \right) d\tau \\ & + \int_{t - \varepsilon_{\text{BW}}}^t (t - \tau)^{\alpha/2-1} (t - \tau + \varepsilon_{\text{BW}})^{\alpha/2} d\tau \\ & \approx \int_{t_0}^{t_0 + \varepsilon_{\text{BW}}} (t - \tau)^{\alpha/2-1} \left((t - t_0)^{\alpha/2} - (t - \tau)^{\alpha/2} + \frac{\alpha}{2} \varepsilon_{\text{BW}} (t - \tau)^{\alpha/2-1} \right) d\tau \\ & + \int_{t_0 + \varepsilon_{\text{BW}}}^{t - \varepsilon_{\text{BW}}} (t - \tau)^{\alpha/2-1} \alpha \varepsilon_{\text{BW}} (t - \tau)^{\alpha/2-1} d\tau \\ & + \int_{t - \varepsilon_{\text{BW}}}^t (t - \tau)^{\alpha/2-1} \left((t - \tau)^{\alpha/2} + \frac{\alpha}{2} \varepsilon_{\text{BW}} (t - \tau)^{\alpha/2-1} \right) d\tau \\ & = \frac{2(t - t_0)^{\alpha/2}}{\alpha/2 - 1} \left((t - t_0)^{\alpha/2} - (t - t_0 - \varepsilon_{\text{BW}})^{\alpha/2} \right) \\ & \quad - \frac{2}{\alpha - 1} \left((t - t_0)^\alpha - (t - t_0 - \varepsilon_{\text{BW}})^\alpha \right) \\ & \quad + \frac{\alpha \varepsilon_{\text{BW}}}{\alpha - 2} \left((t - t_0)^{\alpha-1} - (t - t_0 - \varepsilon_{\text{BW}})^{\alpha-1} \right) \\ & \quad + \frac{\alpha \varepsilon_{\text{BW}}}{\alpha - 1} \left((t - t_0 - \varepsilon_{\text{BW}})^{\alpha-1} - \varepsilon_{\text{BW}}^{\alpha-1} \right) \\ & \quad + \varepsilon_{\text{BW}}^\alpha \left(\frac{2}{\alpha - 1} + \frac{\alpha}{2\alpha - 4} \right) \end{aligned} \tag{5.17}$$

$$\begin{aligned}
 &\approx \frac{\alpha \varepsilon_{\text{BW}}}{\alpha/2 - 1} (t - t_0)^{\alpha-1} \\
 &\quad - \frac{2\alpha \varepsilon_{\text{BW}}}{\alpha - 1} (t - t_0)^{\alpha-1} \\
 &\quad + \frac{\alpha(\alpha - 1) \varepsilon_{\text{BW}}^2}{\alpha - 2} (t - t_0)^{\alpha-2} \\
 &\quad + \alpha \varepsilon_{\text{BW}}^2 (t - t_0)^{\alpha-2} \\
 &= \left(\frac{\alpha \varepsilon_{\text{BW}}}{\alpha/2 - 1} - \frac{2\alpha \varepsilon_{\text{BW}}}{\alpha - 1} \right) (t - t_0)^{\alpha-1} \\
 &\quad + \left(\frac{\alpha(\alpha - 1) \varepsilon_{\text{BW}}^2}{\alpha - 2} + \alpha \varepsilon_{\text{BW}}^2 \right) (t - t_0)^{\alpha-2}
 \end{aligned} \tag{5.18}$$

□

For the case of $\alpha = 2$ the variance can be directly calculated and leads to

$$\mathbb{V} [X_2(t, \omega) | X_2(t_0, \omega)] \leq \frac{(\sigma^2 + \mu^2) \varepsilon_{\text{BW}}}{2} (2(t - t_0) - \varepsilon_{\text{BW}}) - ((t - t_0) \mu)^2 \tag{5.19}$$

Lemma 5.6. *If the PSD of $X_\alpha(\cdot, \omega)$ exists, it is proportional to $1/\xi^\alpha$, i.e., it is a $1/f^\alpha$ -noise process.*

Proof. The PSD of $X_0(\cdot, \omega)$ is almost white (c.f. Lemma 4.7). Using Theorem 1.55 (Fourier Transform of Fractional Integrals) we get

$$\begin{aligned}
 |\mathcal{F}_\infty \{X_\alpha\}|^2(\xi) &= \left| \mathcal{F}_\infty \left\{ {}_0I_\infty^{\alpha/2} X_0 \right\} \right|^2(\xi) \\
 &= \frac{1}{(2\pi j \xi)^\alpha} |\mathcal{F}_\infty \{X_0\}|^2(\xi)
 \end{aligned} \tag{5.20}$$

□

PART II

Noise Propagation in Electronics

CHAPTER 6

The Importance of Time in Modern Metrology

Time, and its inverse frequency, is the SI base unit that we can measure with the highest accuracy and precision. Achieving a precision of a few parts in 10^{-17} is common with optical atomic clocks and an uncertainty below 10^{-18} has been demonstrated¹. (see e.g., [Bre+19; Hob+20]) While the next most-precisely measurable unit is the Volt² with a precision of 10^{-11} [Rüf+18].

The accuracy and ease with which a frequency standard can be built and can be transferred from radio frequencies almost continuously up to the ultra-violet range [Pet+09] has, over time, lead to more and more SI base units to depend on the definition of the second. E.g., the meter was redefined using the speed of light in vacuum as a constant and in combination with the length of the second in 1983 [BIP83]. Today, realizations of the meter are either based on time-of-flight measurements of light or counting of wavelength (i.e., interference patterns) of laser beams with a very accurate frequency. With the 2019 revision of the SI base unit system, all but one of the SI units depend directly or indirectly on the definition of the second [BIP18] (see Fig. 6.1).




Figure 6.1: SI units relation graph after the 2019 redefinition of the SI base units. From [Pis16].

1. Currently, the achieved accuracy is limited to 10^{-16} by the second being defined using the hyperfine splitting of the ground state of the Caesium-133 atom and this being a microwave transition. Efforts to redefine the second using an optical transition are underway, though.

2. Though the Volt is not an SI base unit, it is usually favored above the Ampere due to ease of realization. Today, the Ampere is most commonly realized using a voltage standard (a Josephson voltage standard) and a resistance standard (quantum Hall effect standard). Which makes the Ampere a derived unit in practice.

With this in mind, it is no surprise that methods to measure time and frequency continue to be an important topic. Though the used methods did not fundamentally change over the past decades, they got refined quite a bit. Besides sampling based methods (e.g., [SJ16]) frequency measurement is performed by counting cycles of the input and measuring precisely the time difference between the start of the first cycle and the end of the last cycle [Joh14; Rub12; HP97]. Or in other words, measuring frequency can be performed by measuring the time of events occurring and thus, in most cases, time interval measurement is the fundamental operation in time and frequency measurements.

Kalisz gives a nice overview of the methods used for interval measurement in [Kal03]. Of these methods, two are especially noteworthy for being the most commonly implemented methods these days: The time-to-amplitude conversion method and the tapped delay line method (often called time to digital converter (TDC) for it leading directly to a digital time value without an intermediate analog representation step). The time-to-analog conversion method has been the workhorse of precision time interval measurements from the 1970s until the early 2010s and are still the most common method to do measurements with a precision of a few ps. When only 10s of ps or less precision/resolution is needed TDC have overtaken the time-to-amplitude converters due to their compact construction and cheap, mass-market manufacturing.

6.1 The Current State of Art in Low-Noise Amplifier Design

Common to all measurement methods mentioned above is to label a specific point in time with an electric pulse. This is done by means of a rising or falling edge with a steep slope.

Unfortunately, most experiments do not have an output signal with a steep slope. It is quite often that only a sinusoid with low frequency or a slow ramp is provided. Thus the amplification of this slope, and doing so with as little noise as possible, is paramount for precise measurement.

Noise in amplifiers has been studied quite extensively in the past. Probably the most cited work in this area is by Motchenbacher and Conelly [MC93], where most of the book is devoted on how to design low noise amplifiers including modeling of noise sources and how different stages affect the noise. Ambrózy does delve quite a bit into the mathematical details of noise in his book [Amb82], focusing on probability distributions, their spectra and how they change when passing

through linear circuits. The book by Vasilescu [Vas05], while having some chapters on circuit design, focuses more on the physical sources of noise. More practice oriented books like the book by Tietze, Schenk, and Gamm [TSG19] have a section on low-noise design in most chapters, but while the basic theory and how to apply it in an actual circuit design is covered, more advanced models are not mentioned. The often used book by Horowitz and Hill [HH15, Chapter 8] mostly focuses on the practical aspects of low-noise electronics design.

One aspect of all the above mentioned works is that they only deal with linear circuits. But in our case, if the whole signal would be amplified linearly, the result would be a signal with an amplitude that could not be supported by the following measurement circuit. Adding an element that limits the amplitude adds non-linearity to the circuit, which makes analysis more difficult, especially analysis of the noise performance. This problem was first noted by Dick, Kuhnle, and Sydnor in [DKS90] with the proposed solution of splitting the amplification into multiple stages. Later Collins gave an approximation in [Col96] that allowed to calculate an approximately optimal choice for the gain of each stage, but without giving any bound how accurate the approximation was. A model that split the noise into an input related and circuit delay related was proposed by Calosso and Rubiola in [CR14], but could not fully explain the phase noise spectrum seen in experiments.

To the best of our knowledge, noise performance of non-linear amplifiers has only been studied in the context of oscillators. Notably in [HL98] by Hajimiri and Lee have presented an analysis upon which we will build upon in the following chapters. They model the noise sensitivity as a periodic function, depending on the input signal of the amplifier, thus capturing the non-linearity of the amplifier itself. We would like to mention here that the work by Hajimiri and Lee, although intended for oscillators, is flawed in its mathematical modeling for oscillators as it does not properly account for the memory of the resonator when it comes to phase noise and thus only works for circuits without memory or resonators. For a detailed discussion see [DR97]. In [DMR00] Demir, Mehrotra, and Roychowdhury follow a similar, but more intricate approach, properly accounting for the memory of the resonator using Floquet theory (c.f. [Dem00]).

In the following chapters we will first refine the model introduced by Calosso and Rubiola and give a mathematical / physical explanation for the observed scaling in frequency. Then we will use the same technique to refine Collins' model, which leads to an exact solution for the choice of optimal gain per stage.

CHAPTER 7

Noise Propagation in Electronics

In order to measure the time an event happened precisely it is useful to have an as strong signal as possible, relative to the noise floor, to maximize the signal to noise ratio. Unfortunately, there is a practical limit on how strong a signal can be handled in an electronic system. The major limiting factor being the maximum amplitude (or power) of the signal, either driving the electronics into saturation, with all kinds of side effects, or carrying so much power that the losses lead to excessive heating. As only the zero-crossing of the signal is relevant for time measurement, it is common to amplify the signal to increase the slope at the zero-crossing, while deliberately limiting the amplitude, effectively turning a sinusoid signal into a square-wave signal, in such a way that the circuit shows as little non-ideal behavior as possible. Unfortunately, limiting the amplitude makes the circuit non-linear, which in itself is a non-ideal behavior.

In the following sections we will explore how this non-linearity and other non-ideal behavior of real circuits affect measurement precision. While the bulk of the analysis is done for comparators, the analysis technique can be applied to all non-linear circuits, starting from “linear” amplifiers, over limiting amplifiers and mixers, to attenuators and more.

7.1 Circuit Model

A sine-to-square converter can be modeled by a comparator, i.e., an amplifier that saturates to ± 1 . To make the model more realistic, but still keep it simple, we assume here, that the converter consists of an ideal, noiseless and zero-delay comparator with a hysteresis $\pm H$ followed by a single noiseless amplifier with a delay t_{amp} (see Fig. 8.1). We base our circuit model on two assumptions:

This chapter is work done solely by the author and based on an article published as [Kin18].




Figure 7.1: The circuit model of a sine-to-square converter is simplified to a noiseless comparator input stage with some hysteresis $\pm H$ and all input related noise being lumped together into the offset voltage V_{ofs} . Noise due to power supply variation $V_{DD,n}$ is modeled using an amplification stage with a delay of t_{amp} that only depends on the supply voltage.

- (1) That all noise that is directly or indirectly related to the input signal like phase noise of the input, variations of the offset voltage, phase noise of the sine-to-square converter,..., can be folded into a single input referred noise source.
- (2) That all delay variations within the circuit, due to temperature, voltage,..., can be folded into a single, delay related noise source and that his noise source is independent of the input related noise.

Thus we use this split in our model to separate noise contributions due to different processes.

Additionally we make the following assumptions as used later in the text:

- (3) The (integrated) magnitude of phase noise is small enough that the traditionally used small angle approximation holds.
- (4) The delay variations due to noise or variations of circuit parameters are small enough relative to the period of the input signal that the small angle approximation holds.
- (5) Independence of noise sources with different physical origin.

The input signal

$$V_i(t) = (V_0 + V_{i,AM}(t)) \sin(2\pi\nu_0 t + \varphi_i(t)) \quad (7.1)$$

with the two noise parts, the input amplitude noise $V_{i,AM}(t)$ and the input phase noise $\varphi_i(t)$ enters the comparator, which has a hysteresis of $\pm H(t)$ and an input offset voltage of $V_{ofs}(t)$. The output of the comparator gets further amplified and delayed by time $\Delta t_{amp}(t)$ by the following amplifier. We further assume that fluctuations and noise on the power supply $V_{DD,n}$ do not affect the comparator (e.g., it being an ideally symmetrical differential pair) and model the effect of $V_{DD,n}$ as variations in the amplifier delay t_{amp} .

As we are only interested in the phase noise contribution of the amplifier, we will ignore the input phase noise $\varphi_i(t)$ for the further analysis. The amplitude noise $V_{i,AM}(t)$ is included to determine its effect on the output phase noise, due to AM-PM conversion through the hysteresis of the comparator.

Multi-stage converters can easily be modeled by series connection of the elementary stage in Fig. 8.1. In longer chains or with sufficiently large input amplitude, there will be a slew-rate saturation. This can be used to simplify all following stages to single amplifier stages, without the comparator, and fold the input noise into the variation of the amplifier delay t_{amp} and by setting the amplifier gain to 1. If filters are used between stages, like in the case of Collins style sine-to-square converters [Col96] care has to be taken to account for the change in noise properties in each stage.

7.2 Noise Sources

Assuming the hysteresis H of the comparator is symmetric around the zero point with offset voltage V_{ofs} , i.e., $V_{ofs} \pm H$, and both are small enough such that the small angle sine approximation can be used around the zero-crossing point (see Fig. 7.2), then the propagation delay through the comparator is

$$t_{delay}(t) \approx t_{ofs}(t) + t_H(t) + t_{amp}(t) \quad (7.2)$$

$$= \frac{V_{ofs}(t) + H(t)}{2\pi\nu_0(V_0 + V_{i,AM}(t))} + t_{amp}(t). \quad (7.3)$$

As we are interested in the variation of the delay, we will be focusing at Δt_{delay} . Assuming delay variations are small relative to the signal period one can split the contributions:

$$\Delta t_{delay}(t) \approx \Delta t_{ofs}(t) + \Delta t_H(t) + \Delta t_{AM}(t) + \Delta t_{amp}(t) \quad (7.4)$$

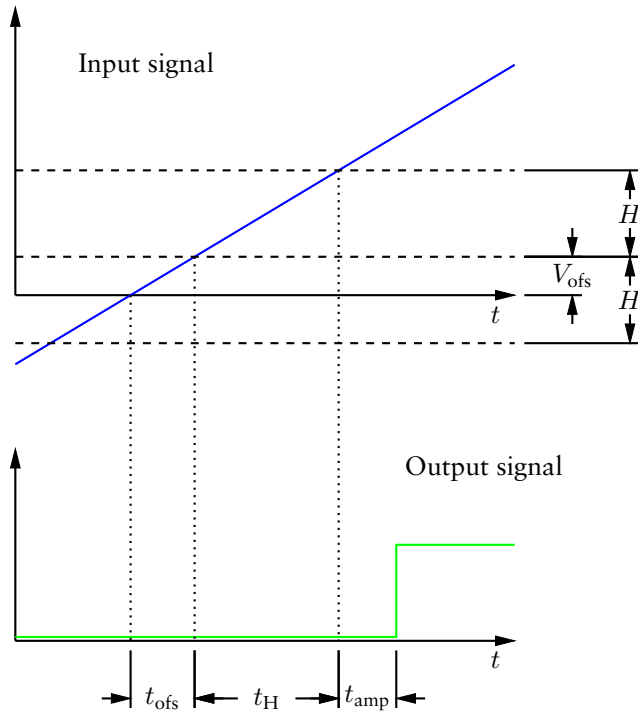


Figure 7.2: The zero crossing of the input signal gets delayed by the offset voltage V_{ofs} and the hysteresis $\pm H$ and the delay t_{amp} through the amplifier stage. The offset voltage and hysteresis related delays t_{ofs} and t_{H} depend not only on the values of V_{ofs} and H respectively, but also on the slew-rate of the input signal. The amplifier delay t_{amp} only depends on the supply voltage V_{DD} .

with the first order Taylor approximations:

$$\Delta t_{\text{ofs}}(t) \approx \frac{\partial \Delta t_{\text{delay}}}{\partial V_{\text{ofs}}} \Delta V_{\text{ofs}}(t) = \frac{1}{2\pi\nu_0 V_0} \Delta V_{\text{ofs}}(t) \quad (7.5)$$

$$\Delta t_{\text{H}}(t) \approx \frac{\partial \Delta t_{\text{delay}}}{\partial H} \Delta H(t) = \frac{1}{2\pi\nu_0 V_0} \Delta H(t) \quad (7.6)$$

$$\Delta t_{\text{AM}}(t) \approx \frac{\partial \Delta t_{\text{delay}}}{\partial V_{\text{i,AM}}} \Delta V_{\text{i,AM}}(t) \approx \frac{H}{2\pi\nu_0 V_0^2} \Delta V_{\text{i,AM}}(t) \quad (7.7)$$

As all these noise sources have physically different origins, we can assume that they are independent. It should be noted, though, that there might be independent factors that lead to correlation between these noise sources. E.g., large variations in supply voltage or temperature can lead to correlated shifts in offset voltage V_{ofs} and hysteresis H . But in most cases, especially in stabilized measurement instruments, these variations should be small and thus independence should be a valid first order approximation.

Although, there are many factors that affect delay through the circuit, for the sake of brevity, we will ignore them all in this discussion but V_{DD} . In case these other factors need to be properly modeled, a simple substitution of V_{DD} or $V_{DD,n}$ is sufficient. As the exact relationship of V_{DD} and the delay is not known, we approximate it by

$$\begin{aligned}\Delta t_{\text{amp}}(t) &= \frac{\partial \Delta t_{\text{amp}}}{\partial V_{DD}} V_{DD,n}(t) + \mathcal{O}(V_{DD,n}^2) \\ &= cV_{DD,n}(t) + \mathcal{O}(V_{DD,n}^2) \\ &\approx cV_{DD,n}(t),\end{aligned}\tag{7.8}$$

with the circuit dependent parameter c . This, of course, removes time dependent delay variations due to e.g., aging or temperature, which can be significant at long time scales. But these contributions are easy to add later in the analysis and are left out, at the moment, for simplicity.

7.3 Noise Translation and Scaling

The phase noise is defined by

$$S_{\varphi}(\xi) = \varphi_{\text{rms}}^2(\xi)\tag{7.9}$$

with the phase fluctuation φ measured over a bandwidth of 1 Hz [Fer+09]. As the phase relates to time with $\varphi = 2\pi\nu_0 t$ we can write

$$S_{\varphi}(\xi) = (2\pi\nu_0)^2 \langle \Delta t^2 \rangle_{\xi}\tag{7.10}$$

with $\langle x \rangle_{\xi}$ denoting the average of the absolute value of the Fourier coefficient in a 1/2 Hz neighborhood of the frequency ξ of the signal $x(t)$:

$$\langle x \rangle_{\xi} = \int_{\xi-1/2}^{\xi+1/2} \left| \int_{-\infty}^{\infty} x(t) e^{-2\pi j \zeta t} dt \right| d\zeta\tag{7.11}$$

For reasons of being concise, we ignore here the mathematical details of integrating over time series of random signals, which might potentially be non-continuous and assume all random signals are of finite bandwidth and thus integrable. We also assume all integrals go over finite time (measurement) intervals in order for them to be defined in case of $1/f^{\alpha}$ -noise which otherwise would lead to infinite signal power. Alternatively, wherever the Fourier transform is used, it could be substituted for the P^p -Fourier transform (c.f. Definition 3.6), leading to the same result for infinite energy, finite power signals. Please refer to Part I for details.

As all discussed noise sources are assumed to be independent, we can write

$$\begin{aligned} S_\varphi(\xi) &= (2\pi\nu_0)^2 \langle \Delta t_{\text{delay}}^2 \rangle_\xi \\ &= \frac{1}{V_0^2} \langle \Delta V_{\text{ofs}}^2 \rangle_\xi + \frac{1}{V_0^2} \langle \Delta H^2 \rangle_\xi \\ &\quad + \frac{H^2}{V_0^4} \langle \Delta V_{\text{i,AM}}^2 \rangle_\xi + (2\pi\nu_0)^2 c^2 \langle V_{\text{DD,n}}^2 \rangle_\xi \end{aligned} \quad (7.12)$$

Different frequency scaling for different noise sources becomes already evident. The input related noise processes do not scale with ν_0 while the delay related noise does scale with ν_0^2 . These are the φ -type and x -type noises, respectively, as discussed in [CR14].

Impulse Sensitivity Function

In [Ega90] Egan noted that white phase noise gets aliased due to sampling. Formally, this can be described by using the impulse sensitivity function (ISF) as introduced by Hajimiri and Lee in [HL98]. We slightly modify it to adapt it for the more general setting of sine-to-square converters¹:

$$\Delta\varphi(t) = \int_{-\infty}^t \Xi(\tau) n(\tau) d\tau \quad (7.13)$$

with $\Xi(t)$ being the ISF and $n(t)$ being the effecting noise. Please note that $\Xi(t)$ is implicitly also a function of the circuit and its parameters, which also include the input signal. In other words, if the shape or amplitude of the input signal changes, this will potentially result in a change of the shape of $\Xi(t)$. For simplicity, we assume that these changes are small and can be neglected. The ISF for a sine-to-square converter can be approximated by a comb of alternating positive and negative rectangular pulses:

$$\Xi(t) = \sum_{k=-\infty}^{\infty} \Pi\left(\frac{t}{\tau_w} - kT_0\right) - \sum_{k=-\infty}^{\infty} \Pi\left(\frac{t}{\tau_w} - kT_0 - \tau_d\right) \quad (7.14)$$

with $\Pi(\cdot)$ being the rectangular pulse function:

$$\Pi(t) = \begin{cases} 1, & |t| < \frac{1}{2} \\ \frac{1}{2}, & |t| = \frac{1}{2} \\ 0, & |t| > \frac{1}{2} \end{cases} \quad (7.15)$$

using a period of $T_0 = 1/\nu_0$ and a pulse width of τ_w . τ_d denotes the phase shift between the positive and the negative pulses and is related to the duty cycle of the output signal and depends, in our circuit model, on the input signal amplitude V_0 and the offset voltage V_{ofs} . In a first

1. Most publications use $\Gamma(t)$ for the ISF, but to avoid confusion with the Γ -function we will use $\Xi(t)$ here

order approximation, for a 50 % duty cycle $\tau_d = T_0/2$. In order to keep the formulas simple at this point, we assume that the positive and negative pulse, which relate to the positive and negative zero crossing respectively, are of the same magnitude and width, which is not necessarily the case. In case they are not, this will manifest itself in even order harmonics and thus can be easily modelled in the frequency domain, as explained later.

It should be noted, that τ_w depends on the output slew-rate of the converter, thus is proportional to T_0 . Conversely, for slew-rate limited converters, τ_w becomes independent of T_0 (in first order). The Fourier series of the function $\Xi(t)$ can be expressed as

$$\begin{aligned} \Xi(t) = & \frac{\tau_w}{T_0} \sum_{k=-\infty}^{\infty} \text{sinc}(\pi k \nu_0 \tau_w) e^{-\pi j k \nu_0 \tau_w} e^{-2\pi j k \nu_0 t} \\ & + \frac{\tau_w}{T_0} \sum_{k=-\infty}^{\infty} \text{sinc}(\pi k \nu_0 \tau_w) e^{-\pi j k \nu_0 \tau_w} e^{-2\pi j k \nu_0 t} e^{-2\pi j k \nu_0 \tau_d} \end{aligned} \quad (7.16)$$

with $\text{sinc}(x) = \frac{\sin(x)}{x}$. The Fourier series directly explains two phenomena reported in [CR14]: The $1/\nu_0$ scaling of white noise and the $1/\nu_0^2$ and $1/\nu_0$ scaling of flicker noise. We will explore and explain this in the following sections.

Scaling of White Noise

Under the assumption that $\tau_d = T_0/2$ the Fourier transform of $\Xi(t)$ in Eq. (7.16) becomes a Dirac comb like structure with Dirac pulses at odd multiples of ν_0 due to its periodic nature and because of the signal symmetry the even harmonics cancel out. These Dirac pulses $\delta(\xi - (2k + 1)\nu_0)$ have approximately constant amplitude a_k up to the frequency $1/\tau_d$ from which on they decay with $1/f$ or 20 dB/dec. The multiplication with noise in Eq. (7.13) results in a mixing process (c.f. [Lap09, §7., pp. 101]) that converts all noise in distance ν_0 to one of the Dirac pulses $a_k \delta(f - (2k + 1)\nu_0)$ down into the signal passband around ν_0 with an amplitude that is proportional to the amplitude of the Dirac pulse. Because the noise in each down-converted frequency region is uncorrelated, the total down converted noise becomes a

geometric sum

$$S_{\varphi,white,total} \propto \sum_{k=1}^{\infty} a_k^2 \quad (7.17)$$

$$= \sum_{k=1}^{\lfloor \frac{2}{\tau_d \nu_0} \rfloor} a_k^2 + \sum_{k=\lfloor \frac{2}{\tau_d \nu_0} \rfloor + 1}^{\infty} a_k^2 \quad (7.18)$$

$$\approx \frac{2}{\tau_d \nu_0} a_1^2 + a_1^2 \sum_{k=1}^{\infty} \frac{1}{k} \quad (7.19)$$

$$= a_1^2 \left(\frac{2}{\tau_d \nu_0} + \sum_{k=1}^{\infty} \frac{1}{k} \right) \quad (7.20)$$

The harmonic series in Eq. (7.20) grows slowly and can be approximated by $H_n = \sum_{k=1}^n \frac{1}{k} = \ln n + \gamma + \mathcal{O}(1/n)$, with γ being the Euler-Mascheroni constant (c.f. [Knu97, §1.2.7, pp. 75 ff.]). Even though $H_{\infty} = \infty$ and thus $S_{\varphi,white,total} = \infty$, the sum is limited in reality. One reason is that the ISF edges have a finite steepness, which adds a second sinc term to $\Xi(t)$ and thus a second corner frequency after which it decays with 40 dB/dec or $1/\xi^2$. Another is the limited bandwidth of the circuit, which acts similarly by adding a cut-off frequency, after which the noise (and signal) decay with an additional 20 dB/dec. The sum $H_{\infty}^{(r)} = \sum_{k=1}^{\infty} \frac{1}{k^r}$, $r > 1$ is bounded by a small constant (e.g., $H_{\infty}^{(2)} = \pi^2/6$) [Knu97, §1.2.7., pp. 76], thus we can express the total white phase noise as:

$$S_{\varphi,white,total} \propto a_1^2 \left(\frac{2}{\tau_d \nu_0} + \frac{c_{\text{BW}}}{\nu_0} \right) \quad (7.21)$$

$$\propto a_1^2 \frac{c'_{\text{BW}}}{\nu_0} \quad (7.22)$$

with c_{BW} and $c'_{\text{BW}} = 2c_{\text{BW}}/\tau_d$ being (noise) bandwidth dependent constants of the circuit². We conclude that the total white noise of the sine-to-square converter gets an additional scaling with a factor of $1/\nu_0$ due to aliasing induced by the periodicity of the ISF. Thus we end up with:

$$S_{\varphi,white}(\xi) \propto \frac{1}{\nu_0 V_0^2} \langle \Delta V_{\text{ofs}}^2 \rangle_{\xi} + \frac{1}{\nu_0 V_0^2} \langle \Delta H^2 \rangle_{\xi} + \frac{H^2}{\nu_0 V_0^4} \langle \Delta V_{\text{i,AM}}^2 \rangle_{\xi} + \nu_0 c^2 \langle V_{\text{DD,n}}^2 \rangle_{\xi} \quad (7.23)$$

The proportionality factor of Eq. (7.23) depends on the equivalent noise bandwidth, respectively the weighted sum of harmonics contributing to aliasing, and on the ratio τ_w/T_0 from $\Xi(t)$.

In case $\tau_d \neq T_0/2$, then $\Xi(t)$ will also have even harmonics. For white noise, the even harmonics will act the same way as the odd

2. To be precise, c_{BW} is the sum over all harmonics, normalized by the amplitude of the fundamental.

harmonics and add to the proportionality factor of Eq. (7.23). For most systems, one can safely assume that the duty cycle will be close to 50% and thus the even harmonics will be small. Hence it is possible to ignore the effects of even harmonics in a first order approximation.

Scaling of Flicker Noise

Flicker noise is, initially, only present around DC, thus the harmonics up-convert the flicker noise due to the multiplication of the ISF with the noise in Eq. (7.13), which acts as a convolution in the frequency domain. But only the first harmonic up-converts the DC flicker noise into the signal band. Hence, the flicker component derives from Eq. (7.12) directly as:

$$S_{\varphi, \text{flicker}}(f) \propto \frac{1}{V_0^2} \langle \Delta V_{\text{ofs}}^2 \rangle_{\xi} + \frac{1}{V_0^2} \langle \Delta H^2 \rangle_{\xi} + \frac{H^2}{V_0^4} \langle \Delta V_{i, \text{AM}}^2 \rangle_{\xi} + (2\pi\nu_0)^2 c^2 \langle V_{\text{DD}, \text{n}}^2 \rangle_{\xi} \quad (7.24)$$

Scaling in a Multi-Stage Sine-to-Square Converter

If multiple gain stages³ are used in a sine-to-square converter, then each stage acts upon the noise and thus the harmonics of the ISF of each stage alias noise into the signal band. Even if all stages are the same, each stage will have a different $\Xi(t)$ as τ_w will change with the slew rate of the input signal of each stage. Thus a simple multiplication of S_{φ} with the number of stages will, in general, not lead to an accurate result. Nevertheless one can derive scaling rules quite easily:

For white noise, the harmonics of the additional stages each convert noise up, which is then down converted into the signal band by the following stage. As the equivalent noise bandwidth of each stage individually is fixed, the scaling rules in Section 7.3 remain unchanged and thus Eq. (7.23) is still valid with only a larger proportionality factor.

For flicker noise, the up-conversion and the following down conversion of consecutive stages change the behavior slightly. If the duty cycle is exactly 50%, then only odd harmonics will exist, and hence none of the present harmonics will see any flicker noise in a distance of ν_0 . Thus only the first harmonic of each stage will up-convert flicker noise into the signal band and, as with white noise, Eq. (7.24) is still valid with a slightly larger proportionality factor. But, due to the presence of V_{ofs} , the duty cycle will deviate from 50% and give rise to even harmonics. Please note, it is not only the DC component of V_{ofs} that leads to even harmonics, but also its higher frequency noise components that modulate the output signal and with it the ISF. Hence the scaling

3. It is important to point out that most common RF amplifier structures, even if they have only a single amplifying transistor, employ additional transistors for bias point stabilization and as current sources. These additional transistors can act like multiple stages within the amplifier

of flicker noise will change its properties depending on the frequency spectrum of V_{ofs} .

Because now there are harmonics at a spacing of ν_0 , the previously up-converted flicker noise is seen by a harmonic at a distance of ν_0 and is thus converted down again into the signal band. Unlike the aliasing of white noise, the aliased flicker noise ultimately has the same origin, thus all down converted flicker noise components are correlated. Thus for S_φ , this leads to a scaling proportional to $(2\pi\nu_0)^2$.

From Eq. (7.14) we see that the power of the even harmonics relates to the power of the odd harmonics with $|1 + \exp(j2\pi\nu_0\tau_d)| = |\sin(2\pi\nu_0\tau_d)|$. Assuming $\tau_d \approx T_0/2$ we can replace τ_d by it's deviation (noise value) from $T_0/2$:

$$\begin{aligned}\tau_{d,n} &= \tau_d - T_0/2 \\ &= \frac{1}{2\pi\nu_0 V_0} V_{\text{ofs}}(t)\end{aligned}\tag{7.25}$$

leading to

$$\begin{aligned}|\sin(2\pi\nu_0\tau_d)| &= |\sin(2\pi\nu_0\tau_{d,n})| \\ &= \left| \sin\left(\frac{1}{V_0} V_{\text{ofs}}(t)\right) \right| \\ &\approx \left| \frac{1}{V_0} V_{\text{ofs}}(t) \right|\end{aligned}\tag{7.26}$$

To evaluate the effects of Eq. (7.26) on the noise spectrum, we have to take into account, that $\tau_{d,n}$ represents a jitter value due to V_{ofs} . As such, it is subject to the same aliasing and thus scaling laws as S_φ . If the amplification of the first converter stage is large, we can safely assume that $\tau_{d,n}$ is dominated by the first stage. If we also assume that the noise equivalent bandwidth is large and thus (the jitter) $\tau_{d,n}$ is dominated by white noise. We then can ignore the contribution and scaling due to flicker noise. Due to aliasing of white noise, we get an additional scaling term of $1/(2\pi\nu_0)$, as we have already seen with Eq. (7.21). Thus the power of the even harmonics will scale approximately with $1/(2\pi\nu_0 V_0) \langle V_{\text{ofs}} \rangle_\xi$.

The scaling due to aliasing will act differently on different types of noise. While all input related noise sources will see the full effect of aliasing, the V_{DD} related noise component will not. As the V_{DD} related noise acts as a delay in each stage of the multi stage converter, it will only see part of the flicker noise aliasing, depending on which stage the source of the noise was. Thus V_{DD} related noise will see an additional scaling factor between 1 and $2\pi\nu_0$ depending on the exact structure of the sine-to-square converter and which stages contribute how much to the output noise.

	input related noise		delay related noise	
	white	flicker	white	flicker
pure noise	const.	const.	ν^2	ν^2
single state	$1/\nu$	const.	ν	ν^2
multi-stage	$1/\nu$	ν	ν	ν^α

Table 7.1: Summary of the frequency scaling of the different noise types for the pure noise, after a single stage amplifier and after a multi-stage amplifier. The frequency scaling for delay (V_{DD}) related flicker noise after a multi-stage amplifier is ν^α with $\alpha \in [2, 3]$.

Putting the arguments above together, we can conclude that $S_{\varphi,\text{flicker}}$ of a multistage sine-to-square converter gets an additional $(2\pi\nu_0)^2$ term due to aliasing of correlated noise and an $1/(2\pi\nu_0)$ term due to the power scaling of the even harmonics:

$$\begin{aligned}
S_{\varphi,\text{flicker,multi}}(f) &\propto \nu_0 S_{\varphi,\text{flicker}}(f) \\
&\propto \frac{\nu_0}{V_0^2} \langle \Delta V_{\text{ofs}}^2 \rangle_\xi + \frac{\nu_0}{V_0^2} \langle \Delta H^2 \rangle_\xi \\
&\quad + \frac{\nu_0 H^2}{V_0^4} \langle \Delta V_{\text{i,AM}}^2 \rangle_\xi + \nu_0^\alpha c^2 \langle V_{\text{DD,n}}^2 \rangle_\xi
\end{aligned} \tag{7.27}$$

with α , the frequency scaling factor of the V_{DD} related noise, being between 2 and 3.

7.4 Conclusion

We have derived formulas for the scaling of noise in sine-to-square converters, for white and flicker noise, for input and delay related noise, and for single and multi-stage circuits and the combinations thereof. These scaling factors explain the noise scaling as seen in [CR14] and other experiments accurately, including the prior unexplained drop in flicker noise at lower frequencies. A summary of the frequency scaling factors for the different noises can be found in Table 7.1.

We would like to note that in case exact numbers for the noise contributions are needed, these can be easily calculated using the derivation here by accounting for the factors in the few steps where we used proportional values.

We also would like to highlight the importance to reduce even order harmonics in the signal and with that in the ISF to a minimum in multi-stage circuits to keep the up- and down-conversion of $1/f^\alpha$ -noise as low as possible.

CHAPTER 8

Noise Propagation in Hard Limiters

In this chapter we will apply the scaling laws and methodology developed in the previous chapter to an important circuit in time metrology: the hard limiter.

As mentioned previously, to accurately measure the timing of events, the input signal has to have an as high slew rate as possible. Especially with techniques such as the dual mixer time difference (DMTD) [AD75] where the input signal is a sinusoid with very low frequency (in the order of 1 Hz), increasing the slew rate with as little added noise as possible is important. The simplest way to build a zero-crossing detector is by use of limiting amplifiers, i.e. by transforming the (most often sinusoidal) signal into a square wave signal. Collins described in [Col96] how to design the slope-gain in multi-stage limiting amplifiers such that the jitter due to the amplification and squaring-up is minimal. Even though, DMTD is not the only application for zero-crossing detectors, their noise properties have received very little attention over the years since Collins' paper.

We will follow Collins' work to some extent, but instead of approximating a sinusoidal signal by a trapezoid, as he has done, which fails at low gain values, we will use the methodology of the previous chapter to derive an exact solution.

8.1 Circuit Model

The circuit model used in this chapter is a slightly adapted version of what Collins used in [Col96]. A single stage consists of a noiseless amplifier, where all its noise is lumped together into a single noise source at its input v_n^2 . For brevity of notation, we also include the signal's noise in v_n^2 and assume an otherwise noiseless signal at each stage. Please note, that we assume here that the amplifier has no offset voltage and the noise voltage v_n^2 has zero mean. This slightly artificial assumption is justified by the need to precisely control the duty cycle to 50 %

This chapter is work done solely by the author and based on an article published as [Kin19].




Figure 8.1: The circuit model, adapted from [Col96] of a single stage zero-crossing detector is simplified to a noiseless amplifier input stage, with all input referred noise being lumped together into the noise voltage v_n^2 . After the amplifier there is a (noise-less) limiter and a first order filter.

and, by extension, the offset voltage in order to eliminate amplifying flicker noise due to the otherwise present even order harmonics, as we have shown in the previous chapter. For sake of brevity, it is further assumed the circuit does not contain any variable time delay term (e.g. due to power supply noise). The effects of any such delay noise can be analyzed in a similar manner as the input referred noise effects.

The switch in Collins circuit model has been replaced by the more natural clamping diodes. One can still assume these diodes to be noise free in a very good approximation of the real circuit performance. On one hand, only the clamping diodes of the last stage will contribute to the output noise. On the other hand, the noise contribution of the diodes can be neglected compared to the input referred noise that has been amplified through even a single amplifier stage, even if said input noise would be as low as the noise of a single diode.

The filter is here modeled by a (noiseless) RC-filter as a stand in of any, more general low pass filter that could be used.

8.2 Circuit Analysis

The assumption that there is no offset voltage and thus no even harmonic components in the signal and impulse sensitivity function (ISF) together with the assumption that there is no variable (noise dependent) delay result in the flicker noise being dominated by the input referred noise (c.f. [CR14] and Section 7.3). I.e. the output flicker noise is dominated by the flicker noise of the first stage (c.f. Friis formula [Fri44]) and the up-conversion of the flicker noise is due to the fundamental of the ISF at each stage. This leaves us the analysis of the white noise propagation through the circuit.

Following Collins, we want to optimize the variance of the jitter $J^2 = N_{\text{out}}^2/\rho_{\text{out}}^2$ with N_{out}^2 being the output noise power density and ρ_{out} being the output slew rate.

Slew Rate

The output slew rate is obviously the slew rate of the output signal $V_{\text{out}}(t)$ at the zero crossing. Assuming a purely sinusoidal input signal of frequency ν_0 with amplitude 1, total gain $G_n = \prod_{i=1}^n g_i$, with g_i being the individual stage gains¹ and clamping back to 1 again, the output signal becomes $V_{\text{out}}(t) = \min\{\max\{G_n \cos(2\pi\nu_0 t), -1\}, 1\}$. This is a periodic function with the Fourier series coefficients (frequency ν_0 normalized to 1) being

$$\hat{V}_{\text{out}}[k] = \frac{1}{\pi} \left(\left(\frac{2}{k} - \frac{2k}{k^2 - 1} \right) \sin(k\tau) + \frac{2G_n}{k^2 - 1} \sin(\tau) \cos(k\tau) \right) \quad (8.1)$$

with $\tau = \arccos(1/G_n)$ being half of the rise/fall time of the output signal. The slew rate can then easily be calculated as

$$\rho_{\text{out}} = \sum_{k=1}^{\infty} k \hat{V}_{\text{out}}[k] \quad (8.2)$$

Following Collins we can approximate V_{out} by a trapezoidal function for large G_n , which leads to $\hat{V}_{\text{out}}[k]$ decaying with $1/k$ up to the frequency $1/(\pi\tau)$ and with $1/k^2$ from then onward. It is important to note that for small G_n this approximation does not hold and the exact Fourier coefficients have to be used.

Noise

The input referred white noise is being sampled by the ISF, down-converting broadband noise from higher frequencies down to the signal band. Assuming the ISF is solely due to the non-linearity of the clamping circuit (i.e. the amplifier is perfectly linear), the input noise becomes amplified by the stage gain g_i and then (additional) noise power is down-converted from higher frequencies. As the exact calculation of the ISF is not possible without knowing the exact circuit we approximate the ISF by a pulse train of rectangular pulses with width τ_w . This results in the ISF's Fourier coefficients $\hat{\Xi}[k]$ being constant up to a frequency of $1/(\pi\tau_w)$ after which they decay with $1/k$. To ensure convergence, we further assume there is an upper frequency ν_{max} of after which the ISF decays faster than $1/k$. This is a safe assumption for real circuits, as invariably all components will have a finite, non-zero switching time, thus limiting the steepness of the ISF and guaranteeing

1. While we use here a similar notation as Collins did, please be aware that g_i and G_n have a subtle but important difference. While Collins uses the slope gain of his circuits, we use here the gain of the amplifier.




Figure 8.2: Comparison of the Fourier transform of the output signal \widehat{V}_{out} and the Fourier transform of the ISF $\widehat{\Xi}$. While \widehat{V}_{out} first decays with $1/f$ up to the frequency of $1/(\pi\tau)$ and then decays with $1/f^2$, the ISF is constant up to the frequency of $1/(\pi\tau_w)$ and decays with $1/f$ afterwards. Generally, $\tau > \tau_w$ can be assumed.

finite and non-zero rise and fall times. Thus the output noise power becomes:

$$N_{\text{out}}^2 = v_n^2 g_i \sum_{k=1}^{\infty} \widehat{\Xi}[k] \quad (8.3)$$

For most practical circuits τ_w will be smaller than τ , i.e. the corner frequency for the ISF will be larger than the corner frequency for the output signal.

Jitter and the Effect of the Filter

Comparing the Fourier transforms of both the output signal V_{out} and the ISF reveals that the frequency range between $2\pi\nu_0$ and $1/(\pi\tau)$ contributes the most to increasing the slew rate ρ . At frequencies higher than $1/(\pi\tau)$, the contribution of each harmonic is still not negligible², but the decay of the harmonics with $1/k^2$ and thus the contribution to the slew rate with $1/k$ quickly becomes much lower than the contribution of the ISF harmonics, which remain constant up to the frequency of $1/(\pi\tau_w)$. From this can be easily concluded, that an optimal filter should cut off the output signal harmonics and ISF at $1/(\pi\tau)$ for optimal jitter J^2 . This is the frequency domain equivalent to Collins' result that the optimum half-level crossing time k is equal to 1. Thus, the same result for optimal distribution of gain between stages, namely $g_{n-1} = \sqrt{2g_n}$ holds true for this frequency domain analysis as well. It also becomes evident that a higher order filter will immediately give

2. $\sum_{k=1}^{\infty} k \cdot 1/k^2 = \sum_{k=1}^{\infty} 1/k = \infty$

benefits in terms of jitter performance, as it will reduce down-sampling of broadband noise much more than it will reduce the slew rate of the output signal.

There are two things to note, though. First, this last, graphical analysis used the trapezoid signal approximation for V_{out} . While this holds true for circuits where the cumulative, total gain G_n is large, it does not for longer chains. E.g. Collins calculated for a 6 stage chain with a total gain of 10^6 the gains for the first three stages to be 2.3, 2.7, and 3.7 respectively. This results in a total gain at each stage of 2.3, 6.2, and 23.0 respectively. These gains are low enough that Collins' trapezoid approximation produces significant errors, i.e., the noise contribution at low gains is overestimated. To reach an optimal solution, the noise contributions of the low gain stages should be calculated by summing up the spectral lines of the ISF Ξ instead.

Second, it becomes evident, that for optimal performance the harmonic components of the ISF have to be limited. Using the circuit model from above, where the ISF is generated due to the non-linearity of the clamping circuit, it can be easily achieved by placing the filter before the clamping circuit instead of after, thus limiting the harmonic contents of the signal reaching the clamping circuit, which in turn reduces the harmonics of the ISF generated. Obviously, if there is no explicit clamping circuit, but the clamping is part of the amplifier itself (e.g. by using a differential pair driven into saturation), then the filter has to become a part of the amplifier in order to be able to modify the ISF.

As mentioned earlier, in a practical design the offset voltage of each amplifier stage should be closely controlled. While once a high G_n has been reached, all remaining stages will have the same duty cycle and thus can be controlled using an earlier stage, the first few stages have to be individually controlled in order to limit even order harmonics generation and thus flicker noise amplification in these stages.

8.3 Conclusion

We have derived an optimal formula for the gains of a multi-stage hard limiter, similar to Collins' result. But instead of using a trapezoidal approximation that produces large errors at the important first stages with their low gain, we get exact results using the spectral analysis of the ISF. At the same time we can calculate the effects of using higher order filters on the noise, which is not possible with Collins' method.

PART III

Faults and Fault-tolerance in Clocking Applications

CHAPTER 9

Clock Synchronization and Reliability

After having looked at noise in clocks and how it propagates through electronic systems, we will now have a look at clock synchronization. Specifically, at clock synchronization in electronic systems when high reliability is required.

With electronic systems, and chips in particular, ever getting larger, faults become more and more a problem [Rad+13], especially as a single fault can potentially knock out the whole system. The industry best practice is to separate the subsystems as much as possible, thus preventing faults in one subsystem propagating into another. Unfortunately, the clocking system is still a single point of failure as it relies on a single, central source as reference [Xan09, §2, p. 9]. But if each subsystem would be run using its own, independent clock, then there will be no known phase or frequency relation between the subsystems, which in turn would require some form of synchronization to allow fault free communication [Kin08, §2., p. 13] [Mar81]. Hence some form of clock synchronization is required, but care must be taken, for this clock synchronization scheme not to become a new single point of failure.

In the following chapters we will look at one approach how to achieve fault-tolerant clock synchronization in a system with multiple independent nodes. We will then look into faults caused by metastability sensitive parts of the clock synchronization implementation and how to mitigate its effects in an efficient way.

CHAPTER 10

Fault-tolerant Clock Synchronization

10.1 Introduction

Reliably clocking complex very large scale integration (VLSI) circuits is a highly challenging problem. The traditional approach of using a global clock tree brings a variety of scalability issues: In high-performance designs, minimizing the clock *skew*, i.e., the time difference between the earliest and latest clock transition arrivals at the clock tree leafs, requires advanced buffer insertion, snaking wires, and wire sizing techniques [Fri01; LLC10; Sha+10]. These techniques typically rely on high precision delay models or symmetry assumptions [SC12] of the involved components. In addition, monolithic clock trees make the system dependent on a single clock source and its system-wide distribution. This introduces a single point of failure, entailing that any dependable architecture clocked in this way requires an extremely robust clock tree design; naturally, this aggravates scalability issues even further.

In light of these obstacles, globally asynchronous locally synchronous (GALS) systems [Cha84] offer a paradigm shift away from centralized clocking. Instead, the system is partitioned into multiple clock domains, each featuring its own clock generation and distribution mechanism. Depending on the relation between the clock domains, such systems are called *mesochronous* (same frequency, bounded phase relation), *plesiochronous* (same nominal frequency), or *heterochronous* (else) [TGL07].

While clock generation and distribution for plesiochronous and heterochronous systems are easily realizable by independent sources and distribution layers, these solutions introduce an entire batch of new problems:

This part is the result of close collaboration with Florian Huemer and Christoph Lenzen. It is based on an article published as [KHL16]. The author's focus was on PCB design, VHDL implementation, and measurement of the system.

- Asynchronous communication across clock domains requires the use of synchronizers for each data path, increasing delays and buffer sizes, and thus decreasing the overall throughput.
- Slight differences in clock speeds may result in different rates of data production and processing, potentially causing buffer overflows.
- Introducing communication (handshaking, etc.) to resolve this issue increases design complexity at the application level and shifts the difficulty of providing strong real-time response guarantees to the application designer.

This advocates re-introducing timing guarantees between the different clock domains, i.e., using mesochronous GALS systems, allowing for higher inter-domain communication throughput and slimmer communication circuits [TGL07; SG03], as well as metastability-free communication [PHS09]. However, while such systems may not suffer from throughput penalties and do not rely on a clock tree as the top-level synchronization mechanism, this top-level synchronization is critical to its operation: unless the inter-domain clocking mechanism itself is fault-tolerant, again a single point of failure has been created.

10.2 Using Fault-Tolerant Algorithms for Clock Synchronization

One way to solve the issue of single point of failures in clocking systems is to use multiple, independent clock sources and then synchronize these with distributed clock algorithms in a fault-tolerant way.

A canonical approach to distributed clock synchronization is to let nodes agree on an approximate common notion of time periodically, and readjusting their local clocks to the value upon which they agreed. Early work on reaching approximate agreement in distributed systems [Dol+86] lead to fault-tolerant clock synchronization algorithms based on this method, see e.g. [Sch87] for an overview.

In this work, we make use of the algorithm by Welch and Lynch [WL88], for the following reasons:

- (1) Its skew is proportional to the delay uncertainty rather than the maximum delay as, e.g., the algorithm proposed in [ST87].
- (2) In contrast to more involved algorithms that require multi-round communication for a single resynchronization [Sch87], it is well suited for an on-chip VLSI implementation.

We modify the algorithm by Lynch and Welch [WL88] to broadcast simple 0-1-0 clock pulses, as opposed to nodes communicating round numbers. In large parts, this modification is inconsequential. However, the algorithm in [WL88] makes use of the round numbers to achieve a powerful recovery property: any node can resynchronize after a transient fault, provided that out of the n nodes never more than $f := \lfloor (n-1)/3 \rfloor$ are faulty or out of sync.

System Model

The system consists of a set V of n nodes that are fully connected by (1-bit) broadcast channels. Each node $v \in V$ is a fault-containment region in the sense defined by Kopetz [Kop03]: a single (physical) fault, such as a gate malfunction, does not directly affect correctness of the components outside the fault-containment region that contains the faulty component. Node v comprises a local physical clock H_v (e.g. a ring oscillator), the circuitry implementing the algorithm's logic for v , and its outgoing links. Note that this means that communication does not use a shared bus, which would be a single point of failure. Any potential application logic clocked by v will be part of its fault containment region as well. Thus, any transient or permanent faults of components (in the fault-containment region) of v affect other nodes only indirectly via communication. A faulty node (i.e., one whose containment region contains faulty components) can behave arbitrarily; in particular, it may send a clock pulse to a subset of the nodes only. We assume that at most $f = \lfloor (n-1)/3 \rfloor$ nodes are faulty, and refer to the set of correct nodes as $C \subseteq V$.

Nodes in C communicate by broadcasts. If $v \in C$ broadcasts at time t_v , any other correct node $w \in C$ has received and processed the respective pulse at some time $t_{wv} \in [t_v + d - U, t_v + d]$, where d is the *maximum delay* and U is the *delay uncertainty*. For faulty senders in $V \setminus C$, such restrictions are irrelevant, as they may deviate from the protocol in an arbitrary way, i.e., send pulses at arbitrary times and independently to different receivers.

A correct node measures the time of arrival of other nodes' pulses relative to the time of arrival of its own pulse of the same round (cf. Algorithm 1). This is done by looping the broadcast signal back and using time-to-digital converters (TDCs) to determine the respective time difference. (cf. Section 10.3). We assume a one-sided¹ worst-case measurement error of our TDCs when comparing signals arriving at times t and t' that fulfills

$$e(|t - t'|) = G + \nu|t - t'|, \quad (10.1)$$

1. By this we mean that we specify the length of the interval around the true value the measurements may come from.

where G is the *granularity* of the time to digital converter (TDC) (i.e., its discretization error) and $\nu \ll 1$ is the maximum relative deviation of the frequency of the TDC's time reference from its nominal frequency.

A node v has no access to real-time, but only to its local clock $H_v: \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$, where $H_v(t)$ is the local clock value at real-time t . For the purpose of a straightforward presentation of the algorithm, we assume that

$$\forall t, t' \in \mathbb{R}_0^+, t > t': t - t' \leq H_v(t) - H_v(t') \leq \vartheta(t - t'), \quad (10.2)$$

2. Naturally, in practice H_v will be discrete and bounded. However, H_v is merely used to control the local logic of the algorithm, rendering this inconsequential to our considerations.

where $\vartheta > 1$ is a constant close to 1.² For the sake of simplicity, we set $\nu = \vartheta - 1$ in the following, i.e., the clock source of a node and its TDCs have the same worst-case phase drift. We assume that $H_v(0) \in [0, F)$ for all $v \in C$, where F is determined by the precision of the booting process. For better readability, we denote real-times with t and local times with τ , with respective indices.

Basic Algorithm

Algorithm 1 gives the pseudocode of the algorithm. Each node $v \in V$ starts round $r \in \mathbb{N}$ at time $t_v(r - 1)$, where $t_v(0) = F$, and ends round r at $t_v(r)$. To fully specify the algorithm, we need to determine τ_1, τ_2 and

Algorithm 1: Synchronization algorithm, code for node v

```

1 //  $H_w(0) \in [0, F)$  for all  $w \in V$ 
2 wait until time  $t_v(0)$  with  $H_v(t_v(0)) = F$ ;
3 foreach round  $r \in \mathbb{N}$  do
4   start listening for messages;
5   wait for  $\tau_1$  local time;           // all nodes are in round  $r$ 
6   broadcast clock pulse to all nodes (including self);
7   wait for  $\tau_2$  local time;         // all messages arrived
8   foreach node  $w \in V$  do
9      $\tau_{vw} := H_v(t_{vw})$ , with reception time  $t_{vw}$  of first message
       from  $w$  ( $\tau_{vw} := \infty$  if no message received from  $w$ );
10   $T_v := \{\tau_{vw} - \tau_{vw} \mid w \in V\}$  (as multiset);
11  let  $T_v^k$  denote the  $k^{\text{th}}$  smallest element of  $T_v$ ;
12   $\delta_v \leftarrow \frac{T_v^{f+1} + T_v^{n-f}}{2}$ ;           // clock correction
13  wait until time  $t_v(r)$  with
        $H_v(t_v(r)) = H_v(t_v(r - 1)) + T_R - \delta_v$ ;           // round ends

```

T_R . The following conditions are sufficient for the algorithm to work

as intended.

$$\begin{aligned}\tau_1 &\geq \vartheta F \\ \tau_2 &\geq \vartheta(F + \tau_1 + d) \\ T_R &\geq \vartheta(\tau_1 + F + U) + \tau_2 + t_{\text{comp}} + G,\end{aligned}$$

where t_{comp} is the time required to compute and apply the phase correction. It is desirable to keep the round length T_R small, unless one seeks to lower the communication frequency. Since any values satisfying these inequalities are acceptable, one may always round up to the next integer multiple of the cycle time of the oscillators controlling the logic, i.e., no constraints on oscillator frequencies are needed. It can be shown that the minimal feasible choices result in a steady-state skew of $E \approx 4(U + G)$ for $\vartheta - 1 \ll 1$. More detailed calculations show that the algorithm can handle frequency offsets of up to $\vartheta - 1 \approx 1\%$ without dramatic impact on E .

Node Recovery

So far we assumed that nodes are initially synchronized and maintain this property. We now address the case that $n - f$ nodes are synchronized, but an additional node is out-of-sync (possibly after a transient fault) and attempts to resynchronize. The modification to the algorithm is extremely simple: whenever a node receives fewer than $n - f$ signals while listening for them in a given round, it will cut this round short. Thus, it quickly catches up with the main field.

Complete proofs for synchronization and node recovery can be found in [KL19].

10.3 Implementation and Experiments

We implemented the algorithm on four Cyclone IV FPGA development boards. We designed a simple additional board to carry the clock oscillator for the field-programmable gate array (FPGA) and the connectors for the coaxial cables between the nodes. In order to allow corrections of the pulse position with sub-clock cycle granularity, we apply phase shifts using a voltage controlled crystal oscillator (VCXO), which supplies the reference frequency for the phase-locked loop (PLL) within the FPGA. The nodes are connected to each other using coaxial cables of the same length (ca. 30 cm), one for each pair of nodes and direction. The FPGA implements four TDCs (see below) to measure the timing of the incoming pulses, implements the logic of the algorithm, and controls the VCXO. An additional pulse output is available for measurements.

Due to limitations of the development board, pulses use 3.3 V LVCMOS signaling. The resulting reflections slightly add to the measurement uncertainties. Furthermore, the FPGA development board only provides two pins for ground connection. This resulted in an involuntary test of the algorithm’s fault-tolerance properties: having many high-speed signals over the same connector, the setup suffered from significant ground bounce of up to 200 mV between the ground potentials of the development board and the interface board; this caused one of the nodes to lose several clock ticks during our experiments.

Cycle Structure and Phase Control

We clock the FPGA with 130 MHz derived from a 20 MHz VCXO on our interface board. As discussed above, to achieve sub-cycle length (i.e. smaller than 7.7 ns) corrections of the phase of the pulse, we control the reference oscillator’s output frequency. We implemented this using a 16-bit, 1 Msps digital to analog converter (DAC) with serial peripheral interface (SPI) interface. This design choice imposed two important restrictions on our implementation. First, the oscillator’s modulation bandwidth of about 10 kHz imposes a lower bound on the round length, as we need to allow for sufficient time for the oscillator to respond to a changed control input. Therefore, we chose a fairly large round length of $T_R = 50 \mu\text{s}$, of which 40 μs are allocated for shifting the clock phase.

Second, the tuning range of the oscillator is roughly 10 ppm limiting the phase correction per round to $\approx 400 \text{ ps}$. This is smaller than the duration of clock cycle of the FPGA ($\approx 7.7 \text{ ns}$), preventing a simple implementation of larger phase shifts by enabling to adjust the (integral) number of clock cycles per round. Fortunately, the convergence analysis shows that the algorithm achieves the same steady-state error with this limitation on phase corrections. However, the number of rounds required for recovering nodes to resynchronize is much larger; with a frequency correction of at most 10 ppm, this takes up to about 10^5 rounds, yielding a worst-case bound on the time to recover in the order of seconds.

Time to digital converter

We employ standard techniques in our TDC design (like in, e.g., [SAL06]). Our implementation is based on the tdc-core by the White Rabbit Project [Bou19] (details are presented in [Bou13] and [Bou11]), which we adapted with minimal changes to Cyclone IV. The TDC uses an adder carry chain as delay line and a coarse counter for the




Figure 10.1: Histogram of the encoded TDL output values during offline calibration. These values correspond to the bin sizes of the delay line.

measurements. Additionally, there is a ring oscillator to measure and compensate for voltage and temperature effects during operation (see [Bou11] for details). We used the internal startup calibration system to get an estimate on the bin size of the TDC and thus its precision (see Fig. 10.1). The largest observed bin size is 140 ps. Estimating a calibration error of up to 20 ps, this yields a single-shot precision of $G \leq 160$ ps.

Parameter Extraction

The performance-critical parameters from the setup are:

- As discussed above, we have $G \leq 160$ ps for the TDC.
- We calibrated the differences in wire delays on the development and interface boards using the TDCs. This results in an uncertainty of $U \leq G + 40$ ps ≤ 200 ps, where 40 ps is an estimated upper bound on the delay variations in equivalent paths between the TDCs.
- We measured a frequency deviation between one pair of oscillators of 1.5 ppm. The manufacturer lists a typical frequency deviation including initial deviation and over temperature range of typical 3 ppm, i.e., $\vartheta - 1 \approx 3 \times 10^{-6}$.

Inserting these values into the bound obtained from the analysis, the estimated worst-case clock skew without faults is $2(G + U) + (\vartheta - 1)T_R = 870$ ps, where $T_R = 50$ μ s is the nominal duration of a round. With faults, this becomes $4(G + U) + 2(\vartheta - 1)T_R = 1740$ ps.




Figure 10.2: Experimental setup. Left: measurement instruments. Center: nodes with FPGA and interface boards, and a stray mouse. Right: recording PC.

Experimental Setup and Results

We fully connected the nodes using cables of length 30 cm. The physical diameter of the whole setup is approximately 50 cm, cf. Fig. 10.2. Measurements are taken by a WaveCrest DTS-2075, which has a single-shot accuracy of ± 25 ps and calibrated the input port skew to achieve better than 2 ps accuracy. To rule out any spurious effects from the instrument, we used two Stanford Research SR620 to verify these bounds.

Skew Measurements

We measured the skew between all pairs of nodes sequentially for at least one hour each, which corresponds to 7.2×10^7 rounds. It turned out that one of the nodes suffers from lost clock cycles due to the aforementioned ground bouncing. In the course of an hour, 34 jumps of over 1 ns were recorded. These jumps have no noticeable effect on the other nodes. We observed a maximum clock skew of 180 ps between correctly working nodes. Figure 10.4 showcases the short-term behavior of the clock skew.

To test the behavior under worst-case faults, we modified one node with the aim to maximize the skew of the remaining nodes. The analysis indicates that the maximum impact of faults is achieved when faulty nodes send early pulses to nodes that are already ahead and none to those that lag behind. After implementing this behavior, we observed an increase in the maximum skew to 270 ps.




Figure 10.3: Long-term evolution of the clock skew of three nodes against the same reference node over a period of an hour, measured sequentially. The thick lines depict the average clock skew over 10s, the light yellow colored fill with the thin lines depict the minimum and maximum in the same interval.




Figure 10.4: Short-term behavior of the clock skew of the "blue" node vs. the reference node from Figure 10.3 over an arbitrarily selected period of 20 s.

Resynchronization

To verify that nodes resynchronize after transient fault conditions, we modified one node to drop out using a manually actuated switch. Triggering the switch every couple of seconds results in randomly distributed restarting times with respect to the clock phase of the correctly synchronized nodes. In 20 measurements, we observed the expected stabilization behavior. In accordance with our earlier discussion, recovery took up to 7 s for our implementation.




Figure 10.5: TDEV between all pairs of nodes, measured sequentially. The colors of pairs match those from Figure 10.3.

Time and Frequency Stability

We analyzed the statistical time and frequency stability (cf. [Sul+90]) of the system in long term measurements. The time deviation (TDEV) plots (Fig. 10.5) are measured between pairs of nodes of the synchronized system. As it can be seen, the noise of the system behaves mostly like white phase noise up to a τ of approximately 10 s.

The results significantly exceed our expectations in the range below 10 s. While the algorithm inherently suppresses effects from outliers, as it drops the largest and smallest measurement value in each round, and subsequently averages between the remaining two, this merely suggests improvements of factor 3 to 5 over a free-running oscillator (TDEV of $\sim 1 \times 10^{-9}$ s @ $\tau = 1$ s). In contrast, uncertainties of parts in 10^{-12} s are already reached above 1 s for the correctly working nodes. These are quite astonishing stability values, especially in light of the crude setup resulting from the employed affordably priced hardware.

As the primary application of the clock synchronization system is to serve as a clock source for circuits, we also analyzed the absolute frequency fluctuations against a Stanford Research FS275 rubidium frequency standard. We show two Allan deviation (ADEV) plots, see Figure 10.6. The first compares a free-running node to the rubidium, i.e., the algorithm is deactivated in order to measure the raw performance of the oscillator. The second depicts the behavior of the same node, but now synchronized to the other nodes, via the algorithm. We observe that the long term stability over $\tau > 10$ s is approximately the same. This is expected, as the long-term behavior is dominated by the temperature-induced frequency fluctuation of the used oscillators.




Figure 10.6: ADEV between a node and a rubidium frequency standard. The top (blue) curve compares to the free-running oscillator, with the algorithm disabled. The bottom (red) curve is the same node with the algorithm enabled and the system fully synchronized. The temperature effects beyond $\tau = 100$ s differ because the traces were recorded on subsequent days with different weather conditions and thus different heating patterns of the building.

Below a τ of 1 s, however, the stability of the synchronized system is higher than the one of the free running node. But as simulations with an accurate noise model of the crystal oscillators have not shown such an increase in stability, we attribute this to the experimental setup, most likely to injection locking through ground bounce. I.e., the ground currents running between the boards injection locking the oscillators to each other, thus effectively forming a clock ensemble. The stability increase of more than a factor of 3 below 1×10^{-1} s is beyond the expectation of $\sqrt{n} = 2$ of such an ensemble of though. Unfortunately we have no explanation for this behavior. Neither were we able to verify whether ground bounce was indeed the cause of this increase in short-term stability.

CHAPTER 11

Metastability in Time to Digital Converters

11.1 Introduction

In the clock synchronization system introduced in the previous chapter, there is one component that exhibits faults often and by design: the time to digital converter. Because the TDC gets pulses from multiple source, some of which could be faulty in arbitrary ways and thus those pulses arriving at arbitrary times, it is impossible to avoid metastable upsets completely [Mar81]. Even more so, by the very construction principle of a TDC where latches are used to measure the arrival time of pulses, it is very likely that one of the latches will become metastable.

In the following chapter, we will look into how to design a TDC such that these metastable upsets can be resolved quickly and with minimal overhead.




Figure 11.1: Tapped delay-line TDC. Latches are initially enabled and output 0. The delayed starting signal iteratively sets latches to 1 until the stopping signal disables them.

This part is the result of close collaboration with Matthias Függer, Christoph Lenzen, and Thomas Polzer. It is based on an article published as [Füg+17]. The author's focus was on circuit simulation.

Fine and Coarse TDC

To keep complexity low and to extend the dynamic range, often a fine-TDC/coarse-TDC approach is employed. There exist many variants of this approach and similar strategies have been invented many times independently. To the best of our knowledge, the earliest example of the fine-TDC/coarse-TDC approach is [Nut68].

The Problem of Metastable Upsets

In his seminal work [Mar81], Marino formally proved that *any* circuit with a bistable output can become *metastable*, i.e. that its output might transition to a stable 0 or 1, an arbitrarily long time after the input stimulus has been applied. In the case of TDCs, the latch input might transition exactly when being captured, violating setup/hold constraints, resulting in metastability.

However, in a TDC of the general structure depicted in Fig. 11.1, it is straightforward to restrict the number of metastable output bits to at most one as follows. Assume that the TDC has $n \geq 1$ latches and each delay element has latency $T_d > 0$. Let T be the time the latches are given before their outputs are further processed by a readout circuit or copied to memory. For a fixed execution of a START transition traversing the ring TDC, each latch has a *critical window* of time: if the STOP signal arrives at a time t from this window, a metastable upset of the latch may not resolve by time $t + T$; if t lies outside of this window, it is guaranteed that the latch is stable at time $t + T$. Larger T results in smaller window sizes [Kin08], but longer readout time and thus lower sample rates. Thus, there is a general trade-off between fast readout (high sample rate) and high reliability (low upset probability). Note that routing delay uncertainties/variations can be accounted for in the critical window sizes and positions.

Note that, for given window sizes and bounds on the unbalanced arrival of the stopping signal, one can choose the delay-element latency T_d large enough for the critical windows to be non-overlapping. In this case, the (thermometer encoded) TDC readout is guaranteed to contain at most one metastable bit: it is of the form $11..100..0$ or $11..1M00..0$, i.e. a consecutive series of 1's followed by a consecutive series of 0's with at most one metastable bit/latch M in between. No matter how M resolves, the final measurement value is off by at most one from the actual readout; we say the TDC *guarantees precision of 1 (in units of stage delays)*. By Marino's result this is optimal: there is no TDC implementation which guarantees only stable output bits.

We stress that reducing T_d further, such that potentially more than one bit may become metastable, may lead to better average precision, but does *not* improve worst-case precision. In other words, the critical window size is a lower bound on the single-shot precision of the TDC. In this work, we are aiming for optimal single-shot precision, which is relevant to applications that provide or rely on guaranteed precision bounds, e.g., within control loops in mission critical hardware. As in most settings the above lower bound can be matched by delay-line TDCs, extending our techniques to Vernier line TDCs is outside the scope of this work.

Dealing with Metastability

As we have seen, metastable upsets cannot be avoided deterministically. Worse, when increasing the accuracy of the TDC, we increase the risk for metastability, as the relative size of the critical windows becomes larger. At the latest when reaching the limit on TDC precision implied by Marino's results, the effects of metastability on the circuit cannot be neglected anymore. Thus, the threat of metastability must be addressed by the circuit design. This can be done in the following ways.

Resolving in TDC: The TDC readout can be given sufficient time for the metastable bit to stabilize with sufficiently high probability. While this solution is the simplest, it prevents the TDC from taking a new measurement during this time, reducing its sample rate.

Resolving in Memory: The TDC readout can be written to memory; this may result in metastability of a written bit, which then is resolved by (chains of) synchronizers [Kin08] or, again, simply waiting for sufficiently long time. Note that the TDC can “keep running,” in this setting, as the “stopping” signal here is used to capture the current TDC state in memory only. This enables the TDC to be reused for additional measurements while time-stamps resolve metastability in memory.

Resolving during/after Computation: Friedrichs et al. propose a new approach [FFL18]: they introduce the concept of metastability-containing circuits. Instead of waiting for the metastable bits to stabilize, such circuits guarantee a bounded degree of uncertainty in outputs, given that the input comprises a bounded amount of metastable bits. For example, one can immediately begin computing the maximum of potentially metastable measurement values of precision 1, yet ensure output of precision 1 [BLM17;

LM16]. This permits to simultaneously use the time to resolve metastability for computations, as well as using an adaptive amount of time for resolving metastability *after* computations (depending on the time until the output needs to be sampled for further processing).

No Resolution: If all operations are metastability-containing, resolving metastability is not necessary at all. For example, this is the case when the result of the computation is used for analog control; the authors of [BLM17; FFL18; LM16] discuss clock synchronization as an application where this is feasible and of interest.

We emphasize that existing TDC designs have severe limitations when it comes to the last two options:

- Delay-line TDCs without coarse counters ensure that only a single bit is metastable (under the same constraints as our designs). However, such designs are practical for a very limited dynamic range due to the inefficient encoding and require subsequent conversion to Gray code for (efficient) follow-up computations.
- Most current counter-based designs directly protect their (binary) counters from metastable upsets by synchronizers, incurring a delay of several clock cycles (i.e., typically nanoseconds) before the measurement becomes available. This also necessitates to stop the TDC oscillator, disallowing multiple measurements pertaining to the same starting signal.
- An exception is the design by Mota et al. [Mot+00]. It employs two coarse counters (driven by rising and falling clock edges, respectively), one of which is guaranteed to not be metastable upon sampling the TDC state with the stop signal (i.e., latching the corresponding registers). This enables waiting for stabilization in memory. However, reading the memory word representing the TDC value correctly requires to determine which counter's sampled value to use, which in turn requires prior stabilization of the thermometer encoding of the sampled fine-TDC value. Apart from stopping the entire TDC, this requires an additional dead time of at least 1 ns to reach a moderate mean time between failures (MTBF) of 8.8 years per thermometer bit (cf. Section 11.3).

Therefore, the techniques presented here are a big step towards an efficient implementation of the clock synchronization scheme proposed by Friedrichs et al. [FFL18].

Memory Efficiency

Achieving precision 1 and the possibility to resolve metastability *outside* the TDC (and thus taking multiple measurements concurrently) are prime design goals. Unfortunately, as mentioned above, they come at the cost of a very large memory footprint for tapped delay-line TDCs: The time resolution of the TDC is given by the delay element latency T_d , and its measurement range is nT_d . Resolving metastability in memory does not alleviate this problem, as the memory consumption for storing a single measurement means that we could just add another delay-line TDC at essentially the same cost!

Note that naively packing the thermometer code measurement values into binary coded values *before* metastability is resolved may result in arbitrarily bad precision (i.e., has arbitrarily large errors). For an example illustrating the problem, assume that 111111M (with M denoting the metastable bit) is written to a FIFO with thermometer encoded stages. After metastability resolution, we end up with either 1111110 or 1111111 — decimal 7 or 8. The binary representation of these numbers are 0111 and 1000, respectively. Note that these strings differ in all bits. As shown in [FFL18], *any* circuit computing the binary representation of unary inputs may thus return MMMM when given input 111111M. This is to be seen as a discrete version of Marino’s impossibility result [Mar81]. The string may thus resolve to any 4-bit string, i.e., the binary representation of *any* number between 0 and 15 — losing all TDC precision at the encoding step! Although the probability for such upsets can be made small (but not zero!) by synchronizer chains, this is at the cost of increased TDC latency and synchronizer chains for *all* measurement bits. Portmann and Meng studied the same issue as it arises for flash ADCs in [PM96].

To reduce both the memory footprint and the size of the TDCs, both for tapped delay-lines [Hen10] and Vernier lines [DSH00] solutions have been proposed where the delay line has been folded into a ring oscillator-like structure with an additional (binary) coarse counter to count the number of rounds the start signal makes in the ring oscillator; see Fig. 11.2. The stopping signal then freezes both the latches and the counter.

However, these designs suffer from potential metastable upsets when incrementing the counter. For the reason laid out above, even waiting for the counter to stabilize after possible metastability does not fix the problem: in case the counter is binary, the resulting TDC output value has arbitrarily bad precision. Note that this analysis does not include the timing violations due to different path lengths for different bits within the adder structure, but just the metastable upsets within




Figure 11.2: Generic ring TDC architecture with coarse counter C . Latches and counter are initialized to 0. The counter increments on rising and falling transitions.

the output register. To address this, one can employ synchronizers more carefully [BG15; Kin08; Vee80], decreasing the probability of a metastable upset of the counter arbitrarily. This trades time (for the synchronizer steps) for increased reliability.

Even if this is acceptable, it requires to stop the counter for taking a measurement, restricting the TDC to taking a single measurement *and* waiting for the synchronizers to stabilize before starting a new one. While the former restriction can be lifted by duplicating the counter logic and synchronizers for each concurrent measurement, this basically amounts to duplicating the entire TDC for each concurrent measurement.

11.2 Binary Cycle Counter Designs

When using a binary counter to keep track of the number of rounds/cycles in the ring oscillator between arrival of `START` and `STOP`, the ring is used as clock signal for the counter, and the `STOP` signal is used as a gate signal (connected to the enable input of the counter's D-flip-flops (D-FFs)). Such an implementation is used quite frequently in synchronous designs. The problem in the case of TDCs is that the `STOP` is outside of the timing closure of the counter and may occur at an arbitrary point in time. To mitigate this problem, synchronizers can be used to align the `STOP` signal with the counter clock. However, this naive solution either dramatically reduces the precision of the TDC, as an additional delay is added to the `STOP` signal that varies depending on the phase offset between the clocking and `STOP` signals, or does little to improve safety, as too little time is allocated to reliably resolve metastability. More involved designs use two fine TDCs, one for the `START` and one for the `STOP` signal (e.g. [Nut68]) such that the counter can be stopped synchronously to its clock, but this entails having two fine TDCs and potentially duplicating the entire counter logic to allow for multisampling (due to the additional delay in the `STOP` signal caused by the synchronizers).




Figure 11.3: Transient simulation of the capture register of the simulated TDC. Even very small variations of the applied input signal increase the time until the output value reaches its final voltage level considerably. Note that a full D-FF was used instead of a single latch.




Figure 11.4: Stabilization time variation due to metastable upset of the simulated TDC capture register. time between application of the input and the clock edge vs. output stabilization time.

To demonstrate that the probability of metastable upsets cannot be neglected, consider a state-of-the-art ring TDC with binary counter as implemented by [PC12; PC13]. Let us replace a ring latch in Fig. 11.2 with a D-FF whose data input is driven by the respective ring segment and that is clocked by the STOP signal. For the above setup assume input transitions to occur with a rate of $f_d = 500$ MHz and a STOP transition occurring every 20 ns (an input rate of $f_c \geq 50$ MHz is typical for high-speed clock synchronization systems). For the mean time between metastable upsets in the synchronous circuit, we obtain [Kin08] $T_{\text{MTBF}} = e^{T_{\text{res}}/\tau} / (f_c f_d T_w)$, with the remaining parameters being determined as follows. To characterize metastability of the ring state flip-flop, we implemented part of the ring (6 ring inverters) as described in [PC12; PC13], driving the data port of the D-FF, and shaped the clock signal by 6 upfront standard inverters. Cell design of ring inverters and D-FF, including transistor sizing were taken from [PC12], for a 130 nm technology. We then ran Spice simulations (Fig. 11.3,

Fig. 11.4) for the above circuit showing a settling time τ for the complete chain from START and STOP to the D-FF outputs of 31.6 ps and a critical window size T_w of 8 ps. Assuming that metastability has $T_{\text{res}} = 1$ ns time to resolve before the measured value is used, we obtain an T_{MTBF} of 8.8 years *per bin and measurement channel* (i.e. assuming a TDC with 64 bins and 8 channels the MTBF reduces to roughly 6 days). Note that the situation is even worse for the coarse counter T_{MTBF} , as it comprises several flip-flops and its increment logic is not perfectly balanced. This shows that using synchronizer chains is inevitable for such designs, leading to large dead times of the TDC. In the following section, we present three alternative solutions.

11.3 Metastability-Aware time to digital converters

In this section, we present our generic approach for devising ring TDCs that do not suffer from reduced precision due to metastable upsets. We showcase our ideas at hand of three alternative designs, each of which allows for taking multiple measurements concurrently (without duplicating the TDC logic) and achieves optimal single-shot precision. Our designs differ in terms of the trade-offs between routing constraints and memory overhead; the output of one satisfies the input specification of the metastability-containing sorting networks from [BLM17; LM16], without the need for further conversion or waiting for stabilization.

For the sake of presentation, we assume throughout this section a simultaneous arrival of the stopping signal at all relevant components, following the tapped delay-line approach. In practice, one may choose other options and compensate for them or switch to a Vernier line approach; these considerations are outside the scope of this work. Moreover, we discuss our designs in the context of a measurement with a single start signal followed by a single stop signal, where the measurement value is captured in the TDC.

The Ring

Consider a ring TDC architecture of length n as depicted in Fig. 11.2. We denote the output of latch i , $0 \leq i \leq n - 2$, at arrival of the STOP signal by $c(i)$. At this point we assume latches with zero switching time. We will discuss this issue shortly. For the moment, assume that the coarse counter C is a (single-bit, binary) latch. We denote its output by $c(n - 1)$.

We denote by cnt the sum of the total number of stage transitions of all stages between the starting and stopping signals. Then either the




Figure 11.5: Solution I: Implementing C by two binary counters and a latch. Counters C_0 and C_1 and latch L are initialized to 0. The increment inputs of C_0 and C_1 are driven via a delay buffer. C_0 increments on rising, C_1 on falling transitions.

ring oscillator made a number of full rounds and all latch outputs are the same, or its state was captured in between and the latch outputs are of the form 00..011..1 or 11..100..0;

This abstracts away the issue that the latches cannot switch from 0 to 1 or vice versa in zero time. Disabling them during their critical window may therefore result in metastability. However, as laid out in Section 11.1, choosing the delay latency T_d large enough to separate the latches' critical windows guarantees that at most one bit is metastable. This, plus the fact that the ring stores the measurement value thermometer encoded, guarantees precision of 1. In other words, if we can assume that the critical windows of the latches are non-overlapping, then, after any potential metastability has been resolved, $\text{cnt} \bmod n$ can be resolved with precision 1.

The Coarse Counter

We now address the issue of counting the number of cycles. Doing this in a naive way incurs a loss in precision of the TDC, as has been shown in Section 11.3. We present three coarse counter variants which do not suffer from this problem, i.e., ensure an optimal precision of 1.

Solution I: Redundant Binary Counters

Consider the coarse counter implementation in Fig. 11.5 (as described in [Mot+00]). The circuit's underlying idea is to use two redundant binary counters C_0 and C_1 and the output b of latch L capturing their common input. Initially, $b = 0$, $C_0 = 0$, and $C_1 = 0$. When the stopping signal arrives, the binary counters and latch L are disabled. Output b serves as a control bit: we use the value stored in C_0 if $b = 0$ and the value in C_1 if $b = 1$. This ensures that the counter b is pointing to had a full cycle of the ring oscillator to stabilize.




Figure 11.6: Critical windows for Solution I as the START signal travels through the ring oscillator. First the first stage “handles” the START signal during the stage delay time T_d and can become metastable during its critical window W_0 , then the second stage during its window W_1 and so on. When the START signal traveled once around the the ring oscillator it will increment the first counter C_0 after the buffer delay T_{buf} and thus cause it to enter its critical window W_{C_0} , while the signal continues to progress through the ring oscillator. After the second round, the second counter C_1 is incremented and can potentially become metastable during its critical window W_{C_1} .

In order for this approach to work, we require the following constraints to hold: The critical windows W_{C_0} and W_{C_1} of counters C_0 and C_1 , as well as the critical window W_L of latch L are mutually disjoint and obey the following order in time: $W_L, W_{C_0}, W_L, W_{C_1}, W_L, W_{C_0}$, and so on. Assuming the same designs for both counters (i.e., $|W_{C_0}| = |W_{C_1}|$), this can be achieved by the following design constraints:

- (1) W_L before W_{C_0} : choose the delay latency T_{buf} sufficiently large.
- (2) W_{C_0} before next W_L : choose the ring size n sufficiently large for the binary counter to complete an increment before the transition traverses the ring.

Fig. 11.6 depicts the alignment of the critical windows with the above constraints fulfilled.

When reading the counter, the number of completed cycles cyc is computed as $2C_b + b$. Since $2C_b + b$ is equal to the value of the concatenated binary counter (C_b, b) , we may view bit b as a shared bit of the binary counters $(C_0, 0)$ and $(C_1, 1)$ that hold the value cyc in case of $b = 0$ and $b = 1$, respectively. This allows for an efficient method for the TDC to directly return the binary encoding of cyc .

The complete TDC is given by combining the methods of counting modulo n discussed in Section 11.3 and determining cyc given in Section 11.3. I.e., After metastability of the latches has been resolved $cnt = (2C_b + b)n + x_{1-b}$, where x_{1-b} is the number of latches having value $1 - b$. Moreover, it is exactly the leading x_{1-b} latches that have value $1 - b$.

Recall from Section 11.3 that waiting for 1 ns before consulting the latch states resulted in a moderate MTBF of 8.8 years per bit and

channel. Further note that, if the binary counters have B bits, we store $2B + n$ bits for a measurement, but the maximum cnt value $2^B n - 1$ (in principle) requires $B + \lceil \log n \rceil$ bits only. This is of concern when taking multiple measurements concurrently. The next sections discuss two TDC variants that address this issue.

Solution II: Gray Code Counter

In this section, we present a circuit that can be used to read the TDC in a way requiring to store optimal $B + \lceil \log n \rceil$ bits per measurement only. As a starting point, observe that we know that no circuit can avoid metastability of an output bit if, given all other inputs, the output value depends on an input bit that may be metastable [FFL18; Mar81]. Therefore, the first key insight is that we must use an encoding guaranteeing that metastability of *any* latch must not cause metastability of more than one output bit: otherwise, for any encoding without redundancy, we must lose precision, as two or more bits resolving to arbitrary values can induce a change of the encoded value larger than 1. In other words, we must use a Gray code, for which any up-count changes exactly one bit.

Look-ahead Gray Code Counter Consider the ring architecture from Fig. 11.2 and replace the coarse counter by a single *Gray code counter with look-ahead* that increments both on rising and falling transitions. When the counter is triggered, it is required (i) that it internally prepares an up-count within a single oscillator cycle, and (ii) that, once ready for an increment and it is triggered, its outputs react fast, such that their critical window is short and does not overlap with any other critical windows of the ring latches; thus the term look-ahead. Fig. 11.7 shows the resulting alignment of critical windows. Note that such a counter can be implemented by a Gray code counter that counts “ahead” by one (performing each increment within an oscillator cycle) and whose outputs are captured upon an “actual” increment. Fig. 11.8 shows such an implementation: the rising edge triggered counter is initialized with (the look-ahead) value 1, and the rising edge triggered register with 0. The XOR transforms falling and rising transitions on inc into pulses, triggering (a) the register to capture the actual (pre-computed) counter value, and (b) the counter to start the next increment. The buffer delay T_{del} makes sure that (a) always happens before (b). The counter now has a complete ring oscillator cycle to finish its computation.

A TDC readout thus comprises the register states (in Gray code) and the ring latches (thermometer encoded).




Figure 11.7: Critical windows for Solution II.




Figure 11.8: Solution II: Implementation with look-ahead Gray code counter and register. The counter C is initialized with 1. Register R_C is of same width as the counter, and initialized with 0.

We next discuss how to interpret a stabilized TDC readout, e.g., assume it was stored in an external memory with sufficient time to stabilize, and now is read by an application.

We determine whether cyc is odd or even from the coarse counter value, i.e., we compute $b = cyc \bmod 2$. From b , we can infer how to correctly interpret the value stored for $cnt \bmod n$, i.e., whether thermometer encoding of cnt is of the form 11..100..0 (even) or 00..011..1 (odd); making us count the leading 1s or 0s. What if the Gray code counter was metastable, and thus b was metastable, too? As before, this does not matter, as the resolution of the counter (and thus b) will induce a difference of 1 in the value of cnt at most (the counter has been triggered or has not).

However, this approach requires us to store TDC readouts as tuples of Gray code and thermometer codes, until they are guaranteed to have stabilized.

We discuss a more efficient encoding in the next section. Observe that this is non-trivial, because the interpretation of the ring latch states depends on the Gray code counter’s value.

Encoding the Ring Latch States Also here, we need to make use of a Gray code for safely “compressing” the $(n - 1)$ -bit thermometer encoding into a $\lceil \log n \rceil$ -bit string. It can be shown that given an arbitrary Gray code, e.g., the binary reflected Gray code (BRGC), there is a circuit that efficiently (i.e., in $\mathcal{O}(\log n)$ -depth) translates the thermometer encoded ring state to this Gray code, inducing metastability of at most one bit that is the currently least significant bit.

0	1	2	3	4	5	6	0	1	2	#1s	#0s
0	0	0	0	0	0	0	0	0	0	0	7
1	0	0	0	0	0	0	0	0	1	1	
1	1	0	0	0	0	0	0	1	1	2	
1	1	1	0	0	0	0	0	1	0	3	
1	1	1	1	0	0	0	1	1	0	4	
1	1	1	1	1	0	0	1	1	1	5	
1	1	1	1	1	1	0	1	0	1	6	
1	1	1	1	1	1	1	1	0	0	7	0
0	1	1	1	1	1	1	1	0	1		1
0	0	1	1	1	1	1	1	1	1		2
0	0	0	1	1	1	1	1	1	0		3
0	0	0	0	1	1	1	0	1	0		4
0	0	0	0	0	1	1	0	1	1		5
0	0	0	0	0	0	1	0	0	1		6

Table 11.1: Encoding the ring latch states in case $n = 8$: original thermometer encoded ($n - 1$ bits $0 \dots 6$), efficient BRGC encoded ($\log_2(n)$ bits $0 \dots 2$), and the relevant decimal counts of 1s and 0s.

As an additional benefit of this approach, note that the same *XOR-tree circuit* can be used to encode the number of 0s in n -bit strings $00..011..1$, for the following reason: Switching between encodings $11..100..0$ and $00..011..1$ is just by negation of the input. Propagating the negations from all inputs through the XOR-tree to the outputs yields that it suffices to negate a fixed subset of the output bits to obtain the Gray code for the complemented input thermometer encoding. In the case of BRGC, only the left-most BRGC bit has to be negated:




Figure 11.9: XOR-tree circuit implementation counting the number of 1s as specified in Table 11.1. Note that a metastable ring latch influences only one, the current least significant, bit (compare lines 4 and 5 in Table 11.1): no precision is lost when encoding metastable data.

e.g, see Table 11.1: counting 1s in input 1000000 yields output 001, counting 0s in 0111111 yields 101.

In particular, one can “delay” the application of the respective output negations until after potential metastability of the counter (and thus how the $\lceil \log n \rceil$ -bit string is to be interpreted) has been resolved.

In summary, we obtain a memory-efficient metastability-containing TDC by using a B -bit Gray code coarse counter and the XOR-tree circuit from above, a measurement can be stored as a tuple of two Gray codes, using $B + \lceil \log n \rceil$ bits, without losing precision.

Metastability-containing Computations Above, we presented a way to efficiently store a, potentially metastable, TDC measurement as a Gray code tuple. While this is memory-optimal, the question arises if we can store a TDC measurement as a *single* Gray code value that does not need any further transformation when being read by an application.

For operations which are metastability-containing, meaning performing the computations *before* metastability has been resolved, having a single BRGC measurement value is of specific interest as there already exist circuits requiring BRGC inputs [BLM17; FFL18; LM16].

Providing such inputs with our approach is not only possible, but straightforward. For n being a power of 2, a single BRGC encoded value of the measurement value can be generated in a very convenient way. If we use a BRGC coarse counter and n is a power of 2, then *just* the concatenation of (a) the output of the above XOR-tree circuit, *without the need to negate* any of its output bits, and (b) the output of the BRGC counter yields a BRGC encoding of the TDC measurement value. Fig. 11.10 depicts the resulting circuit in case of an even coarse counter value. Fig. 11.11 demonstrates that the same circuit works for odd and for metastable coarse counter values.

We stress the utility of this insight. A naive implementation would require to build an adder with inputs (a) BRGC encoded cyc and (b) thermometer encoded latch outputs, which is a complex circuit. Additionally, this requires both inputs to be metastability-free or to design a metastability-containing adder. By contrast, concatenation is trivially metastability-containing: no precision is lost due to possible metastability.

Solution III: Gray Code Counter with Latch

The Gray code counter variant of the previous section requires that the output of the counter is stable at all times outside a time window of size similar to the critical window of the latches (limited by T_d , cf. Fig. 11.7). For high accuracy TDCs with small T_d , this (a) ei-




Figure 11.10: Circuit encoding a TDC measurement value ($n - 1$ thermometer encoded ring latch states and B bit coarse BRGC counter state) as a single BRGC encoded value without losing precision in presence of a metastable input bit. Here $n = 8$ and $B = 2$ and the coarse counter is even (decimal 0), i.e., we have to count 1s in latch states. Observe that the final value 0011M correctly encodes decimal 4 or 5, depending on how M resolves.




Figure 11.11: Left: coarse counter is odd (decimal 1), i.e., we have to count 0s in latch states. Observe that the final value 0101M correctly encodes decimal 12 or 13, depending on how M resolves. Right: coarse counter is metastable. Observe that the final value 0M100 correctly encodes decimal 7 or 8, depending on how M resolves.

ther imposes a harsh constraint on the design of the counter, which is potentially difficult to meet, or (b) requires an implementation with additional registers as a workaround (see Fig. 11.8).

For these cases, we propose a different metastability-aware design depicted in Fig. 11.12: We add one extra (single-bit) latch L with output b only. The XOR gate delays the counter increment input and additionally transforms rising and falling transitions to pulses, allowing the use of standard single-edge triggered flip-flops for the Gray code counter implementation.

Analogously to the redundant binary counter solution, we only require non-overlapping critical windows of latch L and Gray code counter C ; see Fig. 11.13.

Given a (stabilized) TDC readout, again we determine whether cyc is odd or even – this time not from the counter value, but from the explicitly stored b in latch L . This enables to correctly interpret the value stored for $\text{cnt} \bmod n$ as before. Moreover, it is used to account




Figure 11.12: Solution III: Gray code counter variant with single-bit latch.




Figure 11.13: Critical windows for Solution III.

for an incomplete up-count of the Gray code counter: if the parity of the stored counter value differs from b , we know that the counter *should* have been incremented, but has been prevented from doing so by the stopping signal. In this case, we amend this by performing an up-count on the stored value (w.r.t. the used Gray code). This results in a correct value, because metastability of the counter affects only the (unique) bit that is being changed on the respective up-count.

Naturally, it may also happen that b becomes metastable. However, in this case, the delay T_{del} ensures that the counter did not start the increment corresponding to the (incomplete) transition of L . Thus, either resolution of b results in a correct interpretation of the measurement: If L stabilizes as if it had not been triggered, all is well; if it stabilizes as if it had been triggered, we fix the counter value accordingly.

11.4 Conclusion

In TDC applications that require *high-speed sampling* together with *deterministic guarantees on precision*, one must account for metastable upsets and memory-efficient storage of sampled data. Using Spice simulations, we demonstrated that a state-of-the-art ring TDC with coarse binary counter indeed suffers from significant upset rates. As a solution, we compare and formally prove correct three variants of ring TDCs that provide *high-speed sampling*, *optimal guaranteed precision* of 1 delay unit, and *memory-efficient storage* of TDC values. The proposed variants offer different trade-offs between ease of implementation and memory overhead. We stress that variant II uses an encoding of measurement values with zero memory overhead that can be computed *without* waiting for metastability to resolve. This renders it of particu-

lar interest for use in metastability-containing circuits [BLM17; FFL18; LM16].

CHAPTER 12

Efficient Metastability-Containing Multiplexers

12.1 Introduction

Multiplexers (MUXes) are ubiquitous in digital circuits. A MUX accepts a *select bit* s and input words a, b , and outputs a if $s = 0$ and b otherwise. If $a = b$, the MUX must output a , regardless of s , as s represents the choice between two identical options. Unfortunately, common MUX implementations—see Figures 12.1a, 12.2a, and 12.2b for gate-level, transistor-level, and pass-gate MUXes—violate that specification if s is degraded.

In this chapter, we address this issue by specifying and efficiently implementing a Metastability-Containing Multiplexer (CMUX), and simulating it using Simulation Program with Integrated Circuit Emphasis (SPICE). A CMUX behaves like a regular MUX with one additional property: If $a = b$, the output is a , regardless of how degraded the select bit is.

CMUXes are useful in the presence of degraded inputs in general, but our prime motivation is metastability-containment [FFL18]. While degraded select bits can be avoided in standard digital circuits, this is not the case for metastability-containing circuits, which typically mask possibly metastable bits as follows: If s may be metastable but one can guarantee that this only happens if $a = b$, a CMUX is fed with inputs s, a, b and outputs a , thereby masking metastability of s [FFL18]. A standard MUX fails in this situation, violating timing constraints and yielding an undefined output, thus risking further metastable upsets. We detail on the concept of metastability-containment in the following section.

Specifically in a Lynch-Welch based clock synchronization scheme like presented in Chapter 10, MUXes play a crucial role. After the TDC

This part is the result of close collaboration with Stephan Friedrichs. It is based on an article published as [FK17]. The author's focus was on circuit simulation and circuit optimization.

captured the pulse arrival time, there will be some signals left that are metastable (c.f. Chapter 11). Either time has to be spent to resolve metastability, which can be arbitrary long, depending on the MTBF target, or any further processing has to be done in a way that can contain metastability from degrading the output, or worse “infecting” the whole computation. In Lynch-Welch based systems the first, and most involved computation step after time measurement is a sort operation. It has been shown in [BLM17] that this can be done in such a way that metastability can be contained in an optimal way, i.e. that the output precision is not any worse than the input precision. But they rely on the used MUX to be CMUX.

Most metastability-containing circuits proposed to date critically depend on CMUXes [FFL18; BLM17; Füg+17; LM16]. However, the only implementation of a CMUX previously proposed is an 18-transistor gate-level implementation that has not been subject to simulations [FFL18]. We propose an implementation that reduces the transistor count to 8, matching that of a standard MUX, demonstrate its correctness with SPICE simulations, and hence improve a core component of metastability-containing circuits.

Motivation: Metastability-Containing Circuits

We define a *metastable* [Mar81] signal as having a voltage between logical 0 and 1. Such a signal can be stuck between 0 and 1, oscillate, or show any unspecified behavior. In short, it can obtain any voltage between logical 0 and 1, and perform arbitrary and unclear transitions. Metastability may break the abstraction of Boolean logic in digital circuits. Unfortunately, one inherently takes the risk of metastability whenever communicating signals across unsynchronized clock domains, and when using analog to digital converters (ADCs) or TDCs [Mar81].

As it is well-known that no digital circuit can deterministically avoid, resolve, or detect metastability [Mar81], the established strategy is to (1) store potentially metastable signals in synchronizers [BG15; Gin03] for a predefined time and (2) only afterwards use them in the main circuit. This isolates metastable signals from the main circuit by restoring them to proper levels; the probability of maintaining metastability in a synchronizer decreases exponentially over time [KBY02; Vee80].

Recently, the radically different approach of *metastability-containing circuits* emerged [FFL18; BLM17; Füg+17; LM16]. Instead of avoiding metastability altogether, one can fortify a circuit — using fine-grained logical masking — such that metastability only “infects” a limited part of it. Such techniques are necessary when the delay required for reliable synchronization is unacceptable, e.g., in highly time-critical

applications like high-frequency clock synchronization [FFL18], or in ADCs [Cha+17; GHL12; WMM16] and TDCs [Füg+17], i.e., before synchronizers can be applied. Using metastability-containing techniques, it is possible to completely avoid synchronizers if one accepts a certain degree of metastability in the output [FFL18].

An example are metastability-containing Gray code sorting networks [BLM17; LM16]. Suppose ADCs output Gray code but, due to their analog input, a possibly metastable bit decides whether to output x or $x + 1$. As x and $x + 1$ only differ in a single bit, their metastable superposition has at most one metastable bit; all other bits are stable in an appropriate implementation. The property of all possible stabilizations lying in $\{x, x + 1\}$,¹ is called *precision-1* [FFL18; LM16]. It has been shown that it is possible to sort inputs of precision-1 such that the output still has precision-1 using metastability-containing Gray code comparators [BLM17; LM16].

In addition to Gray code sorting networks, a metastability-containing TDC [Füg+17], several operations on unary numbers and their Gray code conversion [FFL18], and a synchronizer-free clock synchronization scheme [FFL18] have been proposed.

To the best of our knowledge, all metastability-containing circuits proposed to date have one thing in common: At some point, a control bit s that decides whether to output a or b may become metastable. In that case, the circuits are specified such that a and b only differ in a limited number of bits and that the output is a superposition of a and b *that can only become metastable in the differing bits*. While a standard MUX fails for this operation, a CMUX ensures that the identical bits in a and b are stable in the output. Regarding Gray code comparators, s may e.g. be metastable if comparing x to a superposition of x and $x + 1$ [BLM17; LM16].

Related Work

Metastability-containing circuits has been established by Friedrichs et al. [FFL18]. These techniques have been successfully applied to TDCs by Függer et al. [Füg+17], and to Gray code sorting networks by Lenzen and Medina [LM16] and Bund et al. [BLM17] Tarawneh et al. presented the related technique of speculative computations [TY12; TYM14].

Metastability-containing circuits are related to glitch-free/hazard-free circuits, which have been extensively studied since Huffman [Huf57] and Unger [Ung59] introduced them. Both show how to identify and prevent hazards resulting from a single transition of a single input. Eichelberger [Eic65] extended these results to multiple switching inputs

1. This is not the case for binary numbers, e.g., 7 and 8 differ in 4 bits.

and dynamic hazards, Brzozowski and Yoeli extended the simulation algorithm [BY79], Brzozowski et al. surveyed techniques using higher-valued logics [BEI01], and Mendler et al. studied delay requirements needed to achieve consistency with simulated results [MSB12].

A common assumption in hazard detection is that inputs only perform well-defined, clean transitions, i.e., the assumption of a hazard-free input-generating circuitry is made. This is the key difference to metastability-containment: Metastability encompasses much more than inputs that are in the process of switching; metastable signals may or may not be in the process of completing a transition, may be oscillating, and may get “stuck” at an intermediate voltage.

Another common assumption in hazard detection is that circuits have a constant delay. This is no longer the case in the presence metastability. In fact, we demonstrate in Section 12.4 that circuit delays can deteriorate in the presence of metastable input signals, even if the circuit generates a stable output.

Model

Friedrichs et al. model metastability in digital circuits on the gate level [FFL18]. In addition to logical 0 and 1, they introduce M as a third signal value; M is a placeholder for arbitrary voltages between 0 and 1, including 0 and 1 as well as any voltage in between. M includes the output of a metastable latch, as well as the output of any circuit or process that may degrade the signal voltage. Metastability breaks the abstraction of Boolean logic: Under input $x = M$, the output voltage of a circuit implementing $x \vee \neg x$ is undefined, i.e., M . Metastable signals can be contained using *logical masking*; for example, $M \wedge 0 = 0$ and $M \vee 1 = 1$, but $M \wedge 1 = M$ and $M \vee 0 = M$. If one can limit the number of metastable input bits, one can use these masking techniques to limit the number and the position of metastable bits in the output, see Section 12.1.

As our goal is an efficient transistor-level implementation of a CMUX, we transfer this model from the Register-Transfer Level (RTL) to the transistor level. To this end, we examine the behavior of transistors under undefined input voltages. A transistor exposed to an input of M might not be completely switched on or off, resulting in an undefined output voltage, i.e., M . Hence, when applying stable voltages to a transistor’s source and drain but M to its gate, it acts as an undefined source–drain resistance.

Note that while combinational logic may output stable values under metastable input, its timing behavior may still be degraded; in fact, we observe this effect in Section 12.4. It is well-known that violated timing




Figure 12.1: Gate-level MUX behavior under inputs $a = b = 1$; in this case, the output should be 1 regardless of the select bit s . For $s = M$, however, the standard MUX (a) can become metastable. The CMUX implementation (b) uses logical masking and outputs 1. Figure adapted from [FFL18].

constraints may result in metastability in latches and flip-flops. Hence, we denote violated timing constraints by M as well.

12.2 Problem Statement

A *Metastability-Containing Multiplexer (CMUX)* has three single-bit² inputs s, a, b , one output, and behaves as follows.

- (1) If $s = 0$ the output is a ,
- (2) if $s = 1$ the output is b ,
- (3) if $s = M$ but $a = b$ the output is a , and
- (4) the output is unspecified if $s = M$ and $a \neq b$.

Condition (3) reflects that even if the choice between a and b is unclear due to a metastable select bit, this must not affect the output in the case of $a = b$. Omitting Condition (3) yields a standard MUX which is oblivious to degraded select bits.

Figure 12.1a shows how a gate-level MUX can become metastable under input $s = M$ and $a = b = 1$. The gate-level CMUX in Figure 12.1b mitigates the effect of $s = M$ using logical masking, i.e., a third AND-gate. Friedrichs et al. show that the circuit in Figure 12.1b satisfies (1)–(4) [FFL18]. However, they do not present an efficient implementation or simulations.

We propose efficient transistor-level implementations which maintain Conditions (1)–(4) and simulate them to verify their behavior. Using 8 and 10 transistors, our implementations require a fraction of the 18 transistors for a naive implementation³ of the CMUX in

2. Specification and circuits directly generalize to arbitrary bit-widths.

3. 4 and 6 transistors for a binary and a ternary NAND, respectively; The latter is the de Morgan equivalent of the ternary OR.

Figure 12.1b.

Our focus is a proof of concept using a generic process and node size. Hence, we do not analyze process variations, leakage current, or area overhead as they vary with technology. Further, note that the actual likelihood of degenerate input signals heavily depends on the input-generating circuit.

12.3 Circuits

We propose two transistor-level implementations of a CMUX, CMUX-A and CMUX-B, depicted in Figures 12.2c and 12.2d, respectively. Both implement Conditions (1)–(4) and invert the output. CMUX-A is a conservative implementation that requires 10 transistors. CMUX-B needs only 8 transistors—the same as the conventional MUX-C—but has a slightly higher peak current and a slightly degraded output voltage under Condition (3), i.e., if $s = M$ and $a = b$.

CMUX-A

Clearly, CMUX-A (Figure 12.2c) satisfies Conditions (1) and (2) if all inputs s, a, b are logical 0 or 1, up to inversion of the output. Regarding Condition (3), recall Section 12.1: A transistor with an undefined input voltage produces an undefined output voltage. So if $s = M$ and $a = b = 0$, there is a low-resistance path from \bar{o} to V_{DD} at the top right of CMUX-A, but only high-resistance paths from \bar{o} to GND, so $\bar{o} = 1 = \bar{a}$. Note that all transistors with voltage M at the gate are bypassed and do not influence the output voltage. CMUX-A behaves symmetrically if $s = M$ and $a = b = 1$. Together, we see that CMUX-A satisfies Conditions (1)–(4). We discuss in Section 12.4 that SPICE simulations confirm the correct behavior of CMUX-A.

CMUX-B

It is easy to check that CMUX-B (Figure 12.2d) fulfills Conditions (1) and (2): If all inputs s, a, b are logical 0 or 1, CMUX-B outputs \bar{a} if $s = 0$ and \bar{b} if $s = 1$. As for Condition (3), consider the the case $s = M$ and $a = b$. Then transistors with gate input s or \bar{s} act as resistors of undefined resistance, hence, there is a current from V_{DD} to GND along the left path of CMUX-B. This results in a higher peak current under $s = M$. If $a = b = 1$ the n-Field-Effect Transistors (FETs) pull the output low while the p-FETs on the left represent an undefined resistance, which may become low enough to cause an undefined output level. This constitutes the key difference to CMUX-A, which masks this

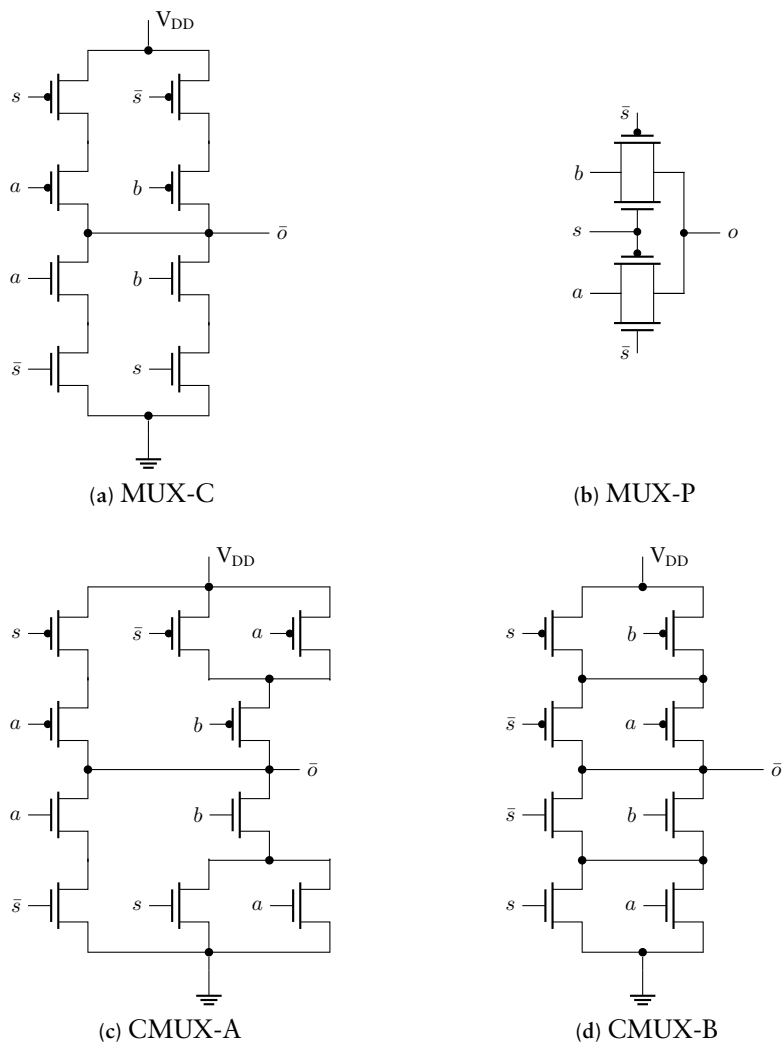


Figure 12.2: Transistor-level MUX and CMUX implementations. MUX-C (a) and MUX-P (b) show conventional combinational and pass-gate MUXes, respectively. Both are vulnerable to metastability, especially the MUX-P: A metastable (or in any other way degraded) select bit causes an undefined pass resistance and connecting them in series severely degrades the output signal. CMUX-A (c) and CMUX-B (d) implement CMUXes, i.e., metastability-containing alternatives. CMUX-A is a conservative implementation that requires 10 transistors and has a low peak current. CMUX-B saves 2 transistors—matching the transistor count of MUX-C—at the expense of a slightly increased peak current in the case of $s = M$ and $a = b$.

case. To ensure correct output levels we use transistors of double width on the right branch with inputs a and b . Our simulations confirm that with this tweak CMUX-B behaves correctly under Condition (3), see Section 12.4. In the case of $a = b = 0$, CMUX-B behaves symmetrically. Altogether, CMUX-B satisfies Conditions (1)–(4), but has a higher peak current than CMUX-A under Condition (3). See Section 12.4 for SPICE simulations confirming our claims regarding correct behavior and peak current.

On Pass-Gate Implementations

MUX-P (Figure 12.2b) is a standard pass-gate implementation. Such implementations have a non-negligible pass-resistance, even for stable inputs, degrading the output signal. This problem severely intensifies in the context of metastability-containment: $s = M$ effectively turns all transistors in MUX-P into undefined resistances or — depending on the Complementary Metal Oxide (CMOS) process — even worse, into source-followers with an input voltage close to mid-voltage. This effect is especially problematic when connecting MUXes in series. Hence, the output signal has to be cleaned up, surrendering the advantage of a pass-gate implementation in the first place. Therefore, we do not base our CMUX implementations on pass gates.

12.4 Simulation

We simulated CMUX-A and CMUX-B, and the reference implementations MUX-C and MUX-P, in SPICE, using a 180 nm process with transistors sized $2.41 \mu\text{m}/0.18 \mu\text{m}$ and $4.90 \mu\text{m}/0.18 \mu\text{m}$ for N- and P-FETs respectively. As discussed in Section 12.3, the transistors on the right branch of CMUX-B were twice the minimal width to guarantee a stable output voltage under all operating conditions.

In each simulation, we applied a fixed voltage to s and simultaneously switched a and b , maintaining $a = b$, from 0 to 1 or from 1 to 0. To obtain realistic waveforms, a and b passed through two inverters. For each circuit, we measured the delay between the first change of the input and the stabilization of its output to 0.18 V (10 %) or to 1.62 V (90 %).

Figure 12.3 shows the input waveform applied to a and b (blue) and the output behavior under a wide range of voltages of s . Figures 12.3a–12.3d show the behavior of the conventional reference implementations MUX-C and MUX-P and Figures 12.3e–12.3h that of the proposed implementations CMUX-A and CMUX-B. Table 12.1 summarizes the best-case (nominal) and worst-case delays, the end

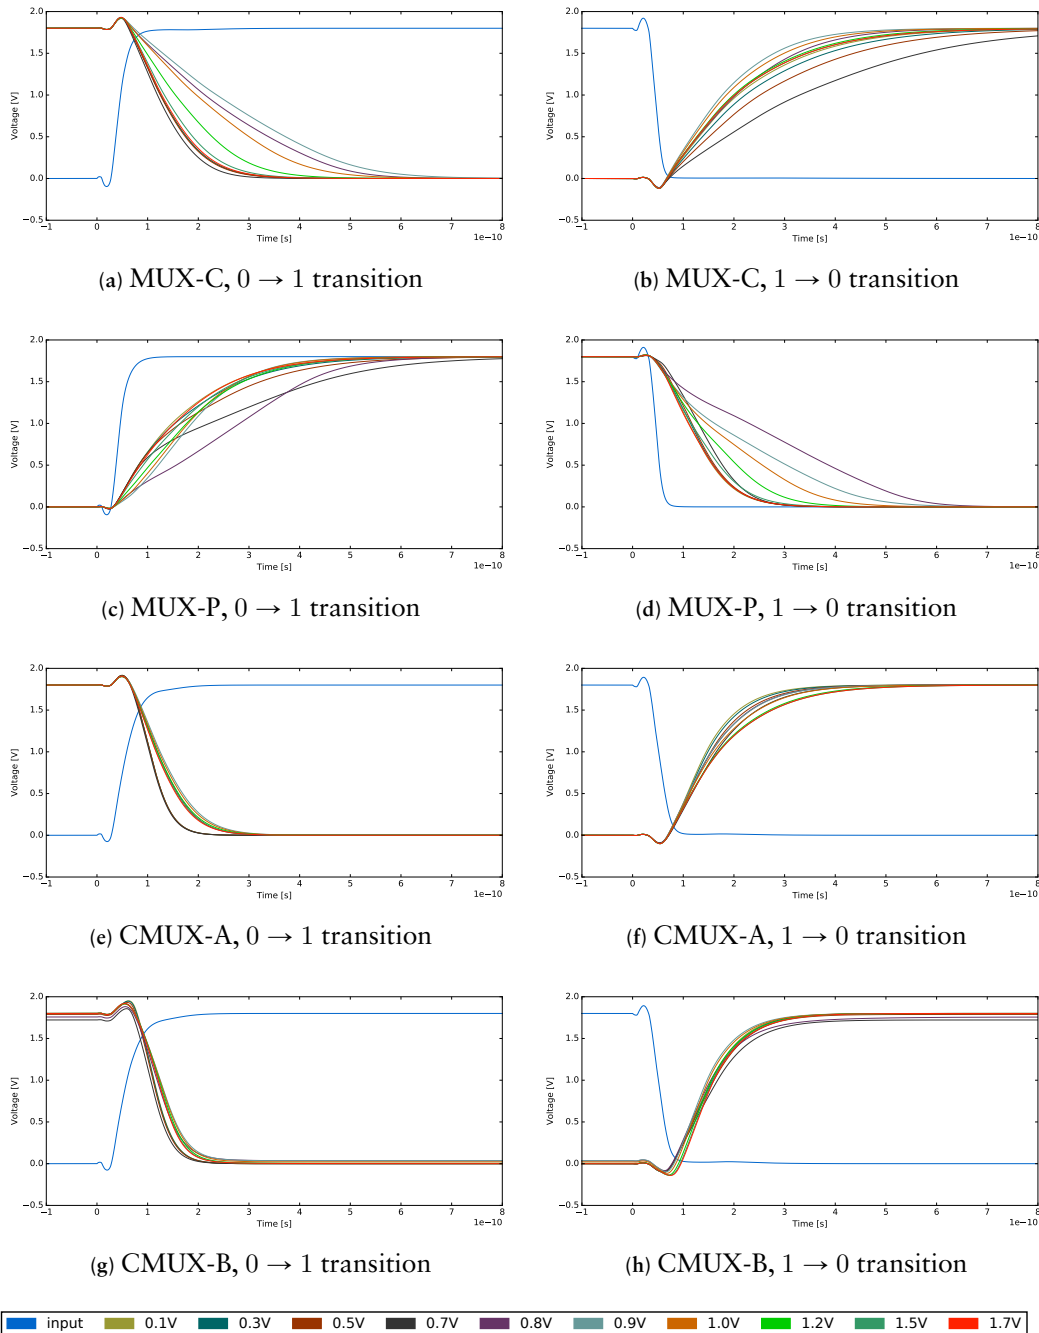


Figure 12.3: SPICE simulation results. The graphs show the input signal (blue) for $a = b$, and the obtained output signal for different voltages (between 0.1 V and 1.7 V) applied to s . The timing performance of the conventional MUX implementations MUX-C (a)–(b) and MUX-P (c)–(d) severely degrades under a wide variety of these voltages. In contrast, CMUX-A (e)–(f) and CMUX-B (g)–(h) hardly show any degradation.

Type	MUX-C Fig. 12.2a	MUX-P Fig. 12.2b	CMUX-A Fig. 12.2c	CMUX-B Fig. 12.2d
Transistor Count	8	4	10	8
Prop. Delay 0 → 1 (nominal) 1 → 0	215 ps 319 ps	315 ps 206 ps	155 ps 237 ps	155 ps 231 ps
Prop. Delay 0 → 1 (worst-case) 1 → 0	600 ps 728 ps	607 ps 532 ps	222 ps 326 ps	192 ps 300 ps
End Voltage 0 → 1 (worst-case) 1 → 0	0.001 V 1.799 V	1.799 V 0.001 V	0.001 V 1.799 V	0.049 V 1.709 V
Avg. Power 0 → 1 (nominal) 1 → 0	13.5 μ W 42.6 μ W	30.6 μ W 19.4 μ W	14.9 μ W 44.9 μ W	18.4 μ W 51.6 μ W
Avg. Power 0 → 1 (worst-case) 1 → 0	429 μ W 458 μ W	447 μ W 436 μ W	431 μ W 461 μ W	657 μ W 709 μ W

Table 12.1: SPICE simulation results. *Nominal* and *worst-case* refer to the condition of the select bit.

voltages after 8 ns, and the average power consumption (including the input inverters) over 10 ns.

Clearly, the output waveforms of the reference implementations MUX-C and MUX-P degrade heavily under the influence of metastability at s . Furthermore, the propagation delays become large as the voltage of s approaches 0.8 V. Hence, despite eventually reaching a well-defined output voltage, they do not contain metastability: The severely degraded propagation delay of MUX-C and MUX-P leads to violated timing constraints when used in larger circuits. This in turn can lead to the violation of setup/hold times of registers further downstream, increasing the risk of metastable upsets.

In stark contrast, Figures 12.3e–12.3h demonstrate CMUX-A and CMUX-B are indeed metastability-containing: they only show negligible degradation. In terms of power consumption, CMUX-B requires slightly more than CMUX-A and MUX-C in the nominal and about 50 % more in the worst case. This is due to the cross links between the branches that cause an additional current flow through the FETs controlled by s and \bar{s} .

The bottom line is that CMUX-A and CMUX-B cope with even the most severely degraded select bits, and thus fulfill Conditions (1)–(4). As the conventional implementations fail to fulfill Condition (3), the proposed CMUXes form a solid foundation for metastability-containing circuits capable of keeping timing guarantees without undue guard bands or comprising the MTBF.

bits	4-sort			7-sort			10-sort		
	[LM16]	we	saved	[LM16]	we	saved	[LM16]	we	saved
2	370	340	8.1 %	1010	1010	0.0 %	1750	1750	0.0 %
4	2990	1620	45.8 %	7654	5184	32.3 %	12986	9396	27.6 %
8	21060	7060	66.5 %	67392	22592	66.5 %	122148	40948	66.5 %
16	95460	29460	69.1 %	305472	94272	69.1 %	553668	170868	69.1 %

bits	4-sort			7-sort			10-sort		
	[BLM17]	we	saved	[BLM17]	we	saved	[BLM17]	we	saved
2	960	400	58.3 %	3072	1280	58.3 %	5568	2320	58.3 %
4	4440	1640	63.1 %	14208	5248	63.1 %	25752	9512	63.1 %
8	13920	4960	64.4 %	44544	15872	64.4 %	80736	28768	64.4 %
16	37080	13000	65.0 %	118656	41600	65.0 %	215064	75400	65.0 %

Table 12.2: Transistors saved by using CMUX-B in the sorting networks of Lenzen and Medina [LM16] and Bund et al. [BLM17].

12.5 Impact on Metastability-Containing Circuits

In this section, we demonstrate how our CMUX implementations drastically reduce the complexity of existing metastability-containing circuits. This is achieved by using CMUX-A or CMUX-B as a drop-in replacement wherever a CMUX is required. As an example, consider the metastability-containing Gray code sorting networks proposed by Lenzen and Medina [LM16] and the more recent improvement by Bund et al. [BLM17]; Using CMUX-B reduces the size of a 16-bit 10-sort by 69 % and 65 %, respectively.

Both papers describe 2-sort implementations for Gray code numbers and metastable superpositions of successive Gray code numbers, as motivated in Section 12.1; Lenzen and Medina require $\mathcal{O}(B^2)$ transistors for a B -bit comparator, and Bund et al. reduce the transistor count to $\mathcal{O}(B \log B)$. Each paper presents the transistor counts required by optimal sorting networks of 4, 7, and 10 channels for $B \in \{2, 4, 8, 16\}$ assembled from their respective 2-sort implementation.

The proposed 2-sort implementations require a 4-CMUX. We require 28 transistors to implement a 4-CMUX from CMUX-B, much less than the 84 transistors of the naive implementation of Lenzen and Medina based on Figure 12.1b.

Table 12.2 demonstrates how replacing the naive CMUX used in [BLM17; LM16] with CMUX-B from Figure 12.2d reduces the transistor count of the sorting networks. As for the comparison with [LM16], note that Lenzen and Medina present a polynomially as well as an exponentially sized circuit. We always plug CMUX-B into the polynomially sized circuit and compare against their smallest circuit, hence the reduced savings for small B . When sorting 4-bit numbers, our

approach saves between 27.6 % and 45.8 % of the transistors; our approach even saves 69.1 % when sorting 16-bit numbers. The savings w.r.t. [BLM17] are equally drastic: Using CMUX-B reduces the size of the sorting network by 58.3 %–65.0 %, removing almost two thirds of the circuit.

12.6 Conclusion

We propose two transistor-level implementations of Metastability-Containing Multiplexers (CMUXes)—MUXes that guarantee correct behavior in the presence of a severely degraded or metastable select bit—and verify them using SPICE simulations. Despite having little to no overhead in transistor count and power consumption w.r.t. a conventional MUX, (1) our implementations perform extremely well in the presence of degraded select bits, (2) show almost no degradation in performance even under worst-case conditions, and (3) outperform conventional MUX architectures in terms of propagation delay under simultaneously switching inputs.

Furthermore, we demonstrate the usefulness of our CMUX: It drastically improves the metastability-containing sorting networks by Lenzen and Medina [LM16] and Bund et al. [BLM17], reducing the transistor count by up to 69 % and 65 %, respectively.

PART IV

Backmatter

List of Figures

6.1	SI units relation graph after the 2019 redefinition of the SI base units. From [Pis16].	51
7.1	The circuit model of a sine-to-square converter is simplified to a noiseless comparator input stage with some hysteresis $\pm H$ and all input related noise being lumped together into the offset voltage V_{ofs} . Noise due to power supply variation $V_{\text{DD},n}$ is modeled using an amplification stage with a delay of t_{amp} that only depends on the supply voltage.	56
7.2	The zero crossing of the input signal gets delayed by the offset voltage V_{ofs} and the hysteresis $\pm H$ and the delay t_{amp} through the amplifier stage. The offset voltage and hysteresis related delays t_{ofs} and t_{H} depend not only on the values of V_{ofs} and H respectively, but also on the slew-rate of the input signal. The amplifier delay t_{amp} only depends on the supply voltage V_{DD}	58
8.1	The circuit model, adapted from [Col96] of a single stage zero-crossing detector is simplified to a noiseless amplifier input stage, with all input referred noise being lumped together into the noise voltage v_n^2 . After the amplifier there is a (noise-less) limiter and a first order filter.	68
8.2	Comparison of the Fourier transform of the output signal \hat{V}_{out} and the Fourier transform of the impulse sensitivity function (ISF) $\hat{\Xi}$. While \hat{V}_{out} first decays with $1/f$ up to the frequency of $1/(\pi\tau)$ and then decays with $1/f^2$, the ISF is constant up to the frequency of $1/(\pi\tau_w)$ and decays with $1/f$ afterwards. Generally, $\tau > \tau_w$ can be assumed.	70
10.1	Histogram of the encoded TDL output values during offline calibration. These values correspond to the bin sizes of the delay line.	83
10.2	Experimental setup. Left: measurement instruments. Center: nodes with FPGA and interface boards, and a stray mouse. Right: recording PC.	84

10.3	Long-term evolution of the clock skew of three nodes against the same reference node over a period of an hour, measured sequentially. The thick lines depict the average clock skew over 10 s, the light yellow colored fill with the thin lines depict the minimum and maximum in the same interval.	85
10.4	Short-term behavior of the clock skew of the "blue" node vs. the reference node from Figure 10.3 over an arbitrarily selected period of 20 s.	85
10.5	TDEV between all pairs of nodes, measured sequentially. The colors of pairs match those from Figure 10.3.	86
10.6	ADEV between a node and a rubidium frequency standard. The top (blue) curve compares to the free-running oscillator, with the algorithm disabled. The bottom (red) curve is the same node with the algorithm enabled and the system fully synchronized. The temperature effects beyond $\tau = 100$ s differ because the traces were recorded on subsequent days with different weather conditions and thus different heating patterns of the building.	87
11.1	Tapped delay-line TDC. Latches are initially enabled and output 0. The delayed starting signal iteratively sets latches to 1 until the stopping signal disables them.	89
11.2	Generic ring TDC architecture with coarse counter C . Latches and counter are initialized to 0. The counter increments on rising and falling transitions.	94
11.3	Transient simulation of the capture register of the simulated TDC. Even very small variations of the applied input signal increase the time until the output value reaches its final voltage level considerably. Note that a full D-FF was used instead of a single latch.	95
11.4	Stabilization time variation due to metastable upset of the simulated TDC capture register. time between application of the input and the clock edge vs. output stabilization time.	95
11.5	Solution I: Implementing C by two binary counters and a latch. Counters C_0 and C_1 and latch L are initialized to 0. The increment inputs of C_0 and C_1 are driven via a delay buffer. C_0 increments on rising, C_1 on falling transitions.	97

11.6	Critical windows for Solution I as the START signal travels through the ring oscillator. First the first stage “handles” the START signal during the stage delay time T_d and can become metastable during its critical window W_0 , then the second stage during its window W_1 and so on. When the START signal traveled once around the the ring oscillator it will increment the first counter C_0 after the buffer delay T_{buf} and thus cause it to enter its critical window W_{C_0} , while the signal continues to progress through the ring oscillator. After the second round, the second counter C_1 is incremented and can potentially become metastable during its critical window W_{C_1}	98
11.7	Critical windows for Solution II.	100
11.8	Solution II: Implementation with look-ahead Gray code counter and register. The counter C is initialized with 1. Register R_C is of same width as the counter, and initialized with 0.	100
11.9	XOR-tree circuit implementation counting the number of 1s as specified in Table 11.1. Note that a metastable ring latch influences only one, the current least significant, bit (compare lines 4 and 5 in Table 11.1): no precision is lost when encoding metastable data.	101
11.10	Circuit encoding a TDC measurement value ($n - 1$ thermometer encoded ring latch states and B bit coarse BRGC counter state) as a single BRGC encoded value without losing precision in presence of a metastable input bit. Here $n = 8$ and $B = 2$ and the coarse counter is even (decimal 0), i.e., we have to count 1s in latch states. Observe that the final value 0011M correctly encodes decimal 4 or 5, depending on how M resolves.	103
11.11	Left: coarse counter is odd (decimal 1), i.e., we have to count 0s in latch states. Observe that the final value 0101M correctly encodes decimal 12 or 13, depending on how M resolves. Right: coarse counter is metastable. Observe that the final value 0M100 correctly encodes decimal 7 or 8, depending on how M resolves.	103
11.12	Solution III: Gray code counter variant with single-bit latch.	104
11.13	Critical windows for Solution III.	104

- 12.1 Gate-level MUX behavior under inputs $a = b = 1$; in this case, the output should be 1 regardless of the select bit s . For $s = M$, however, the standard MUX (a) can become metastable. The CMUX implementation (b) uses logical masking and outputs 1. Figure adapted from [FFL18]. 111
- 12.2 Transistor-level MUX and CMUX implementations. MUX-C (a) and MUX-P (b) show conventional combinational and pass-gate MUXes, respectively. Both are vulnerable to metastability, especially the MUX-P: A metastable (or in any other way degraded) select bit causes an undefined pass resistance and connecting them in series severely degrades the output signal. CMUX-A (c) and CMUX-B (d) implement CMUXes, i.e., metastability-containing alternatives. CMUX-A is a conservative implementation that requires 10 transistors and has a low peak current. CMUX-B saves 2 transistors—matching the transistor count of MUX-C—at the expense of a slightly increased peak current in the case of $s = M$ and $a = b$ 113
- 12.3 SPICE simulation results. The graphs show the input signal (blue) for $a = b$, and the obtained output signal for different voltages (between 0.1 V and 1.7 V) applied to s . The timing performance of the conventional MUX implementations MUX-C (a)–(b) and MUX-P (c)–(d) severely degrades under a wide variety of these voltages. In contrast, CMUX-A (e)–(f) and CMUX-B (g)–(h) hardly show any degradation. 115

List of Tables

7.1	Summary of the frequency scaling of the different noise types for the pure noise, after a single stage amplifier and after a multi-stage amplifier. The frequency scaling for delay (V_{DD}) related flicker noise after a multi-stage amplifier is ν^α with $\alpha \in [2, 3]$	65
11.1	Encoding the ring latch states in case $n = 8$: original thermometer encoded ($n - 1$ bits $0 \dots 6$), efficient BRGC encoded ($\log_2(n)$ bits $0 \dots 2$), and the relevant decimal counts of 1s and 0s.	101
12.1	SPICE simulation results. <i>Nominal</i> and <i>worst-case</i> refer to the condition of the select bit.	116
12.2	Transistors saved by using CMUX-B in the sorting networks of Lenzen and Medina [LM16] and Bund et al. [BLM17].	117

List of Acronyms

CMUX	Metastability-Containing Multiplexer (extra information)
ADC	analog to digital converter
ADEV	Allan deviation
a.e.	almost everywhere
BRGC	binary reflected Gray code
D-FF	D-flip-flop
DAC	digital to analog converter
DMTD	dual mixer time difference
ESD	energy spectral density
fBm	fractional Brownian motion
FIFO	first in - first out
FPGA	field-programmable gate array
GALS	globally asynchronous locally synchronous
iid	independent and identically distributed
LVC MOS	low-voltage CMOS
MTBF	mean time between failures
PLL	phase-locked loop
PSD	power spectral density
SPI	serial peripheral interface
TDC	time to digital converter
TDEV	time deviation
VCXO	voltage controlled crystal oscillator
VLSI	very large scale integration

Bibliography

- [ACH97] Leif Arkeryd, Nigel Cutland, and Ward Henson. “Non-standard Analysis.” Kluwer Academic Publishers, 1997. ISBN: 978-94-010-6335-7. DOI: 10.1007/978-94-011-5544-1.
- [AD75] D. W. Allan and H. Daams. “Picosecond Time Difference Measurement System.” In: *29th Annual Symposium on Frequency Control*. May 1975, pp. 404–411. DOI: 10.1109/FREQ.1975.200112.
- [Amb82] András Ambrózy. “Electronic Noise.” McGraw-Hill, 1982. ISBN: 0-07-001124-9.
- [BA66] J. A. Barnes and D. W. Allan. “A statistical model of flicker noise.” In: *Proceedings of the IEEE* 54.2 (1966), pp. 176–178. DOI: 10.1109/PROC.1966.4630.
- [Bat53] Harry Bateman. “Higher Transcendental Functions.” McGraw-Hill, 1953.
- [BEI01] J.A. Brzozowski, Z. Esik, and Y. Iland. “Algebras for hazard detection.” In: *Proceedings 31st IEEE International Symposium on Multiple-Valued Logic*. 2001, pp. 3–12. DOI: 10.1109/ISMVL.2001.924548.
- [Ber37] J. Bernamont. “Fluctuations de potentiel aux bornes d’un conducteur métallique de faible volume parcouru par un courant.” In: *Ann. Phys.* 11.7 (1937), pp. 71–140. DOI: 10.1051/anphys/193711070071.
- [BG15] Salomon Beer and Ran Ginosar. “Eleven Ways to Boost Your Synchronizer.” In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 23.6 (2015), pp. 1040–1049. DOI: 10.1109/TVLSI.2014.2331331.
- [BIP18] BIPM. “Comptes rendus de la 26^e réunion de la Conférence générale des poids et mesures.” In: *26^e Conférence Générale des Poids et Mesures*. Bureau international des poids et mesures, 2018. URL: <https://www.bipm.org/fr/CGPM/db/26>.

- [BIP83] BIPM. “Comptes rendus de la 17^e réunion de la Conférence générale des poids et mesures.” In: *17^e Conférence Générale des Poids et Mesures*. Bureau international des poids et mesures, 1983. URL: <https://www.bipm.org/fr/CGPM/db/17>.
- [BLM17] Johannes Bund, Christoph Lenzen, and Moti Medina. “Near-optimal metastability-containing sorting networks.” In: *Design, Automation & Test in Europe Conference (DATE), 2017*. 2017, pp. 226–231. DOI: 10.23919/DATE.2017.7926987.
- [Boc33] Salomon Bochner. “Integration von Funktionen, deren Werte die Elemente eines Vektorraumes sind.” In: *Fundamenta Mathematicae* 20 (1933), pp. 262–277. DOI: 10.4064/fm-20-1-262-176.
- [Bog07] Vladimir Bogachev. “Measure Theory.” Springer, 2007. ISBN: 978-3-540-34513-8. DOI: 10.1007/978-3-540-34514-5.
- [Bou11] Sébastien Bourdeauducq. “Time to Digital Converter Core for Spartan-6 FPGAs.” Nov. 2011. URL: <http://www.ohwr.org/documents/98>.
- [Bou13] Sébastien Bourdeauducq. “A 26 ps RMS time-to-digital converter core for Spartan-6 FPGAs.” Mar. 2013. URL: <http://arxiv.org/pdf/1303.6840v1.pdf>.
- [Bou19] Sébastien Bourdeauducq. “OHWR TDC core.” <http://www.ohwr.org/projects/tdc-core/wiki>. 2019.
- [Bre+19] S. M. Brewer et al. “ $^{27}\text{Al}^+$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18} .” In: *Phys. Rev. Lett.* 123.033201 (3 July 2019). DOI: 10.1103/PhysRevLett.123.033201.
- [BY79] Brzozowski and Yoeli. “On a Ternary Model of Gate Networks.” In: *IEEE Transactions on Computers* C-28.3 (1979), pp. 178–184. DOI: 10.1109/TC.1979.1675317.
- [Cam18] Evan Camrud. “A Novel Approach to Fractional Calculus: Utilizing Fractional Integrals and Derivatives of the Dirac Delta Function.” In: *Progress in Fractional Differentiation and Applications* 4 (2018), pp. 463–478. DOI: 10.18576/pfda/040402.

- [Cha+17] Chi-Hang Chan et al. “Metastability in SAR ADCs.” In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 64.2 (2017), pp. 111–115. DOI: 10.1109/TCSII.2016.2554798.
- [Cha84] Daniel M. Chapiro. “Globally-Asynchronous Locally-Synchronous Systems.” PhD thesis. Stanford University, 1984.
- [Çin11] Erhan Çinlar. “Probability and Stochastics.” Springer, 2011. ISBN: 978-0-387-87858-4. DOI: 10.1007/978-0-387-87859-1.
- [Col96] O. Collins. “The design of low jitter hard limiters.” In: *IEEE Transactions on Communications* 44.5 (May 1996), pp. 601–608. ISSN: 0090-6778. DOI: 10.1109/26.494304.
- [CR14] C. E. Calosso and E. Rubiola. “Phase noise and jitter in digital electronics.” In: *2014 European Frequency and Time Forum (EFTF)*. 2014, pp. 374–376. DOI: 10.1109/EFTF.2014.7331514.
- [Dem00] Alper Demir. “Floquet theory and non-linear perturbation analysis for oscillators with differential-algebraic equations.” In: *International Journal of Circuit Theory and Applications* 28.2 (2000), pp. 163–185. DOI: 10.1002/(SICI)1097-007X(200003/04)28:2<163::AID-CTA101>3.0.CO;2-K.
- [DK10] J.J. Duistermaat and J.A.C. Kolk. “Distributions.” 1st ed. Birkhäuser, 2010. ISBN: 978-0-8176-4672-1.
- [DKS90] G. John Dick, Paul F. Kuhnle, and Richard L. Sydnor. “Zero-crossing detector with sub-microsecond jitter and crosstalk.” In: *22nd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*. May 1990, pp. 269–282. URL: http://time.kinali.ch/ptti/1990papers/Vol%2022_20.pdf.
- [DMR00] A. Demir, A. Mehrotra, and J. Roychowdhury. “Phase noise in oscillators: a unifying theory and numerical methods for characterization.” In: *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* 47.5 (May 2000), pp. 655–674. ISSN: 1057-7122. DOI: 10.1109/81.847872.

- [Dol+86] Danny Dolev et al. “Reaching Approximate Agreement in the Presence of Faults.” In: *Journal of the ACM* 33 (3 1986), pp. 499–516. DOI: 10.1145/5925.5931.
- [DR89] Francine Diner and Georges Reeb. “Analyse Non Standard.” Hermann, 1989. ISBN: 978-2-7056-6109-0.
- [DR97] Alper Demir and Jaijeet Roychowdhury. “On the Validity of Orthogonally Decomposed Perturbations in Phase Noise Analysis.” Tech. Memo. Bell Laboratories, 1997.
- [DSH00] P. Dudek, S. Szczepanski, and J.V. Hatfield. “A high-resolution CMOS time-to-digital converter utilizing a Vernier delay line.” In: *Solid-State Circuits, IEEE Journal of* 35.2 (Feb. 2000), pp. 240–247. ISSN: 0018-9200. DOI: 10.1109/4.823449.
- [Ega90] W. F. Egan. “Modeling phase noise in frequency dividers.” In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 37.4 (July 1990), pp. 307–315. ISSN: 0885-3010. DOI: 10.1109/58.56498.
- [Eic65] E. B. Eichelberger. “Hazard Detection in Combinational and Sequential Switching Circuits.” In: *IBM Journal of Research and Development* 9.2 (1965), pp. 90–99. DOI: 10.1147/rd.92.0090.
- [Fer+09] Eva S. Ferre-Pikal et al. “IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities.” In: *IEEE Std 1139-2008* (Feb. 2009). DOI: 10.1109/IEEESTD.2009.6581834.
- [FFL18] Stephan Friedrichs, Matthias Függer, and Christoph Lenzen. “Metastability-Containing Circuits.” In: *IEEE Transactions on Computers* 67.8 (2018), pp. 1167–1183. DOI: 10.1109/TC.2018.2808185.
- [FK17] Stephan Friedrichs and Attila Kinali. “Efficient Metastability-Containing Multiplexers.” In: *2017 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. 2017, pp. 332–337. DOI: 10.1109/ISVLSI.2017.65.
- [Fri01] E.G. Friedman. “Clock distribution networks in synchronous digital integrated circuits.” In: *Proceedings of the IEEE* 89.5 (May 2001), pp. 665–692. ISSN: 0018-9219. DOI: 10.1109/5.929649.

- [Fri44] H. T. Friis. “Noise Figures of Radio Receivers.” In: *Proceedings of the IRE* 32.7 (July 1944), pp. 419–422. ISSN: 0096-8390. DOI: 10.1109/JRPROC.1944.232049.
- [Füg+17] Matthias Függer et al. “Metastability-Aware Memory-Efficient Time-to-Digital Converters.” In: *2017 23rd IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*. 2017, pp. 49–56. DOI: 10.1109/ASYNC.2017.12.
- [GHL12] Shankar Guhadós, Paul J. Hurst, and Stephen H. Lewis. “A Pipelined ADC With Metastability Error Rate $< 10^{-15}$ Errors/Sample.” In: *IEEE Journal of Solid-State Circuits* 47.9 (2012), pp. 2119–2128. DOI: 10.1109/JSSC.2012.2198773.
- [Gin03] R. Ginosar. “Fourteen ways to fool your synchronizer.” In: *Ninth International Symposium on Asynchronous Circuits and Systems, 2003. Proceedings*. 2003, pp. 89–96. DOI: 10.1109/ASYNC.2003.1199169.
- [Gol98] Robert Goldblatt. “Lectures on the Hyperreals.” Springer, 1998. ISBN: 978-1-4612-0615-6. DOI: 10.1007/978-3-0348-5567-9.
- [Gra14] Loukas Grafakos. “Classical Fourier Analysis.” 3rd ed. Springer, 2014. ISBN: 978-1-4939-1193-6. DOI: 10.1007/978-1-4939-1194-3.
- [Hal50] Paul R. Halmos. “Measure Theory.” Springer, 1950. ISBN: 978-1-4684-9442-6. DOI: 10.1007/978-1-4684-9440-2.
- [Hen10] Stephan Henzler. “Time-to-Digital Converters.” 1st ed. Springer, 2010. ISBN: 978-90-481-8628-0. DOI: 10.1007/978-90-481-8628-0.
- [HH15] Paul Horowitz and Winfield Hill. “The Art of Electronics.” 3rd ed. Cambridge University Press, 2015. ISBN: 978-0-521-80926-9. URL: <https://artofelectronics.net/>.
- [HL98] A. Hajimiri and T. H. Lee. “A general theory of phase noise in electrical oscillators.” In: *IEEE Journal of Solid-State Circuits* 33.2 (1998), pp. 179–194. DOI: 10.1109/4.658619.
- [Hob+20] Richard Hobson et al. “A strontium optical lattice clock with 1×10^{-17} uncertainty and measurement of its absolute frequency.” In: *Metrologia* 57.6 (Oct. 2020), p. 065026. DOI: 10.1088/1681-7575/abb530.

- [Hof02] Josef Hofbauer. “A Simple Proof of $1 + 1/2^2 + 1/3^2 + \dots = \pi^2/6$ and Related Identities.” In: *The American Mathematical Monthly* 109.2 (2002), pp. 196–200. ISSN: 00029890, 19300972. DOI: 10.2307/2695334.
- [Hol65] Hjalmar Josef Holmgren. “Om differentialkalkylen med Indices af hvad Nator som helst.” In: *Vetenskapssakademiens handlingar* (1865).
- [HP97] HP. “Fundamentals of the Electronic Counter.” Application Note. Hewlett Packard, 1997. URL: https://www.hpmemoryproject.org/an/pdf/an_200.pdf.
- [Huf57] David A. Huffman. “The Design and Use of Hazard-Free Switching Networks.” In: *J. ACM* 4.1 (Jan. 1957), pp. 47–62. ISSN: 0004-5411. DOI: 10.1145/320856.320866.
- [Joh14] Staffan Johansson. “Modern frequency counting principles.” Tech. rep. Pendulum Instruments, 2014.
- [Joh25] J. B. Johnson. “The Schottky Effect in Low Frequency Circuits.” In: *Phys. Rev.* 26.1 (July 1925), pp. 71–85. DOI: 10.1103/PhysRev.26.71.
- [Kal03] Józef Kalisz. “Review of methods for time interval measurements with picosecond resolution.” In: *Metrologia* 41.1 (Dec. 2003), pp. 17–32. DOI: 10.1088/0026-1394/41/1/004.
- [KBY02] D.J. Kinniment, A. Bystrov, and A.V. Yakovlev. “Synchronization circuit performance.” In: *IEEE Journal of Solid-State Circuits* 37.2 (2002), pp. 202–209. DOI: 10.1109/4.982426.
- [KHL16] Attila Kinali, Florian Huemer, and Christoph Lenzen. “Fault-Tolerant Clock Synchronization with High Precision.” In: *2016 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. 2016, pp. 490–495. DOI: 10.1109/ISVLSI.2016.88.
- [Kin08] David J. Kinniment. “Synchronization and Arbitration in Digital Systems.” Wiley Publishing, 2008. ISBN: 978-0-470-51082-7.
- [Kin18] Attila Kinali. “A Physical Sine-to-Square Converter Noise Model.” In: *2018 IEEE International Frequency Control Symposium (IFCS)*. 2018, pp. 1–6. DOI: 10.1109/IFCS.2018.8597525.

- [Kin19] Attila Kinali. “A Fresh Look at the Design of Low Jitter Hard Limiters.” In: *2019 Joint Conference of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (EFTF/IFC)*. 2019, pp. 1–4. DOI: 10.1109/FCS.2019.8855997.
- [KL19] Pankaj Khanchandani and Christoph Lenzen. “Self-stabilizing Byzantine Clock Synchronization with Optimal Precision.” In: *Theory of Computing Systems* 63 (Feb. 2019). DOI: 10.1007/s00224-017-9840-3.
- [Knu97] Donald E. Knuth. “The Art of Computer Programming, Volume 1: Fundamental Algorithms.” 3rd ed. Redwood City, CA, USA: Addison Wesley Longman Publishing Co., Inc., 1997. ISBN: 0-201-89683-4.
- [Kop03] H. Kopetz. “Fault containment and error detection in the time-triggered architecture.” In: *Autonomous Decentralized Systems, 2003. ISADS 2003. The Sixth International Symposium on*. Apr. 2003, pp. 139–146. DOI: 10.1109/ISADS.2003.1193942.
- [Lap09] Amos Lapidot. “A Foundation in Digital Communication.” 1st ed. Cambridge University Press, 2009. ISBN: 978-0-521-19395-5. URL: https://www.afidc.ethz.ch/A_Foundation_in_Digital_Communication/Home.html.
- [Lév65] Paul Lévy. “Processus stochastiques et mouvement Brownien.” 2nd ed. Gauthier-Villars, 1965.
- [Lio32] Joseph Liouville. “Sur quelques Questions de Géométrie et de Mécanique, et sur un nouveau genre de Calcul pour résoudre ces Question.” In: *Journal de l'École polytechnique* 13 (1832), pp. 1–70. URL: <https://gallica.bnf.fr/ark:/12148/bpt6k4336778/f2.item.r=Joseph%20Liouville>.
- [LLC10] Wen-Hao Liu, Yih-Lang Li, and Hui-Chi Chen. “Minimizing clock latency range in robust clock tree synthesis.” In: *2010 15th Asia and South Pacific Design Automation Conference (ASP-DAC)*. 2010, pp. 389–394. DOI: 10.1109/ASPDAC.2010.5419849.
- [LM16] Christoph Lenzen and Moti Medina. “Efficient Metastability-Containing Gray Code 2-Sort.” In: *2016 22nd IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*. 2016, pp. 49–56. DOI: 10.1109/ASYNC.2016.18.

- [LW15] Peter Loeb and Manfred Wolff. “Nonstandard Analysis for the Working Mathematician.” 2nd ed. Springer, 2015. ISBN: 978-94-017-7326-3. DOI: 10.1007/978-1-4939-1194-3.
- [Man02] Ron Mancini, ed. “Op Amps For Everyone.” Texas Instruments, 2002. URL: <http://www.ti.com/lit/an/slod006b/slod006b.pdf>.
- [Mar81] Leonard R. Marino. “General theory of metastable operation.” In: *IEEE Transactions on Computers* C-30.2 (1981), pp. 107–115. DOI: 10.1109/TC.1981.6312173.
- [MC93] C.D. Motchenbacher and J.A. Conelly. “Low-Noise Electronic System Design.” Wiley, 1993. ISBN: 978-0-471-57742-3.
- [Mik78] Jan Mikusinski. “The Bochner Integral.” Springer, 1978. ISBN: 978-3-0348-5569-3. DOI: 10.1007/978-3-0348-5567-9.
- [MN68] Benoit B. Mandelbrot and John W. Van Ness. “Fractional Brownian Motions, Fractional Noises and Applications.” In: *SIAM Review* 10.4 (1968), pp. 422–437. DOI: 10.1137/1010093.
- [Mot+00] M. Mota et al. “A flexible multi-channel high-resolution time-to-digital converter ASIC.” In: *Nuclear Science Symposium Conference Record, IEEE*. Vol. 2. 2000, 9/155–9/159 vol.2. DOI: 10.1109/NSSMIC.2000.949889.
- [MSB12] Michael Mendler, Thomas R. Shiple, and Gérard Berry. “Constructive Boolean circuits and the exactness of timed ternary simulation.” In: *Formal Methods in System Design* 40.3 (2012), pp. 283–329. DOI: 10.1007/s10703-012-0144-6.
- [Nut68] Ronald Nutt. “Digital Time Intervalometer.” In: *Review of Scientific Instruments* 39.9 (June 1968), pp. 1342–1345. ISSN: 0013-5194. DOI: 10.1063/1.1683667.
- [Øks03] Bernt Øksendal. “Stochastic Differential Equations.” 6th ed. Springer, 2003. ISBN: 978-3-540-04758-2. DOI: 10.1007/978-3-642-14394-6.
- [OMS09] Keith Oldham, Jan Myland, and Jerome Spanier. “An Atlas of Functions.” 2nd ed. Springer, 2009. ISBN: 978-0-387-48806-6. DOI: 10.1007/978-0-387-48807-3.

- [PC12] L. Perktold and J. Christiansen. “A Flexible 5 ps Bin-Width Timing Core for Next Generation High-Energy-Physics Time-to-Digital Converter Applications.” In: *PRIME’12*. June 2012, pp. 1–4.
- [PC13] L. Perktold and J. Christiansen. “A fine time-resolution (< 3 ps-rms) time-to-digital converter for highly integrated designs.” In: *IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*. May 2013, pp. 1092–1097. DOI: 10.1109/I2MTC.2013.6555583.
- [Pet+09] E. Peters et al. “A deep-UV optical frequency comb at 205 nm.” In: *Opt. Express* 17.11 (May 2009), pp. 9183–9190. DOI: 10.1364/OE.17.009183.
- [PHS09] Thomas Polzer, Thomas Handl, and Andreas Steininger. “A Metastability-Free Multi-synchronous Communication Scheme for SoCs.” English. In: *Stabilization, Safety, and Security of Distributed Systems*. Ed. by Rachid Guerraoui and Franck Petit. Vol. 5873. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2009, pp. 578–592. ISBN: 978-3-642-05117-3. DOI: 10.1007/978-3-642-05118-0_40.
- [Pis16] Emilio Pisanty. “Unit relations in the new SI.” 2016. URL: https://commons.wikimedia.org/wiki/File:Unit_relations_in_the_new_SI.svg.
- [PM96] C.L. Portmann and T.H.Y. Meng. “Power-efficient metastability error reduction in CMOS flash A/D converters.” In: *IEEE Journal of Solid-State Circuits* 31.8 (1996), pp. 1132–1140. DOI: 10.1109/4.508260.
- [Pod99] Igor Podlubny. “Fractional Differential Equations.” Academic Press, 1999. ISBN: 0-12-558840-2.
- [Rad+13] Martin Radetzki et al. “Methods for Fault Tolerance in Networks-on-Chip.” In: *ACM Comput. Surv.* 46.1 (July 2013). ISSN: 0360-0300. DOI: 10.1145/2522968.2522976.
- [Rie47] Bernhard Riemann. “Versuch einer allgemeinen Auffassung der Integration und Differentiation.” 1847.
- [Rub12] Enrico Rubiola. “High resolution time and frequency counters.” In: *IEEE International Frequency Control Symposium (IFCS)*, IEEE. IEEE, 2012. URL: <http://rubiola.org/pdf-slides/2012T-IFCS-Counters.pdf>.

- [Rud76] Walter Rudin. “Principles of Mathematical Analysis.” McGraw-Hill, 1976. ISBN: 0-07-085613-3.
- [Rüf+18] Alain Rüfenacht et al. “Automated direct comparison of two cryocooled 10 volt programmable Josephson voltage standards.” In: *Metrologia* 55.4 (July 2018), pp. 585–596. DOI: 10.1088/1681-7575/aacbeb.
- [SAL06] Jian Song, Qi An, and Shubin Liu. “A high-resolution time-to-digital converter implemented in field-programmable-gate-arrays.” In: *Nuclear Science, IEEE Transactions on* 53.1 (Feb. 2006), pp. 236–241. ISSN: 0018-9499. DOI: 10.1109/TNS.2006.869820.
- [SC12] Xin-Wei Shih and Yao-Wen Chang. “Fast Timing-Model Independent Buffered Clock-Tree Synthesis.” In: *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 31.9 (Sept. 2012), pp. 1393–1404. ISSN: 0278-0070. DOI: 10.1109/TCAD.2012.2191554.
- [Sch18] W. Schottky. “Über spontane Stromschwankungen in verschiedenen Elektrizitätsleitern.” In: *Annalen der Physik* 362.23 (1918), pp. 541–567. DOI: 10.1002/andp.19183622304.
- [Sch26] W. Schottky. “Small-Shot Effect and Flicker Effect.” In: *Phys. Rev.* 28 (1926), pp. 74–103. DOI: 10.1103/PhysRev.28.74.
- [Sch87] Fred B. Schneider. “Understanding Protocols for Byzantine Clock Synchronization.” Tech. rep. Ithaca, NY, USA, 1987.
- [SG03] Yaron Semiat and Ran Ginosar. “Timing measurements of synchronization circuits.” In: *Ninth International Symposium on Asynchronous Circuits and Systems, 2003. Proceedings.* 2003, pp. 68–77. DOI: 10.1109/ASYNC.2003.1199167.
- [Sha+10] Guru Shamanna et al. “Scalable, sub-1W, sub-10ps clock skew, global clock distribution architecture for Intel® Core™ i7/i5/i3 microprocessors.” In: *2010 Symposium on VLSI Circuits.* 2010, pp. 83–84. DOI: 10.1109/VLSIC.2010.5560279.
- [Sim15] Barry Simon. “Real Analysis.” American Mathematical Society, 2015. ISBN: 978-1-4704-1099-5.

- [SJ16] Jeff A. Sherman and Robert Jördens. “Oscillator metrology with software defined radio.” In: *Review of Scientific Instruments* 87.5 (2016), p. 054711. DOI: 10.1063/1.4950898.
- [SL95] V M Sithi and S C Lim. “On the spectra of Riemann-Liouville fractional Brownian motion.” In: *Journal of Physics A: Mathematical and General* 28.11 (June 1995), pp. 2995–3003. DOI: 10.1088/0305-4470/28/11/005.
- [ST87] T. K. Srikanth and Sam Toueg. “Optimal Clock Synchronization.” In: *J. ACM* 34.3 (July 1987), pp. 626–645. ISSN: 0004-5411. DOI: 10.1145/28869.28876.
- [Str03] Robert S Strichartz. “A Guide to Distribution Theory and Fourier Transforms.” 1st ed. World Scientific, 2003. ISBN: 978-981-238-430-0.
- [Sul+90] D.B. Sullivan et al. “Characterization of Clocks and Oscillators.” NIST Technical Note 1337. National Institute of Standards and Technology, Mar. 1990. URL: <http://tf.nist.gov/general/pdf/868.pdf>.
- [TGL07] P. Teehan, M. Greenstreet, and G. Lemieux. “A Survey and Taxonomy of GALS Design Styles.” In: *Design Test of Computers, IEEE* 24.5 (Sept. 2007), pp. 418–428. ISSN: 0740-7475. DOI: 10.1109/MDT.2007.151.
- [TSG19] Ulrich Tietze, Christoph Schenk, and Eberhard Gamm. “Halbleiter-Schaltungstechnik.” 16th ed. Springer, 2019. ISBN: 978-3-662-48553-8. URL: <http://www.tietze-schenk.de>.
- [TY12] Ghaith Tarawneh and Alex Yakovlev. “An RTL method for hiding clock domain crossing latency.” In: *2012 19th IEEE International Conference on Electronics, Circuits, and Systems (ICECS 2012)*. 2012, pp. 540–543. DOI: 10.1109/ICECS.2012.6463557.
- [TYM14] Ghaith Tarawneh, Alex Yakovlev, and Terrence Mak. “Eliminating Synchronization Latency Using Sequenced Latching.” In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 22.2 (2014), pp. 408–419. DOI: 10.1109/TVLSI.2013.2243177.
- [Ung59] S. Unger. “Hazards and Delays in Asynchronous Sequential Switching Circuits.” In: *IRE Transactions on Circuit Theory* 6.1 (1959), pp. 12–25. DOI: 10.1109/TCT.1959.1086527.

- [Vas05] Gabriel Vasilescu. “Electronic Noise and Interfering Signals.” Springer, 2005. ISBN: 978-3-540-82915-7. URL: <https://www.springer.com/978-3-540-82915-7>.
- [Vee80] Harry JM Veendrick. “The Behaviour of Flip-Flops Used as Synchronizers and Prediction of their Failure Rate.” In: *Solid-State Circuits, IEEE Journal of* 15.2 (1980), pp. 169–176. DOI: 10.1109/JSSC.1980.1051359.
- [WJ +10] Frank W.J. Olver et al. “NIST Handbook of Mathematical Functions.” Cambridge University Press, 2010. ISBN: 978-0-521-14063-8. URL: <https://www.cambridge.org/de/academic/subjects/mathematics/abstract-analysis/nist-handbook-mathematical-functions>.
- [WL88] Jennifer Lundelius Welch and Nancy A. Lynch. “A New Fault-Tolerant Algorithm for Clock Synchronization.” In: *Information and Computation* 77.1 (1988), pp. 1–36. DOI: [https://doi.org/10.1016/0890-5401\(88\)90043-0](https://doi.org/10.1016/0890-5401(88)90043-0).
- [WMM16] Allen Waters, Jason Muhlestein, and Un-Ku Moon. “Analysis of Metastability Errors in Conventional, LSB-First, and Asynchronous SAR ADCs.” In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 63.11 (2016), pp. 1898–1909. DOI: 10.1109/TCSI.2016.2594256.
- [Xan09] Thucydides Xanthopoulos. “Clocking in Modern VLSI Systems.” Springer, 2009. ISBN: 978-1-4419-0260-3. DOI: 10.1007/978-1-4419-0261-0.