

Title	Interaction of additive and quantization noises in digital PLLs
Authors	Ó Tuama, Cillian
Publication date	2014
Original Citation	Ó Tuama, C. 2014. Interaction of additive and quantization noises in digital PLLs. PhD Thesis, University College Cork.
Type of publication	Doctoral thesis
Rights	© 2014, Cillian Ó Tuama - http://creativecommons.org/licenses/ by-nc-nd/3.0/
Download date	2024-09-01 05:12:01
Item downloaded from	https://hdl.handle.net/10468/1831



University College Cork, Ireland Coláiste na hOllscoile Corcaigh

Interaction of additive and quantization noises in digital PLLs

by

Cillian Ó Tuama

A thesis submitted to National University of Ireland, Cork for the degree of Doctor of Philosophy

DEPARTMENT OF APPLIED MATHEMATICS

Supervisor: Prof. James P. Gleeson

Head of Department: DR. KIERAN MULCHRONE

May 2014

Contents

A	Abstract				
A	ckno	wlegen	nents	6	
1	Intr	roducti	on	7	
2	Background theory				
	2.1	Proba	bility theory recap	11	
		2.1.1	Basics	12	
		2.1.2	Random variables and distribution functions	15	
		2.1.3	Density functions	18	
		2.1.4	Moments and expectations	22	
		2.1.5	Multivariate distributions and densities	26	
	2.2	Stocha	astic processes	30	
		2.2.1	Random vectors and random functions	30	
		2.2.2	Probability distribution for random processes	32	
		2.2.3	Basic properties of random processes	33	
		2.2.4	Power spectral density	37	
		2.2.5	Markov processes	43	
	2.3	Linear	filters	48	
		2.3.1	Analogue linear filters	49	
		2.3.2	Digital linear filters	56	
3	PLI	L theor	'y	60	
	3.1	Analog	gue PLL summary	61	
		3.1.1	Basics	61	

		3.1.2 Linear approximation	65
		3.1.3 Non-linear PLLs in the absence of noise \ldots .	70
		3.1.4 Additive noise in PLLs	79
		3.1.5 Additive noise in the linear model \ldots \ldots \ldots	82
		3.1.6 Analysis of additive noise in non-linear model \therefore	85
	3.2	Digital PLL summary	96
		3.2.1 Basics	98
		3.2.2 Linear approximation	105
		3.2.3 Quantization effects	107
		3.2.4 Circle rotation map preliminaries	116
		3.2.5 Circle rotation map and DPLL	118
	3.3	Summary	126
4	Results for the first-order DPLL without loop delay		128
	4.1	Circle map and the first-order DPLL with additive input	
		noise	129
	4.2	Numerical results for circle map and the first-order DPLL	
		with noise	133
	4.3	Analysis of circle map with additive noise	137
	4.4	Approximate solutions for steady-state	
		circle rotation map	144
		4.4.1 Large- L approximation	144
		4.4.2 Small- α approximation	147
	4.5	Summary	154
5	Res	sults for the first-order DPLL with loop delay	155
	5.1	Preliminaries	156
	5.2	Analysis	159
	5.3	Small- α results	165
	5.4	Summary	172
6	\mathbf{Res}	sults for the second-order DPLL	174
	6.1	Preliminaries	174
	6.2	Numerical results	176

	6.3	Analy	sis	178			
		6.3.1	Second-order limit cycles	178			
		6.3.2	Comparison to first-order case	183			
		6.3.3	Future work	186			
	6.4	Summ	ary	186			
7	Con	iclusio	ns and future work	188			
Bi	Bibliography						

I declare that this thesis is my own work, and has not been submitted for another degree, at National University of Ireland, Cork or any other institution.

Date

Abstract

Phase-locked loops (PLLs) are a crucial component in modern communications systems. Comprising of a phase-detector, linear filter, and controllable oscillator, they are widely used in radio receivers to retrieve the information content from remote signals. As such, they are capable of signal demodulation, phase and carrier recovery, frequency synthesis, and clock synchronization.

Continuous-time PLLs are a mature area of study, and have been covered in the literature since the early classical work by Viterbi [1] in the 1950s. With the rise of computing in recent decades, discrete-time digital PLLs (DPLLs) are a more recent discipline; most of the literature published dates from the 1990s onwards. Gardner [2] is a pioneer in this area.

It is our aim in this work to address the difficulties encountered by Gardner [3] in his investigation of the DPLL output phase-jitter where additive noise to the input signal is combined with frequency quantization in the local oscillator. The model we use in our novel analysis of the system is also applicable to another of the cases looked at by Gardner, that is the DPLL with a delay element integrated in the loop. This gives us the opportunity to look at this system in more detail, our analysis providing some unique insights into the variance 'dip' seen by Gardner in [3].

We initially provide background on the probability theory and stochastic processes. These branches of mathematics are the basis for the study of noisy analogue and digital PLLs. We give an overview of the classical analogue PLL theory as well as the background on both the digital PLL and circle map, referencing the model proposed by Teplinsky et al. [4, 5].

For our novel work, the case of the combined frequency quantization and noisy input from [3] is investigated first numerically, and then analytically as a Markov chain via its Chapman-Kolmogorov equation. The resulting delay equation for the steady-state jitter distribution is treated using two separate asymptotic analyses to obtain approximate solutions. It is shown how the variance obtained in each case matches well to the numerical results. Other properties of the output jitter, such as the mean, are also investigated. In this way, we arrive at a more complete understanding of the interaction between quantization and input noise in the first order DPLL than is possible using simulation alone.

We also do an asymptotic analysis of a particular case of the noisy first-order DPLL with delay, previously investigated by Gardner [3]. We show a unique feature of the simulation results, namely the variance 'dip' seen for certain levels of input noise, is explained by this analysis.

Finally, we look at the second-order DPLL with additive noise, using numerical simulations to see the effects of low levels of noise on the limit cycles. We show how these effects are similar to those seen in the noisefree loop with non-zero initial conditions.

Keywords: Nonlinear dynamics, phase jitter, phase-locked loops, noise, difference equations.

Acknowlegements

- I would like to thank my supervisor, Prof. James Gleeson for his patience and unwavering support over the many years I have spent compiling this thesis.
- Without the steady support from my wife, Marian, and my parents, this work would not have been possible.
- Helpful discussions with Stefan Tertinek and Orla Feely are gratefully acknowledged.
- This research is funded by Science Foundation Ireland under programmes 02/IN.1/IM062 and 06/IN.1/I366.
- Note that some of the background material in Chapter 3 and the numerical results in Section 4.2 have appeared previously in preliminary form in [6].

Chapter 1

Introduction

Phase-locked loops (PLLs) have been widely used for many years as components in electronics, and in communications systems in particular. A PLL is a closed loop consisting of a phase-detector, linear filter, and a voltage-controlled oscilllator (VCO) as described and analysed in detail in [1]. When arranged in the above series within the loop, the feedback signal entering the oscillator acts as an error signal, and drives the phase of the oscillator to match that of the input signal. Because of this property, PLLs are often used to synchronize clock signals across devices as well as to extract carrier signals in communications receivers. With some straightforward modifications to the circuit, a PLL can also be used as a frequency multiplier or divider.

Most modern communications systems are digital in nature, i.e., rely on discrete-time sampling of the incoming signal, which are then fed into the digital PLL (DPLL). In this case the PLL is simply a logic device, or perhaps even an algorithm operating on a digital computer, operating over discrete time-steps, and with discretized state values at each point in the loop. The VCO becomes a number-controlled oscillator (NCO) where the output frequency is determined by the control word x_q that is presented at the input. Since x_q is quantized, i.e., can only take on a finite number of values, the output frequency will also be quantized, and, in general, this will prevent the output signal from matching the input signal exactly. This quantization jitter was examined in detail in [4, 5]. A block diagram for the DPLL is shown in Figure 3.19. We assume, initially at least, a first-order DPLL as examined by Gardner in [3], with delay D = 1 and no downsampling of phase-detector samples (M = 1).

Gardner [3] examined numerically the case where the input signal was embedded in additive white Gaussian noise, and attempted to derive some rules of thumb to indicate how this additive noise interacted with the quantization jitter. In particular, expressions were found for the variance of the output jitter when no noise was added to the input signal. It was also noted in numerical simulations that, in the large noise limit, the variance of the output phase jitter was identical to what would be expected for an analogue PLL, i.e., without quantization. However, no general understanding of the overall problem was arrived upon.

This thesis examines the effect of the interaction between the input noise and the quantization jitter on the DPLL output using a number of asymptotic methods within the realm of stochastic calculus. We see that the DPLL operation in its steady-state is a circle rotation map, and show that the input noise affects the dynamics of this map in a unique and novel way. We also examine another case considered by Gardner, that is the DPLL with loop delay. We perform an analysis of this particular nonlinear dynamical system and derive results for a specific case using an asymptotic approximation. Finally, we look at the second-order DPLL and see how the limit cycles are disrupted by additive noise. We illustrate the similarities between the first- and second-order systems. In all cases, we reproduce earlier numerical results and show how our asymptotic solutions match to these results.

Note that some texts on PLLs [7] refer to the DPLL we examine in this work as an *all-digital PLL (ADPLL)*, reserving the DPLL acronym for a standard analogue PLL modified to use digital sampling of the input signal and a digital phase-detector. This derives from the early development of digital loops in the 1970s [7]; the first attempts at digitization took place in the phase-detector only, the rest of the loop following in later decades. In this thesis we use the term "DPLL" to refer to an all-digital loop as described above and is standard in [2], [3], [4], [5]. This thesis is organized as follows: Chapters 2 and 3 provide background material needed for an understanding of our results, which are presented from Chapter 4 onwards. Chapter 2 recaps some of the fundamental mathematical theory used throughout the thesis: probability theory, stochastic processes, linear systems and filter theory. In Chapter 3 we summarize the published literature in order to provide an introduction to both analogue and digital PLLs, deriving basic equations of operation, and behaviour of the loop under noisy conditions. We look at both the full non-linear loop equations as well as the classical linear approximation. In all cases we draw heavily on the theories introduced in Section 2, in particular linear filters and stochastic processes. In Chapter 3 we also introduce the basic properties of the circle map and then show that the DPLL output phase error obeys a circle map under certain conditions. Additionally, it is shown how the circle map parameters are related to those of the DPLL.

Our new results are presented from Chapter 4 onwards. Firstly, in Section 4.1, we write the circle map equation in a novel way and extend it to include the case corresponding to the DPLL in the presence of additive input noise. In Section 4.2 the main simulation results for the circle map output variance are presented and compared to corresponding plots for the DPLL. From this we can extrapolate ranges of validity for the approximation of the DPLL by the circle map. Next, in Section 4.3, by considering the circle map as a Markov process, we use a Chapman-Kolmogorov equation to derive a delay equation for the steady-state PDF of the circle map output for a given additive noise type N(t). Several properties of this PDF are derived, both for general N(t) and for noise with specific distributions. In Section 4.4 we use asymptotic analysis of this equation to obtain several approximate solutions, thus giving explicit results for the circle map output variance in the various ranges of interest of the DPLL parameters. In Chapter 5 the first-order DPLL with delay (D > 1 in [3]) is presented, and analyzed as a delayed circle map. In particular, the initial decrease in the output jitter variance with increasing additive noise noted in [3] is examined in detail. In our final results section, Chapter 6, we look at the second-order DPLL with additive input noise. Using numerical simulation we see how the limit cycles are broken up by the noise. We also note that some of the features of the new steady-state behaviour is similar to the noise-free loop with non-zero initial conditions, previously investigated in [4, 5]. Finally, conclusions and opportunities for future work are discussed in Chapter 7.

Both the background material in Section 3.2.4 and the numerical results in Section 4.2 have appeared previously in preliminary form in [6]. In addition, a special case of Theorem 7 on page 144 of this work, for $r = \frac{1}{2}$, appeared in [6]. This result is generalized here.

Chapter 2

Background theory

This chapter provides the background on the underlying mathematical material needed for the modelling of analogue and digital PLLs. We start in Section 2.1 with a recap of basic probability theory, taking in the concepts of random variables, distribution and density functions, moments and expectations, and the extension from scalar to multivariate random variables and densities. Section 2.2 then introduces the time-dependent random variable or stochastic process. We cover the important properties of random processes, such as autocorrelation, power spectral density, spectrum as well as introducing the concept of a Markov process. Finally in Section 2.3, we provide some material on analogue and digital linear filters, which are crucial to the operation of phase-locked loops. The material in this chapter serves to introduce Chapter 3, which provides the background material on analogue and digital PLLs.

2.1 Probability theory recap

In this section we provide an overview on the theory of probability, a fundamental branch of mathematics that is important for the understanding of later material. A useful introduction to probability theory for signal detection and processing in noisy environments can be found in Chapter 1 of [8]. We follow roughly the approach taken in that publication in the overview provided here. We start by introducing the concept of probability from first principles, before formalizing the approach with random variables, distribution and density functions. We go on to cover moments and expections (averages) as well as multivariable random variables and distributions.

2.1.1 Basics

Probability theory was devised [9] in order to formalize the notion of the 'likelihood' of a particular outcome of a random experiment, e.g., that 'heads' would be uppermost on a coin toss. For an unbiased coin it is found that for a large number of tosses, the ratio of the number of heads to the total number of tosses was generally very close to $\frac{1}{2}$. Thus, it could be supposed the the probability of getting heads (or tails) on a particular coin toss was $\frac{1}{2}$, on a scale where 0 represented certainty that the event would *not* occur, and 1 certainty that it would occur.

To introduce the concepts fundamental to probability theory we start with a set Ω , called the *sample space* of the random experiment. This is a set containing all possible *elementary outcomes* of the experiment. The elementary outcomes are denoted ζ_i for discrete (countable) outcomes, or $\zeta(\omega)$ for uncountable outcomes. In the case of a single coin toss, the elementary outcomes could be labelled $\zeta_1 = H$, $\zeta_2 = T$, and for the measurement (with infinite accuracy) of the phase of a random signal, $\zeta(\omega) = \Phi \in \mathbb{R}$. For the case of the coin toss the sample space could be $\Omega = \{H, T\}$. Similarly, for an experiment involving two coin tosses we could have $\Omega = \{HH, HT, TH, TT\}$.

The next component of probability theory is σ , the set of *events* of the random experiment, which is the set of all subsets of Ω . In the case of the double coin toss above, one such event would be $\{HT, TH, TT\}$, the event where at least one tail occurs.

In an experiment where we throw a die, we could have $\Omega = \{1, 2, 3, 4, 5, 6\}$. Then, $A = \{1, 3, 5\}$ is the event that an odd number is thrown, $B = \{2, 4, 6\}$ an even number, and $C = \{4, 5, 6\}$ a number greater than 3. The event that the number thrown is odd AND greater than 3 can be denoted $A \cap C$ and is $\{5\}$. Similarly, the event that the number is odd OR even is $A \cup B = \{1, 2, 3, 4, 5, 6\}$. This is just Ω and thus represents the *certain event*. The event that the number is odd AND even, $A \cap B$, is just the empty set, \emptyset , reflecting the sensible intuition that a number cannot be both odd and even. We define the *complement* (or *negation*) of an event $E \in \sigma$, \overline{E} , as the set of all outcomes in Ω that are not in E. In this example, taking E = C, we have $\overline{C} = \{1, 2, 3\}$, which is the set of outcomes where a number less than or equal to 3 is thrown. In general, for any event $E \subset \Omega$, $E \cup \overline{E} = \Omega$. Finally, we can define the *difference* between two events E_1 and E_2 , $E_1 \setminus E_2$, as the set of outcomes in E_1 that are *NOT* in E_2 , which is just $E_1 \cap \overline{E_2}$. In this example $A \setminus C = \{1, 3\}$, odd numbers that are *not* greater than 3.

Readers familiar with set or measure theory will recognise that σ as defined above represents a σ -algebra [10] on Ω in that it is closed under the operations of union and negation. That is, for any $E \in \sigma$, $\overline{E} \in \sigma$ also, and for any collection of events $E_i \in \sigma$, their union, $\bigcup_i E_i \in \sigma$.

The final component of probability theory is a real-valued function, P, that maps each element of σ to a value such that the following axioms are satisfied:

$$P(E) \ge 0 \ \forall E \in \sigma,$$

$$P(\Omega) = 1,$$

$$P(\cup_i E_i) = \sum_i P(E_i) \text{ for } E_i \in \sigma \text{ and } E_i \cap E_j = \emptyset \ \forall i \neq j.$$
(2.1)

The last axiom states that the probability of a union of events is simply the sum of the individual probabilities if the events are disjoint. From these axioms follow some of the basic properties of probability that will be familiar to many readers, such as:

$$P(\emptyset) = 0,$$

$$P(\overline{E}) = 1 - P(E).$$
(2.2)

In our die-throwing example above, for each event $E \in \sigma$ and denot-

ing the number of outcomes in E as |E|, suppose we define $P(E) = \frac{|E|}{6}$. The function P defined this way satisfies both the basic axioms of probability (check!) and also our intuitive notion of a 'fair die'. That is, $P(\{1,3,5\}) = P(\{2,4,6\}) = \frac{1}{2}$, so that, on average, an odd number will be thrown one half of the time and an even the other half. Also, $P(\{1\}) = \frac{1}{6}$, i.e., one throw in every 6 on average results in a 1. In fact, we could define an 'unfair die' as one having a probability function Pother than the one just described! In general, for a given random experiment $\{\Omega, \sigma, P\}$ we will define our probability function P to be such that, for an arbitrarily large number of runs of the experiment and for any fixed event $E \in \sigma$, the fraction of outcomes that satisfy event E will be approximately P(E).

Building on the basic concepts of probability already outlined, we can next define the *conditional probability* of the event E_1 given that event E_2 has occurred as:

$$P(E_1|E_2) = \frac{P(E_1 \cap E_2)}{P(E_2)}.$$
(2.3)

To see that this corresponds to our intuitive notion of conditionality, take $E_1 = A$ and $E_2 = C$ in our example above. Then, $P(A) = P(\{1,3,5\}) = \frac{1}{2}$ and $P(C) = P(\{4,5,6\}) = \frac{1}{2}$. Also, $P(A \cap C) = P(\{5\}) = \frac{1}{6}$. Then, equation (2.3) gives $P(A|C) = \frac{1}{6}/\frac{1}{2} = \frac{1}{3}$. This makes sense as C is the event $\{4,5,6\}$ and so if this has occurred, the sample space has effectively been reduced to this set, and the new probability of A occurring is $\frac{|\{5\}|}{|\{4,5,6\}|} = \frac{1}{3}$.

Two events $E_1, E_2 \in \sigma$ are called *independent* if

$$P(E_1 \cap E_2) = P(E_1)P(E_2), \tag{2.4}$$

which is to say that the probability of both E_1 and E_2 occurring is the product of the individual probabilities of each. From equation (2.3) above we have $P(E_1 \cap E_2) = P(E_1|E_2)P(E_2)$. From (2.4), if the events are independent, we then must have

$$P(E_1|E_2)P(E_2) = P(E_1)P(E_2).$$

This is just

$$P(E_1|E_2) = P(E_1), (2.5)$$

which says that the probability of event E_1 occurring is unaffected by the knowledge, or lack thereof, that event E_2 has occurred, which does correspond to our standard notion of independence. In the example above, we found $P(A) = P(\{1,3,5\}) = \frac{1}{2}$, $P(C) = P(\{4,5,6\}) = \frac{1}{2}$, but $P(A \cap C) = P(\{5\}) = \frac{1}{6}$. This means events A and C are not independent; knowledge, for example, that event C has occurred changes the probability that event A has also occurred.

2.1.2 Random variables and distribution functions

Probability distribution functions allow us to work analytically with the abstract concepts introduced in Section 2.1.1 above, e.g., to represent graphically or analytically the relative probabilities of the occurrence of various events. In particular, we consider the probability that the outcome of an experiment is 'less than' a particular value.

To put this on a firm footing, we introduce the concept of a random variable. This is a real-valued function $X(\zeta)$ on the elementary outcomes in the sample space Ω . For example, for our single die-throw experiment we could just set $X_i = X(\zeta_i)$, where $X_i = i \in \mathbb{R}$. Then, we can consider the event, for example, that the result of the die-throw is less than or equal to four, $D = \{\zeta : X(\zeta) \le 4\} = \{1, 2, 3, 4\}$. Then, for a fair die, $P(X \le 4) = \frac{2}{3}$.

In general we define the *(cumulative)* distribution function of the random variable X as

$$P_X(x) = P(X \le x). \tag{2.6}$$

Note that in equation (2.6), X is a real-valued function on the elementary outcomes ζ and, for each real value x, the inequality $(X \leq x)$ represents an event $E \in \sigma$. We sometimes abbreviate the term "distribution function" to simply "distribution" or *CDF*.



Figure 2.1: Distribution function $P_X(x)$ for the random variable X = i in the throw of a fair die.

For the single throw of a fair die, the distribution function is shown in Figure 2.1 for the random variable X = i, where *i* is the outcome of the throw. A number of characteristics are noticeable:

- For x < 1, P_X(x) = 0. This reflects the fact that the result of a die throw cannot be less than 1. In general, for the distribution function of a single, real-valued random variable, P_X, we have lim_{x→-∞} P_X(x) = 0. This follows directly from the basic axioms of probability 2.1 and 2.2, and the fact that distribution functions are based on probabilities as in Equation 2.6.
- For $x \ge 6$, $P_X(x) = 1$. This follows from the fact that a die throw always shows a number less than or equal to 6. For the distribution function of a single real-valued random variable we generally have $\lim_{x\to\infty} P_X(x) = 1$.
- For this particular distribution, P(X ≤ 3) = ¹/₂ so that P_X(3) = ¹/₂. Also for any x in the range 3 ≤ x < 4, P(x) = ¹/₂. This is because, for example, the events (X ≤ 3) and (X < 3.9) are exactly the same, i.e., ζ ∈ {1,2,3}. The regions between each integer value of x in Figure 2.1 are constant for this same reason.

Another characteristic of distributions functions is that they are monotonically non-decreasing, from a value of 0 at $-\infty$ to 1 at $+\infty$. This is another consequence of the axioms of probability. To see this, consider the quantity $P_X(x + dx)$, which is:

$$P_X(x + dx) = P(\{\zeta : X \le x\} \cup \{\zeta : x < X \le x + dx\})$$

= $P(\{\zeta : X \le x\}) + P(\{\zeta : x < X \le x + dx\})$ (2.7)
 $\ge P(\{\zeta : X \le x\}) = P_X(x).$

Here we have used the axioms that state that $P_X(x) \ge 0$ always, and that the probability of a union of disjoint events is the sum of the individual probabilities of each event. The manipulations in (2.7) also reveal another feature of distribution functions:

$$P(\{\zeta : a_1 < X \le a_2\}) = P(\{\zeta : X \le a_2\}) - P(\{\zeta : X \le a_1\})$$

= $P_X(a_2) - P_X(a_a).$ (2.8)



Figure 2.2: Distribution function $P_X(x)$ for a continuous random variable X.

2.1.3 Density functions

The distribution function shown in Figure 2.1 is piecewise continuous, and of constant value wherever it is continuous; it has a number of discontinuities and its derivative does not exist everywhere. This is because the underlying random variable is discrete rather than continuous. Many of the random variables we deal with in this thesis are continuous and have CDFs of the type shown in Figure 2.2. For such a distribution, the derivative exists everywhere. Since the CDF is a non-decreasing function, its derivative is greater than or equal to zero always.

From the distribution shown in Figure 2.2 it is clear that the probability of the random variable X being less than or equal to 2 is quite small, since $P_X(2) \ll 1$. Also, the probability that X > 4 is $P_X(+\infty) - P_X(4) =$ $1 - P_X(4)$ which is also much less than 1. Therefore, on a series of experiments with outcome X one would expect a result between 2 and 4 to be more likely than values outside of this range. In particular, $P_X(x)$ is increasing most rapidly around x = 3 so one would expect most of the probability to be concentrated around this value.



Figure 2.3: Density function $p_X(x)$ for the continuous random variable X with CDF shown in Figure 2.2.

The derivative of the CDF shown in Figure 2.2 is plotted in Figure 2.3, and bears this out: the probability is most densely distributed around x = 3. In general, the derivative of the distribution function is known as the *(probability) density function* or *PDF* of the random variable X.

By the definition of the density function we have

$$P_X(x) = P(X \le x) = \int_{-\infty}^x p_X(x') dx', \qquad (2.9)$$

which means that

$$P(a_1 < X \le a_2) = P_X(a_2) - P_X(a_1) = \int_a^b p_X(x) dx.$$
 (2.10)

This implies that

$$p_X(x)dx = P(x < X \le x + dx).$$
 (2.11)

This last equation represents a basic property of PDFs for continuous random variables: the quantity $p_X(x)dx$ is the probability that the random variable X lies in the interval (x, x + dx]. The value of any $p_X(x)$ represents a probability density rather an absolute probability value in the range [0, 1]. In fact, for a continuous random variable, the probability that $X = a_1$ for any arbitrary value a_1 is always zero. This can be seen by setting $x = a_1$ and dx = 0 in equation (2.11). This may seem counterintuitive until the following example is considered: consider an experiment where a point is chosen at random on the real line in the interval [0, 1] with a uniform distribution over this interval, i.e., the probability of the point being in any interval of length dx < 1 is exactly dx. Suppose the distance of the point from the origin is measured with infinite accuracy. Then, the probability that the point is at exactly 0.5 (or any other particular value in [0, 1]) must be zero. If it weren't, then it would not be possible to satisfy the axioms of probability: the relationship $\sum_{i} P(E_i) < 1$ could not hold for any collection of the *i* disjoint events $X = a_i$ as the sum taken over *i* could be an infinite number of finite values.

Another important property of PDFs is

$$\int_{-\infty}^{\infty} p_X(x) dx = P_X(\infty) = 1, \qquad (2.12)$$

i.e., the area under the PDF is unity always.



Figure 2.4: Representation of density function $p_X(x)$ for the discrete random variable X with CDF shown in Figure 2.1. Each impulse represents a Dirac delta function of weight $\frac{1}{6}$.

For discrete random variables, such as the die-throw experiment with CDF shown in Figure 2.1, the PDF, defined as the derivative of the CDF, does not not exist. However, it can be formally defined as:

$$p_X(x) = \sum_i P_i \delta(x - x_i), \qquad (2.13)$$

where the x_i are the points of discontinuity on the CDF and the P_i are the absolute probabilities of the occurrence of each $X = x_i$, i.e.,

$$P_i = \lim_{h \to 0} \{ P_X(x_i) - P_X(x_i - h) \}.$$

For the die toss experiment $\{x_i\} = \{1, 2, 3, 4, 5, 6\}$ and $P_i = \frac{1}{6} \forall i$. A representation of the density function is shown in Figure 2.4. Note that in equation (2.13) $\delta(x)$ is the standard Dirac delta function.¹

The discrete PDF as defined in (2.13) is more usually referred to as a

¹The Dirac delta is a generalized function satisfying $\int_{-\infty}^{\infty} f(x)\delta(x-a)dx = f(a)$ for every continuous f(x). In particular $\int_{-\infty}^{\infty} \delta(x)dx = 1$. $\delta(x)$ is zero everywhere except at x = 0 where it is undefined. It may be thought of as the limit of a series of tall, narrow functions at x = 0 with constant integral. No true function exists with these properties, but for most purposes we can use the Dirac delta as though it were a standard function. It was originally introduced as a notation by Dirac in his 1930 work [12].

probability mass function; see Chapter 26 of [11], for example. It satisfies two of the basic properties of general PDFs, i.e., $\int_{-\infty}^{\infty} p_X(x) dx = 1$ and $\int_{-\infty}^{x} p_X(x') dx' = P_X(x)$. A random variable may also be of mixed type, its PDF containing both non-constant continuous sections as well as delta functions. In this case we can write for the PDF

$$p_X(x) = \frac{dP_X^c}{dx} + \sum_i P_i \delta(x - x_i), \qquad (2.14)$$

where P_X^c is the continuous part of the distribution, undefined at the isolated points of discontinuity, x_i , and the train of delta functions are as in (2.13).



Figure 2.5: Wheel-spin experiment with mixed random variable result.

As an example of an experiment with a mixed random variable result, consider a disc with unit radius with a pointer anchored at the centre point, as shown in Figure 2.5. The pointer is spun around so that it can eventually come to rest pointing at a random spot on the circumference. If the pointer is spun fast enough it may land at any point on the circumference with equal probability. Thus the random variable X, the angular displacement from 0 of the pointer after each spin, would be uniformly distributed on the interval $[0, 2\pi]$. Therefore, the PDF of X is a constant $p_X \equiv 1/2\pi$ on the interval $[0, 2\pi]$. Note that because X is a continuous random variable on the real line segment $[0, 2\pi]$ the probability of the pointer landing at any particular point with offset x (X = x) is zero. If we now modify the experiment so that we add a groove or other device at $x = \pi$, so that, on average, one spin in every ten the pointer will come to rest at exactly $X = \pi$, then there is now a non-zero probability that the pointer will land at $x = \pi$, i.e., $P(X = \pi) = 0.1$ and $P(X \neq \pi) = 0.9$. So now the PDF for X is changed so that it is a uniform density with height $0.9/2\pi$ at all points in $[0, 2\pi]$ except at $x = \pi$, where it is a delta function with weight 0.1, as shown in Figure 2.6.



Figure 2.6: PDF for the grooved wheel experiment, where X is the angular displacement along the circumference of the wheel where the pointer comes to rest.

2.1.4 Moments and expectations

The probabilities for all outcomes of a particular random experiment are described fully by either the CDF or PDF, detailed in Section 2.1.3. However, we are often interested in more basic properties, such as the average value of the experiment's outcome, or how large the spread of outcomes is. For example, in the die-toss experiment, we would expect the average value of X, the value facing upward on a die-throw, to be around 3.5. We would also expect the values to be spread evenly spread across the set $\{1, 2, 3, 4, 5, 6\}$. Note that the terms "average" and "expectation" used here are *ensemble averages* and *ensemble expectations*, in that they are the values predicted a priori, using the known probability distribution of the experiment's outcome. Over a sufficiently large number of runs of the experiment, we would expect the measured value of the average to be close to the expected value.

First we define the *moments* of a probability distribution for a random variable X as

$$m_n = \mathscr{E}(X^n) = \int_{-\infty}^{\infty} x^n p_X(x) dx$$
 for $n = 0, 1, 2, \dots$ (2.15)

Note that $m_0 = \int_{-\infty}^{\infty} p_X(x) dx = 1$ for all PDFs p_X . We also define the *central moments* of a distribution as

$$\mu_n = \mathscr{E}\left[(X - \mathscr{E}(X))^n \right] = \int_{-\infty}^{\infty} (x - m_1)^n p_X(x) dx \text{ for } n = 0, 1, 2, \dots$$
(2.16)

Clearly, $\mu_0 = 1$ and $\mu_1 = 0$. Next we define the *mean* as

$$\mathscr{E}(X) = \int_{-\infty}^{\infty} x p_X(x) dx = m_1.$$
(2.17)

The mean is also referred to as the ensemble average or expected value of the random variable X. To see that this corresponds to our usual notion of an average consider again the die-toss experiment, for which the PDF is plotted in Figure 2.4 and is $\sum_{i=1}^{6} \frac{1}{6}\delta(x-i)$. Then, the mean of the distribution is

$$\mathscr{E}(X) = \int_{-\infty}^{\infty} x \sum_{i=1}^{6} \frac{1}{6} \delta(x-i) dx = \frac{1}{6} \sum_{i=1}^{6} i = 3.5.$$
(2.18)

[Aside: Generalizing this, it can easily be seen that a discrete-uniform distribution of equal-weight delta functions over the first n natural numbers, $\{1, 2, ..., n\}$ has mean $\frac{1}{2}(n+1)$. This result is used later in Section 3.2.4.]

We next define the *variance* of a distribution as the second central moment,

$$\mu_2 = \sigma^2 = \mathscr{E}\left[\left(X - \mathscr{E}(X) \right)^2 \right].$$
(2.19)

Expanding and taking expectations inside the brackets, this can also be written as $\sigma^2 = \mathscr{E}(X^2) - (\mathscr{E}(X))^2$. The square root of the variance, σ , is known as the *standard deviation*. Both σ and σ^2 are measures of how widely spread a distribution is around its mean.

Using a similar approach to above, the variance of the discrete-uniform distribution over the first n natural numbers can be calculated as

$$\sigma^2 = \mathscr{E}(X^2) - (\mathscr{E}(X))^2 = \frac{1}{n} \sum_{i=1}^n i^2 - \left[\frac{1}{2}(n+1)\right]^2 = \frac{n^2 - 1}{12}.$$
 (2.20)

Again, for the die-toss this is just $\frac{35}{12}$.

Consider now the continuous analog of the die-toss random variable: the (continuous) uniform distribution on the real line segment [1, L] for L > 1. This distribution has

$$p_X(x) = \begin{cases} \frac{1}{L-1} & \text{for } x \in [1, L] \\ 0 & \text{for } x \notin [1, L] \end{cases}.$$

This means that the random variable X is equally likely to lie within any sub-segment of length dx contained fully within [1, L]. The mean of the mapping is

$$\int_{1}^{L} x \frac{1}{L-1} dx = \frac{L+1}{2},$$
(2.21)

while the variance is

$$\int_{1}^{L} x^{2} \frac{1}{L-1} dx - \left[\frac{L+1}{2}\right]^{2} = \frac{1}{3}(L^{2}+L+1) - \left[\frac{L+1}{2}\right]^{2} = \frac{(L-1)^{2}}{12}.$$
(2.22)

Comparing equations (2.18) and (2.20) to these last two, it can be seen that the discrete-uniform and continuous-uniform distributions over a given range have the same means and variances; a discrete-uniform distribution for an *L*-sided die over values $\{1, 2, ..., L\}$ also has mean (L+1)/2 and variance $(L-1)^2/12$. We will use both discrete and continuousuniform distributions when considering the effects of (continuous-uniform density) additive noise on a (discrete-uniform density) circle rotation map in Section 4.1.

As a final example consider the Gaussian (or normal) density:

$$p_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$
 (2.23)

The first two moments of this distribution can be calculated analytically using elementary methods:

$$m_1 = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} x e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx = \mu,$$
$$m_2 = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} x^2 e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx = \sigma^2 + \mu^2$$

That is, the mean of the distribution is μ and the variance is σ^2 . Note that the distribution is defined fully by its mean and variance; all higher moments will be functions of the first two. The Gaussian density given by (2.23) is often written simply as $\mathcal{N}(\mu, \sigma^2)$. The PDF shown in Figure 2.3 is a Gaussian with mean 3 and variance $\frac{1}{2}$. The CDF is

$$P_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^x e^{-\frac{(y-\mu)^2}{2\sigma^2}} dy,$$

and cannot be expressed in terms of elementary functions. It must instead be computed numerically for each value of x or represented using the *error function*, defined as

$$\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-y^2} dy$$

In this case the CDF can be written as

$$P_X(x) = \frac{1}{2\sigma} \left[1 + \operatorname{erf}\left(\frac{x-\mu}{\sqrt{2\sigma^2}}\right) \right].$$

The Gaussian CDF is that plotted in Figure 2.2.

Figure 2.7 shows several Gaussian distributions, each with the same mean, but with different variances. It is evident that the smaller the variance, the more 'bunched' the probability about the mean, while for



Figure 2.7: PDFs for Gaussian distributions with mean 3 and variance (a) 0.1 (small dash), (b) 0.5 (solid), (c) 2 (large dash).

large variances there is a greater probability that the random variable will be found further from the mean value. For the uniform distribution of length (L-1), we found in equation (2.22) that the variance is $\frac{(L-1)^2}{12}$, so that, again, larger variances imply a larger range of the random variable (about the mean). This is true of probability distributions in general.

2.1.5 Multivariate distributions and densities

From a single random experiment we can derive any number of random variables. In particular suppose that we have two random variables, $X(\zeta)$ and $Y(\zeta)$. Then we can define the *joint (bivariate) distribution function* as

$$P_{XY}(x,y) = P\left(\{\zeta : X \le x\} \cap \{\zeta : Y \le y\}\right) = P\left(X \le x, Y \le y\right).$$
(2.24)

Clearly, this can be extended to any number of variables for the multivariate distribution

$$P_{X_1...X_n}(x_1,...,x_n) = P(X_1 \le x_1,...,X_n \le x_n).$$
 (2.25)

Similar to the single-variable case, we can define a *multivariate probability density function* for distributions that are continuous everywhere as

$$p_{X_1\dots X_n}(x_1,\dots,x_n) = \frac{\partial^n P_{X_1\dots X_n}}{\partial x_1\dots \partial x_n}.$$
(2.26)

From this, just as in the single-dimensional case, we can derive absolute probabilities from the density function. So, for example, the probability that the random variable X is within a certain range AND the variable Y is within another range is

$$P(a_1 < x \le a_2, \ b_1 < y \le b_2) = \int_{a_1}^{a_2} \int_{b_1}^{b_2} p_{XY}(x, y) dy \, dx.$$
 (2.27)

The most prominent example found in applications is the *multivariate* Gaussian, which is just the multi-dimensional analog of equation (2.23):

$$p_{X_1...X_k}(x_1,...,x_k) = p_{\mathbf{X}}(\mathbf{x})$$

= $\frac{1}{(2\pi)^{k/2} |\mathbf{\Sigma}|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \mathbf{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right).$
(2.28)

Here, Σ is the *covariance matrix* of \mathbf{x} and $|\Sigma|$ is its determinant. Note that the quantity inside the exponential involves matrix multiplication, and Σ is a symmetric $k \times k$ positive definite matrix. For the bivariate case we have

$$\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}, \ \boldsymbol{\mu} = \begin{pmatrix} \mu_x \\ \mu_y \end{pmatrix}, \text{ and } \mathbf{\Sigma} = \begin{pmatrix} \sigma_x^2 & \rho \sigma_x \sigma_y \\ \rho \sigma_x \sigma_y & \sigma_y^2 \end{pmatrix},$$

where μ_x , μ_y , σ_x^2 , σ_y^2 are the means and variances of X and Y respectively, and ρ is their *correlation*. In this case equation (2.28) for the bivariate Gaussian becomes:

$$p_{XY}(x,y) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}} \exp\left(-\frac{1}{2(1-\rho^2)} \left[\frac{(x-\mu_x)^2}{\sigma_x^2} + \frac{(y-\mu_y)^2}{\sigma_y^2} - \frac{2\rho(x-\mu_x)(y-\mu_y)}{\sigma_x\sigma_y}\right]\right).$$
(2.29)

We can relate the multivariate distribution and density functions to

those for a single variable by noting, in the bivariate case for example, that

$$P_X(x) = P_{XY}(x, +\infty) = \int_{-\infty}^x \int_{-\infty}^\infty p_{XY}(x', y') dy' dx'.$$
 (2.30)

From this follows the marginal density

$$p_X(x) = \frac{dP_X(x)}{dx} = \int_{-\infty}^{\infty} p_{XY}(x, y')dy',$$
 (2.31)

and similarly for y,

$$p_Y(y) = \frac{dP_Y(y)}{dy} = \int_{-\infty}^{\infty} p_{XY}(x', y) dx'.$$
 (2.32)

For multivariate densities we are often interested in getting the statistics of one variable conditioned on the others. That is, fixing one variable, say Y = y, we would like to find the resulting density for the other variable X. To do this, we start with equation (2.3) and take E_1 as the event X = x and E_2 as $y < Y \le y + dy$. Then, from equation (2.3) we have

$$P_{X|Y}(x \mid y < Y \le y + dy) = \frac{P(X \le x, y < Y \le y + dy)}{P(y < Y \le dy)}$$
$$= \frac{P_{XY}(x, y + dy) - P_{XY}(x, y)}{P_Y(y + dy) - P_Y(y))}.$$

Now if we let $dy \to 0$ we get

$$P_{X|Y}(x|Y=y) = \frac{\partial P_{XY}(x|Y=y)/\partial y}{p_Y(y)}.$$

Finally, taking partial derivatives with respect to x on each side, this becomes

$$p_{X|Y}(x \mid Y = y) = \frac{p_{XY}(x, y)}{p_Y(y)}.$$
(2.33)

Two random variables X and Y are said to be *independent* if the events $x < X \leq x + dx$ and $y < Y \leq y + dy$ are independent (as in

equation (2.4) for all x, y. The joint probability in question is

$$P(x < X \le x + dx, y < Y \le y + dy) = P_{XY}(x, y) \, dx \, dy,$$

where we assume that $dx, dy \to 0$. However, since the events are independent, we have from (2.4):

$$P(x < X \le x + dx, y < Y \le y + dy) = P(x < X \le x + dx)P(y < Y \le y + dy)$$

Again, under the assumption that $dx, dy \to 0$, this last is just

$$P_X(x) dx P_Y(y) dy,$$

so that we have

$$P_{XY}(x,y) = P_X(x)P_Y(y),$$
 (2.34)

which is the analog of equation (2.4) for random variables. Taking partial derivatives $\frac{\partial^2}{\partial x \partial y}$ of each side, we get the equivalent relation for densities:

$$p_{XY}(x,y) = p_X(x)p_Y(y).$$
 (2.35)

Now, for independent random variables equation (2.33) becomes

$$p_{X|Y}(x \mid Y = y) = p_X(x).$$
(2.36)

This is just the analog of equation (2.5) for random variables, and says that the probability density for X = x is not affected by knowledge that a particular value of Y has occurred, as would be intuitively expected for independent variables.

From Section 2.2 onwards, we frequently dispense with some of the cumbersome notation used above when dealing with probabilities in applications. So, for example, we use p(x) to mean $p_X(x)$, and p(x|y) for $p_{X|Y}(x|Y = y)$. For CDFs we often write, for example, P(x) as shorthand for $P_X(x) = P(X \le x)$. In each case, all references to the random variable X are suppressed to simplify the notation; however this does

result in the P (or p) symbol being overloaded somewhat. In general, it should be clear from the context which function is involved and what the underlying random variable is. This notation is standard in almost all texts on probability theory and its applications.

2.2 Stochastic processes

In this section we extend the basic concepts of probability theory to processes that are ongoing in time. The random variables of Section 2.1 become random functions or random vectors, and quantities like mean, variance, and moments become time-dependent. A dynamical system, such a a phase-locked loop, can often be represented by a difference equation for a discrete-time process, or a differential equation for a continuous-time system. Often of interest in the engineering world is where the operation of these systems is dependent on some randomness, either from external noises or inherent in the system. Thus the state variables of the system become time-dependent random variables, or stochastic processes. We again follow the approach taken in [8], where the material on stochastic processes can be found in Chapter 2. Another useful reference on stochastic processes, particularly for stochastic differential equations, is the very readable work by C.W. Gardiner [13].

2.2.1 Random vectors and random functions

In Section 2.1 we had a random experiment with sample space Ω and elementary outcomes ζ . Random variables were then real values $X(\zeta)$ on which probability distributions were based. In many cases we need to deal with time processes that are random, such as a dynamical system described by the solution x(t) of the differential equation

$$\frac{dx(t)}{dt} = f\left[x(t), \eta(t)\right],$$

where $\eta(t)$ is a random noise process. In this scenario, we can consider a random function $X(t,\zeta)$, where ζ is the outcome of the experiment, or simply the outcome nature has chosen by chance. Each outcome ζ determines a standard function of time, $X(t,\zeta)$, called a *realization* or *sample function* of the random, or stochastic, process. The collection of all such time functions for all possible ζ is called the *ensemble* of sample functions. Indeed, the stochastic process itself can be defined as $X(t,\zeta)$, the ensemble of all possible sample functions of the process.

Often we deal with discrete-time processes. These may be sampled versions of the continuous time functions described above, or processes that are inherently time-discrete such as a system described by a difference equation in time:

$$x(t_i) = x_i = f[x(t_{i-1}), \eta(t_{i-1})]$$

In this example, the value of the scalar process x at a particular time is a function of x and the value of the random noise process at the previous time step. In general, the evolution of the process x_i may be a function of all previous x_i as well as all previous noise values, i.e.,

$$x(t_i) = f[x(t_{i-1}), x(t_{i-2}), \dots, x(t_0), \eta(t_{i-1}), \eta(t_{i-2}), \dots, \eta(t_0)].$$

In either case, each outcome of the chance experiment, ζ , defines a multidimensional random variable $\mathbf{X}(\zeta) = \{X_i(\zeta)\}$, where each $X_i(\zeta)$ is the value of the random process at time instant t_i . Hence, each realization of the discrete-time random process can be thought of as a random vector $\mathbf{X}(\zeta)$, and the random process itself is the ensemble $X(t_i, \zeta)$.

In both discrete and continuous time, as with basic random variables, the realizations of a random process may be finite, countably infinite, or uncountably infinite. In the case of the first two, the outcomes of the random experiment may be denoted ζ_j , so the random process may be written as $X(t, \zeta_j)$ or $X(t_i, \zeta_j)$. In general, where the process takes place in the analogue domain its values tend to be real numbers and thus the number of possible realizations is uncountably infinite. For example, an analogue signal in a noisy communications channel may be received as $X(t,\zeta)$ at the analogue front-end, and sampled thereafter as $X(t_i,\zeta)$. In both cases the number of possible ζ is uncountably infinite, under the assumption that the measurement of X can be performed with infinite precision at the receiver.

2.2.2 Probability distribution for random processes

At any particular time t (or $t = t_i$) a random process $X(t, \zeta)$ has a value x(t) that is determined by the particular realization, ζ . That is, fixing a particular $t = \tau$, $X(\tau, \zeta)$ can be treated as a random variable as in Section 2.1. As such, $X(\tau, \zeta)$ has a CDF, PDF, moments etc. in the same way as for any simple random variable.

For stochastic processes we usually use a shorthand to simplify the notation, and write just x(t) instead of $X(t,\zeta)$ on the understanding that x(t) is a random quantity for each t. We can also write p(x(t),t), or simply p(x,t), for the density function of the process at time t. Similarly, the distribution function of the process at time t can be denoted P(x,t). Finally, since for each fixed $t = \tau$, $x(\tau)$ is just a random variable (denoted X in Section 2.1), we can just write p(x) and P(x) for the density and distribution functions of simple random variables for the remainder of this thesis.

As with random variables, we can now define the time-dependent moments of a random process as

$$m_n(t) = \mathscr{E}(x(t)) = \int_{-\infty}^{\infty} [x(t)]^n p(x,t) dx.$$
(2.37)

Similarly, the central moments are

$$\mu_n(t) = \mathscr{E}\left[(x(t) - \mathscr{E}(x(t)))^n \right] = \int_{-\infty}^{\infty} (x(t) - m_1(t))^n p(x, t) dx, \quad (2.38)$$

and, as before, the mean of the process is the first moment $m(t) = m_1(t)$

and the variance is $\sigma(t) = \mu_2(t)$.

For both continuous-time and discrete-time random processes, we are often interested in the statistics of the process at more than one time instance. So, for two particular time instances, we have a joint PDF $p(x_1, x_2; t_1, t_2)$, which should be read as "the probability density for the process having value x_1 at time t_1 and value x_2 at time t_2 ". This is just a standard joint probability distribution as in equation (2.26) and can be extended to n time instants $p(x_1, \ldots, x_n)$, where we omit the time indexes for brevity. All of the usual properties of multivariate distributions apply, such as

$$\int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} p(x_1, \dots, x_n) dx_1 \dots dx_n = 1, \qquad (2.39)$$

$$\int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} p(x_1, \dots, x_n) dx_{k+1} \dots dx_n = p(x_1, \dots, x_k).$$
(2.40)

2.2.3 Basic properties of random processes

For a stochastic or random process, how the statistics of the process varies over time is of primary importance. Of particular interest is how the value of the process at one time instant is related to its value at another, if at all.

As with multivariate distributions in general, the values of a random process over a certain time interval may be *independent* if its values over another; in particular if

$$p(x_1, \dots, x_n) = p(x_1, \dots, x_k)p(x_{k+1}, \dots, x_n),$$
 (2.41)

then it can be said that the process over the time interval $\{x_{k+1}, \ldots, x_n\}$ is independent of its values over $\{x_1, \ldots, x_k\}$. Equation (2.41) is equivalent to

$$p(x_{k+1},\ldots,x_n \,|\, x_1,\ldots,x_k) = p(x_{k+1},\ldots,x_n). \tag{2.42}$$



Figure 2.8: Sample realizations for random processes consisting of (a) random real numbers in the range [-1, 1], and (b) the cumulative sum of random real numbers in the range [-1, 1].

Next we define the *autocorrelation function* of the process x(t) as

$$R_x(t_1, t_2) = \mathscr{E}\left[x(t_1)x(t_2)\right] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_1 x_2 \, p(x_1, x_2; t_1, t_2) dx_1 \, dx_2.$$
(2.43)

An important property of a stochastic process is whether its statistics change over time. In particular, if

$$p(x_1, \dots, x_n; t_1, \dots, t_n) = p(x_1, \dots, x_n; t_1 - \tau, \dots, t_n - \tau)$$
(2.44)

holds true for all possible $n, x_1, \ldots, x_n, t_1, \ldots, t_n$, and $\tau \in \mathbb{R}$, then the process x(t) is said to be *(strictly) stationary*. The mean and variance of stationary processes are constant for all time, and the statistics of the process across multiple time instants is dependent only on the difference between the time instants, i.e., is independent of the absolute times, or position of the origin. For such a process, if we take two time instants, t_1 and $t_2 = t_1 - \tau$, then we have the joint density

$$p(x_1, x_2; t_1, t_2) = p(x_1, x_2; t_1, t_1 - \tau) = p(x_1, x_2; \tau),$$

so that equation (2.43) becomes

$$R_x(t_1, t_1 - \tau) = \mathscr{E}[x(t_1)x(t_1 - \tau)] = R_x(\tau), \qquad (2.45)$$
where we have used the definition of stationarity, (2.44). $R_x(\tau)$ can equivalently be written as $R_x(t_1 - t_2)$. This means that the autocorrelation is a function only of the time difference rather than any absolute times. In Figure 2.8 sections from realizations for two different random processes are plotted; the first is a stationary process while the second is non-stationary.

For real processes found in applications, strict stationarity is quite a difficult property to satisfy or prove in many cases. It is often reasonable to use a more limited definition of stationarity that still leads to simplifications in the analysis. A process is said to be *wide-sense stationary* if its mean is constant and its autocorrelation is a function only of the time difference, i.e., $\mathscr{E}[x(t)] = m_x$ and $R_x(t_1, t_2) = R_x(t_1 - t_2)$. Clearly, from (2.44), all moments and the mean of a strictly stationary process are constant, and from the result for the autocorrelation in (2.45), it can be seen that processes that are strictly stationary are also wide-sense stationary, as would be expected. However, the converse is not necessarily true.

For any general random process with autocorrelation function $R_x(t_1, t_2)$ then the variance is

$$\sigma_x^2(t) = \mathscr{E}\left[(x(t) - m(t))^2\right] = \mathscr{E}\left[x^2(t)\right] - m^2(t) = R_x(t, t) - m^2(t).$$
(2.46)

For a stationary or wide-sense stationary process, this simplifies significantly to

$$\sigma_x^2 = R_x(0) - m_x^2, \tag{2.47}$$

where $m_x = m(t)$ is the constant mean value of x(t).

Some random processes we deal with in this thesis have the property that

$$R_x(t_1, t_2) = 0 \quad \forall t_1 \neq t_2$$

Any process that has such a property is said to be *uncorrelated*. In particular, we often have, for a continuous-time process, $R_x(t_1, t_2) \approx \delta(t_1-t_2)$, or for a wide-sense stationary continuous process, $R_x(\tau) \approx \delta(\tau)$.

These are called *white noise processes* for which more details are given in Section 2.2.4. The discrete-time equivalent of the continuous-time white noise is $R_x(t_i, t_j) = \delta_{ij}$, where δ_{ij} is the *Kronecker delta*, which has a value of 1 for i = j and 0 otherwise.

If the values of a zero-mean process at two time instants t_1 and t_2 are independent then they are also uncorrelated, since

$$R_x(t_1, t_2) = \mathscr{E}\left[x(t_1)x(t_2)\right] = \mathscr{E}\left[x(t_1)\right]\mathscr{E}\left[x(t_2)\right] = 0.$$

However, the converse is not necessarily true; the values of a process at two points in time may be uncorrelated without being independent.

Often when dealing with random processes we only have a single, sample realization available for analysis. For example, we may have a single capture of a random waveform, x(t), of arbitrary duration. In this situation we may calculate a time average over a duration T of the realization as $\frac{1}{T} \int_0^T x(t) dt$. One might expect that if T is taken to be large enough we might always have, for a process we assume to be wide-sense stationary:

$$\langle x \rangle = \lim_{T \to \infty} \left(\frac{1}{T} \right) \int_0^T x(t) dt = \mathscr{E} \left[x(t) \right] = m_x.$$
 (2.48)

That is to say, the average of a single realization over a sufficiently long period of time is equal to the ensemble average for the process. In practice, this is not always the case, and is often difficult to prove one way or the other. It is often simply assumed in the case where only a single sample function is available. This assumption is reasonable in many cases, particularly for processes that are expected to be stationary; multiple time segments of a single sample function can double as separate realizations of the same process. Processes for which equation (2.48) hold exactly are said to be *mean ergodic*. If the ensemble average of any functional of the process is equal to the limit of the time average, then

the process is *ergodic*. That is, for any h(x) we have

$$\mathscr{E}\left[h\left(x(t)\right)\right] = \left\langle h\left(x(t)\right)\right\rangle = \lim_{T \to \infty} \left(\frac{1}{T}\right) \int_0^T h\left(x(t)\right) dt.$$
(2.49)

2.2.4 Power spectral density

Translation to the frequency-domain is often instructive in the analysis of random processes. More insight may be gained into the properties of a process by visualization in the frequency domain rather than in the time domain in many cases. Frequency-domain analysis generally starts with the Fourier transform, which for a realization of the random process x(t)is

$$\tilde{X}(\omega) = \int_{-\infty}^{\infty} x(t)e^{-i\omega t}dt.$$
(2.50)

 $\tilde{X}(\omega)$ is a complex quantity and exists provided

$$\int_{-\infty}^{\infty} |x(t)| dt < \infty.$$
(2.51)

To visualise how the components of x(t) are distributed across the frequency range, ω , we would generally plot the magnitude or squaremagnitude of this, e.g.,

$$S(\omega) = \left| \tilde{X}(\omega) \right|^2 = \left| \int_{-\infty}^{\infty} x(t) e^{-i\omega t} dt \right|^2.$$
 (2.52)

The issue with this approach is that for the real random processes we are interested in, condition (2.51) is rarely met; our processes generally continue as $t \to \infty$ with average amplitude that does not decay to zero. If we take instead the requirement that the average power of the process is finite

$$\lim_{T \to \infty} \frac{1}{T} \int_{-T/2}^{T/2} |x(t)|^2 dt < \infty,$$
(2.53)

we can consider a frequency-domain representation of the process as

$$\tilde{X}_T(f) = \int_{-T/2}^{T/2} x(t) e^{-2\pi i f t} dt.$$
(2.54)

Here $f = \omega/2\pi$. This is a truncated Fourier transform of a realization of the process, x(t), and is a random variable for each T. This leads to the definition of the truncated power spectral density of the process as

$$S_T(f) = \mathscr{E}\left[\frac{1}{T} \left| \tilde{X}_T(f) \right|^2 \right].$$
(2.55)

We would like to be able to define the *power spectral density* (*PSD*) of the process, x(t), as

$$S_x(f) = \lim_{T \to \infty} S_T(f). \tag{2.56}$$

However, without further analysis, it isn't clear that the limit (2.56) exists, or what the range of convergence in f might be.

The Wiener-Khintchin theorem [14] states that, for a wide-sense stationary process subject to condition (2.53), the limit always exists for all f, and identifies its value as

$$S_x(f) = \int_{-\infty}^{\infty} R_x(\tau) e^{-2\pi i f \tau} d\tau, \qquad (2.57)$$

where $R_x(\tau)$ is the autocorrelation function of the process. That is to say, the Fourier transform of the autocorrelation function gives the power spectral density of the process as defined by (2.56).

A detailed proof of the Wiener-Khintchin can be found in many texts on signal analysis and communications theory. A sketch of the proof from [14] goes as follows: *Proof (sketch).* Firstly, note that

$$\mathscr{E} \left| \tilde{X}_{T}(f) \right|^{2} = \mathscr{E} \left| \int_{-T/2}^{T/2} x(t) e^{-i\omega t} dt \right|^{2}$$

$$= \mathscr{E} \int_{-T/2}^{T/2} \int_{-T/2}^{T/2} x(t) x(\tau) e^{-i\omega(t-\tau)} dt d\tau$$

$$= \int_{-T/2}^{T/2} \int_{-T/2}^{T/2} R_{x}(t-\tau) e^{-i\omega(t-\tau)} dt d\tau$$

$$= \int_{-T}^{T} (T - |\tau|) R_{x}(\tau) e^{-i\omega\tau} d\tau,$$

(2.58)

where this last step follows after some elementary calculus. This leads to

$$\mathscr{E}\left[\frac{1}{T}\left|\tilde{X}_{T}(f)\right|^{2}\right] = \int_{-\infty}^{\infty} R_{x,T}(\tau)e^{-i\omega\tau}d\tau, \qquad (2.59)$$

which looks almost like the result we need except that we have the modified autocorrelation function

$$R_{x,T}(\tau) = \begin{cases} \left(1 - \frac{|\tau|}{T}\right) R_x(\tau) & \text{if } |\tau| < T\\ 0 & \text{if } |\tau| \ge T \end{cases}$$
(2.60)

As we let $T \to \infty$ across equation (2.59), the functions $R_{x,T}$ become a sequence of functions that converge (pointwise) to R_x and such that

$$|R_{x,T}(\tau)| \le |R_x(\tau)| \ \forall T.$$

In this case we can use Lebesgue's dominated convergence theorem [15] to interchange the limiting and integration operations in the following to obtain the result:

$$S_x(f) = \lim_{T \to \infty} \mathscr{E} \left[\frac{1}{T} \left| \tilde{X}_T(f) \right|^2 \right] = \lim_{T \to \infty} \int_{-\infty}^{\infty} R_{x,T}(\tau) e^{-i\omega\tau} d\tau$$
$$= \int_{-\infty}^{\infty} R_x(\tau) e^{-i\omega\tau} d\tau.$$
(2.61)

Г		
		-

Clearly, the relation (2.57) is invertible so that the autocorrelation function may be obtained as the inverse Fourier transform of the PSD where the latter is known. In particular, the first two moments are related to the PSD according to

$$R_x(0) = \sigma_x^2 + m_x^2 = \int_{-\infty}^{\infty} S_x(f) df.$$
 (2.62)



 $R_x(\tau)$

(a) Sample realization of random process x(t) with small correlation time.



(b) Narrow Gaussian autocorrelation function.





The concept of white noise was introduced briefly in Section 2.2.3. A zero-mean, wide-sense stationary continuous white noise process has autocorrelation function $R_x(\tau) = \delta(\tau)$, where $\delta(\tau)$ is the Dirac delta function as introduced in Section 2.1.3. Inserting into (2.57), this yields S(f) = 1, i.e., the frequency response is flat across all frequencies. This is where the term "white" originates; white light contains an equal energy distributions across all frequencies in the visible spectrum.

In practice, true continuous-time white noise is impossible to synthe-

sise or observe. For example, equation (2.62) yields

$$\delta(0) = \sigma_x^2 = \infty_z$$

which says that the variance of the process is infinite, a consequence of having constant energy across the (infinitely large) frequency domain. However, white noise is a useful concept, and is often approximated by real processes with a sufficiently short correlation time, for example, processes with autocorrelation function

$$R_x(\tau) = \frac{1}{m\sqrt{\pi}} e^{-\tau^2/m^2},$$
(2.63)

for sufficiently small m. Normally we refer to noise as "white" if its observable spectral density is flat, or it is flat across the bandwidth of the medium in question, e.g., the PSD of white light is flat across the portion of the electromagnetic (EM) spectrum that is in the visual band. However, it decays to zero away from this band.

As an example of an approximate while noise process, we can consider a real random process with autocorrelation function given by (2.63) with m = 0.03, which is plotted in Figure 2.9(b). It may have time-domain realization similar to the process shown in Figure 2.8(a), shown again in Figure 2.9(a). Then, using (2.57) the PSD for the process is found to be

$$S_x(\omega) = \frac{1}{2\pi} e^{-\frac{1}{2}m^2\omega^2}.$$

So, while the autocorrelation is a very narrow Gaussian curve centred at $\tau = 0$, the PSD is a very wide Gaussian centred at $\omega = 2\pi f = 0$. As shown in Figure 2.9(c), this is clearly not the constant PSD required for true white noise. However, when viewed over a sufficiently narrow frequency range centred at $\omega = 0$ it appears approximately flat, as seen in Figure 2.9(d).

Often, the level of additive white noise in a system is specified by its one-sided spectral density, N_0 . This means simply that, for the real noise process n(t), its spectral density is given by $S_n(\omega) = N_0/2$. Alternatively,

if we take n(t) to be complex with $S_n(\omega) = 0 \ \forall \omega < 0$, then $S_n(\omega) = N_0$ for $\omega \ge 0$. While true continuous-time white noise does not exist, an approximate white noise process with spectral density N_0 within the band of interest, and negligible power outside this band, has variance that is proportional to N_0 . This is a consequence of (2.62), which gives, for a zero-mean process band-limited to $f \in [B_1, B_2]$

$$\sigma_x^2 \approx \int_{B_1}^{B_2} \frac{N_0}{2} df = \frac{B_2 - B_1}{2} N_0.$$
 (2.64)

Thus, in many cases in this thesis where we use the terms "level" or "intensity" of 'white noise' in a system, we may be referring to either the variance, σ_x^2 , or (constant) spectral density, N_0 , of the approximate white noise process. The actual measurement in question should generally be clear from the context.

It should be noted that specification of the autocorrelation function or PSD of a random process imposes no restrictions on its probability density. In Section 2.2.2 it was observed that at any time t, the value of the process x(t) is a random variable and thus has a probability distribution and density. In many cases we consider, this density is Gaussian, as in equation (2.23). In fact, very often we assume that noise added to a system is, at least approximately, a white, Gaussian stochastic process. Such noise is called *additive white Gaussian noise (AWGN)*. This assumption is necessary, for example, in the derivation of the Fokker-Planck equation in Section 2.2.5.

It is equally possible, notwithstanding the above, to have non-white (correlated, coloured) Gaussian noise. It is also possible to have white, non-Gaussian noise, e.g., uncorrelated noise samples from a uniform density on [-1, 1]. We will have occasion to use such a noise source later in this thesis, in Chapter 5. Finally, noise, or a general process, may be Gaussian and non-stationary, in which case the PSD doesn't even exist and spectral considerations don't arise. The process x(t) shown in Figure

2.8(b), for example, has a probability density that is approximately

$$p(x,t) = \mathcal{N}(0,t/3) = \frac{1}{\sqrt{2\pi t/3}} e^{-\frac{x^2}{2t/3}}.$$
(2.65)

This is a Gaussian distribution with variance that increases linearly with time. Thus, the energy of the process grows unbounded as $t \to \infty$, equation (2.53) is violated and the PSD is not defined.

2.2.5 Markov processes

In general, the future evolution of a random process depends (in a probabilistic sense) on all its previous values. For example, if it is known that the process had values x_i at times t_i for i = 1, ..., n, in general the probability density of the process at future times will be a function of these past values, namely

$$p(x_{n+1}; t_{n+1} | x_n, \dots, x_1; t_n \dots, t_1).$$

However, many real random processes have the property that the future evolution depends only on the current value, rather than all previous values, i.e., the process has no 'memory'. This may be expressed as follows:

$$p(x_{n+1}; t_{n+1} | x_n, \dots, x_1; t_n, \dots, t_1) = p(x_{n+1}; t_{n+1} | x_n; t_n).$$
(2.66)

Another way of interpreting this is that that future development of the process can be predicted (in a probabilistic sense) from the current state alone. Such processes are said to have the *Markov property* and their analysis is greatly simplified as compared to general random processes. A discrete-time process with the Markov property is known as a *Markov chain*. The probability distribution in equation (2.66) is known as the *transition probability* at time t_n ; the function describes the probability of transitioning from value x_n at time t_n to the value x_{n+1} at a later time t_{n+1} .

Suppose next that $\{X_i : i = 1, 2, ...\}$ is a series of *independent identically-distributed* (*i.i.d.*) random variables. That is, $P_{X_j}(x) = P_{X_k}(x) \forall j, k, \forall x$, and $P(x_j | X_k = x_k) = P(x_j) \forall j, k$. Now if we construct a discrete-time sequence

$$x(t) = x_t = X_t$$
 for $t = 1, 2, \dots$

then the process is called a *(discrete-time) white noise process.* A sample realization from such a process is plotted in Figure 2.8(a). Clearly, the process is Markov; in fact, not only does the future evolution of the process not depend on past values, but it doesn't even depend on the current value! That is, the value of the process at any time is dependent only on a probability distribution that is constant over time.

Now let us take another random process, y(t), where y(t) = y(t-1) + x(t), and x(t) as above. Then we have a process that is the summation of a series of i.i.d. random variables. A sample plot of such a process is shown in Figure 2.8(b). This belongs to a class of stochastic processes with *independent increments* and thus is a Markov chain since the distribution of each value of the process is dependent only on its previous value, as well as the distribution of each increment. As it happens, in this case the distribution of each increment is identical, but in general the distribution may vary over time. For a process to have independent increments, the increments need only be independent of the past values of the process. By construction, all processes with independent increments are Markov. Note finally that, while the increments y(t) - y(t-1) = x(t)are i.i.d., the random variables $Y_t = y(t)$ are in no way i.i.d.! Indeed, one might guess this from a perusal of Figure 2.8(b).

It is typical for Markov chains to be defined via their increments. For example, the random process

$$w(t+1) = \begin{cases} w(t) + 1 & \text{with probability } p \\ w(t) - 1 & \text{with probability } 1 - p \end{cases} \quad \text{for } t = 1, 2, \dots \quad (2.67)$$

is a Markov chain known as a random walk. If $p = \frac{1}{2}$ it is called a symmetric random walk. As constructed, the random walk process has independent increments. In fact, the increments are i.i.d. with probability density $p(x) = p \,\delta(x-1) + (1-p)\delta(x+1)$. Random walks are a good model for many real processes, such as stock market movements, particle dispersion and are also used in search algorithms on computer networks [16].



Figure 2.10: Sample start, intermediate and end points for a Markov process x(t).

Suppose that a Markov process has values x_s , x_i , and x_e at discrete time instants t_s , t_i , and t_e respectively, as illustrated in Figure 2.10. Then, the joint probability density is

$$p(x_e, x_i, x_s) = p(x_e \mid x_i, x_s)p(x_i, x_s).$$

Here, the time indexes have been omitted for brevity but it should be understood that the density functions refer to the process having each value at the corresponding time index. Because the process has the Markov property, the density at any time depends on the density at the previous known time only, and so the expression for the joint distribution can be simplified to

$$p(x_e, x_i, x_s) = p(x_e \mid x_i)p(x_i, x_s)$$

Finally, if we divide across by $p(x_s)$ and integrate both sides over x_i , we

get

$$p(x_e \mid x_s) = \int_{-\infty}^{\infty} p(x_e \mid x_i) p(x_i \mid x_s) \, dx_i.$$
 (2.68)

This equation states that the transition probability from x_s at time t_s to x_e at t_e can be found by choosing an arbitrary time between t_s and t_e and integrating the product of the individual transition probabilities across all possible values of the process at that intermediate time. Equation (2.68) is known variously as the *Smoluchowski equation* or *Chapman-Kolmogorov equation* [17] and is effectively an integral equation for the time evolution of p(x, t) given an initial condition.

Let us next assume that an initial condition for the process is given, i.e., that at t = 0 we have $x = x_0$ so that the initial condition on its density is

$$p(x;0 \mid x_0;0) = \delta(x - x_0).$$
(2.69)

For brevity we now use the notation p(x,t) for the transition probability from the initial condition, i.e., p(x,t) is shorthand for $p(x;t | x_0; 0)$, so that (2.69) can be written more simply as

$$p(x,0) = \delta(x - x_0). \tag{2.70}$$

Now, taking equation (2.68) as a starting point we can obtain, using some non-trivial manipulations, a partial differential equation for p(x, t):

$$\frac{\partial p(x,t)}{\partial t} = \sum_{n=1}^{\infty} \frac{(-1)^n}{n!} \frac{\partial^n}{\partial x^n} \left[A_n(x) p(x,t) \right].$$
(2.71)

In this, each A_n is the growth rate of the *n*th moment of the process x(t). That is,

$$A_n(x) = \lim_{\Delta t \to 0} \int_{-\infty}^{\infty} (x' - x)^n p(x'; t + \Delta t \mid x; t) dx'$$

=
$$\lim_{\Delta t \to 0} \frac{\mathscr{E}\left[(\Delta x)^n \mid x\right]}{\Delta t}.$$
 (2.72)

In (2.72), the density and expectation are taken for the increment of the

process from x at time t to x' at time $t + \Delta t$.

It can be shown (see [13], for example) that, under mild assumptions about the process x(t), namely that it has continuous paths, all but the first two moments, A_1 and A_2 , in equation (2.71) are zero, i.e., the resulting process x(t) is slowly-varying enough that moments higher than the second vanish more quickly than Δt . In this case (2.71) reduces to the *Fokker-Planck equation*, also called the *differential Chapman-Kolmogorov* equation in [13]:

$$\frac{\partial p(x,t)}{\partial t} = -\frac{\partial}{\partial x} \left[A_1(x) p(x,t) \right] + \frac{1}{2} \frac{\partial^2}{\partial x^2} \left[A_2(x) p(x,t) \right].$$
(2.73)

A full derivation of the Fokker-Planck equation as well as methods of solution are given in [18]. This text also deals with the multi-dimensional Fokker-Planck equation and its applications in the physical sciences.

Random processes often result from a deterministic system with an additional driving noise source n(t). We will encounter such situations later in this thesis. These systems are often modelled as differential equations with a noise term, known as a *Langevin equation* [19], for example

$$\frac{dx}{dt} = a(x,t) + b(x,t)n(t).$$
 (2.74)

In the theory of *stochastic differential equations (SDEs)* this is interpreted as the SDE:

$$dx(t) = a[x(t), t]dt + b[x(t), t]dW(t),$$

which is shorthand for

$$x(t) = x(0) + \int_0^t a[x(\tau), \tau] d\tau + \int_0^t b[x(\tau), \tau] dW(\tau).$$
 (2.75)

In these, dW(t) represents the increment of a Wiener noise process [20], and the last term in (2.75) is either an Itô or Stratonovich *stochastic integral*, depending on the interpretation. Both Itô and Stratonovich interpretations of the Langevin equation as an SDE are self-consistent – the ambiguity arises because of the way the idealized white noise process enters the Langevin equation (2.74) – and the theory used in a particular case generally depends on the application. Itô calculus is most often used for proofs and rigorous theory, financial mathematics, cases where the noise is indeed assumed to be white, and where b[x(t), t] is constant. For a good grounding in SDE theory using Itô calculus, see [21]. The Stratonovich interpretation naturally arises in physical applications, where white noise is considered as the limit of a coloured (correlated) noise as the correlation time vanishes to zero. Stratonovich SDEs are generally used in engineering applications, for coloured noise sources, and where b[x(t), t] is not constant ('multiplicative noise'). Stratonovich theory is well covered in many texts on random processes in the physical sciences, see, for example, [22].

Using either Itô or Stratonovich interpretations, a similar Fokker-Planck equation as in (2.73) can be derived ([13], [18], [19], [20]) as was found using the Markov assumption. Moreover, the quantities $A_1(x)$ and $A_2(x)$ in (2.73) can be found in terms of the Langevin coefficients a[x,t] and b[x,t] in (2.74). Finally, a Fokker-Planck equation found using the Stratonovich interpretation can relatively easily be transformed to an equivalent Itô version and vice versa; the equations are identical in the case where b[x,t] is constant, but differ in the case of so-called multiplicative noise.

2.3 Linear filters

Linear filters are of great importance in communications systems, phaselocked loops, and dynamical systems in general. Many systems are naturally linear, or designed to be linear for easy analysis and predictable behaviour. More complicated systems are often linear or approximately linear over a limited range of parameters of interest, thus making them amenable to analysis under certain conditions.

This section is intended only as an introduction to the basics of linear

filters. There are several engineering texts covering this extensive subject in more detail; [23], for example, is a commonly-used reference for electronic engineers.

2.3.1 Analogue linear filters

If x(t) is the input to a linear system, the output y(t) is related to it according to

$$y(t) = \int_{-\infty}^{\infty} h(t,\tau) x(\tau) d\tau.$$
 (2.76)

Here, for example, the linear system may be a linear radio-frequency (RF) amplifier, x(t) and y(t) the time-varying voltages of the input and output RF signals. In equation (2.76) the function $h(t, \tau)$ is known as the *impulse response* of the linear system in that, if the input signal is the Dirac delta function applied at time t_1 , then $x(t) = \delta(t - t_1)$ and

$$y(t) = \int_{-\infty}^{\infty} h(t,\tau)\delta(\tau - t_1)d\tau = h(t,t_1).$$
 (2.77)

That is, the output of the system in response to the impulse at t_1 is $h(t, t_1)$.

If we denote the action of the linear system by the operator \mathscr{F} , equation (2.76) above can be written more compactly as $y(t) = \mathscr{F}\{x(t)\}$. Clearly \mathscr{F} is a linear operator in that $\mathscr{F}\{ax_1(t)+bx_2(t)\}=a\mathscr{F}\{x_1(t)\}+b\mathscr{F}\{x_2(t)\}$ for inputs $x_1(t), x_2(t)$, and $a, b \in \mathbb{R}$.

In the usual case, and all cases we will be concerned with in this thesis, the response at any time of the linear system to an impulse depends only on the difference between the current time and the time of the impulse. In this case the system is said to be *linear time-invariant (LTI)* with impulse response $h(\tau)$ and (2.76) may be rewritten as the convolution integral

$$y(t) = \int_0^t h(t - \tau) x(\tau) d\tau.$$
 (2.78)

Here, we have restricted all signals to be defined for t > 0 only. It can easily be seen that the impulse response is sufficient to define the action of an LTI filter if we consider the following: any input x(t) can be written as $x(t) = \int_0^t \delta(t - \tau) x(\tau) d\tau$. Then,

$$y(t) = \mathscr{F}\left\{\int_0^t \delta(t-\tau)x(\tau)d\tau\right\} = \int_0^t \mathscr{F}\left\{\delta(t-\tau)\right\}x(\tau)d\tau$$

=
$$\int_0^t h(t-\tau)x(\tau)d\tau.$$
 (2.79)

The Laplace transform of a function z(t) is defined as

$$Z(s) = \int_0^\infty e^{-st} z(t) dt,$$

where $s \in \mathbb{C}$. The transform is usually interpreted as an operator that converts a *time-domain* function of t to a complex *frequency-domain* function of s. This is simply a generalization of the Fourier transform from equation (2.50). The relationship between the two is straightforward: $\tilde{Z}(\omega) = Z(i\omega)$.

We can take Laplace transforms of equation (2.78) to obtain an equation of operation for the linear system in the Laplace domain:

$$Y(s) = H(s)X(s).$$
 (2.80)

Note: we can equivalently use Fourier transforms to obtain

$$\tilde{Y}(\omega) = \tilde{H}(\omega)\tilde{X}(\omega),$$

for $\omega \in \mathbb{R}$ in the frequency domain.] In this equation, H(s) [or $H(\omega)$] is known as the *transfer function* or *frequency response* of the filter.

It is often of interest to be able to describe the action of an LTI filter on a random process. For example, when we write the action of a filter as

$$y(t) = \int_0^t h(t-\tau)x(\tau)d\tau,$$

the functions x(t) and y(t) in question may be either deterministic, or sample functions of a random process. We can take expectations across this equation to find that, in the case that x is wide-sense stationary, for example

$$\mathscr{E}\left[y(t)\right] = \int_0^t h(t-\tau)\mathscr{E}\left[x(\tau)\right] d\tau = H(0)m_x = m_y.$$
(2.81)

That is, the mean of the output signal is a multiple of the (constant) mean of the input signal. Here, H(0) is the transfer function of the filter evaluated at the zero frequency (also known as DC or direct-current frequency). For a wide-sense stationary input process there is a simple relationship between the power spectral density of the steady-state filtered output signal and input signal. To see this, first take the general case of a non-stationary input signal x(t). Then, the filtered output signal y(t) has autocorrelation function

$$R_{y}(t,t+\tau) = \mathscr{E}\left\{y(t)y(t+\tau)\right\}$$

= $\mathscr{E}\left\{\int_{0}^{t}\int_{0}^{t+\tau}x(t-t_{1})x(t+\tau-t_{2})h(t_{1})h(t_{2}) dt_{2} dt_{1}\right\}$
= $\int_{0}^{t}\int_{0}^{t+\tau}R_{x}(t-t_{1},t-t_{2}+\tau)h(t_{1})h(t_{2}) dt_{2} dt_{1}.$
(2.82)

If we now take the case where x(t) is wide-sense stationary and let $t \to \infty$, then the right-hand side of (2.82) becomes independent of t, so the the autocorrelation becomes

$$R_y(\tau) = \int_0^\infty \int_0^\infty R_x(\tau + t_1 - t_2)h(t_1)h(t_2) dt_2 dt_1.$$
(2.83)

Thus, from (2.81) and (2.83), for a wide-sense stationary input, the steady-state output of an LTI filter is also wide-sense stationary. Note that this is not necessarily true if the input has been driving the output for only a finite time. The PSD of the steady-state output can now be found by inserting (2.83) into equation (2.57):

$$S_{y}(f) = \int_{-\infty}^{\infty} R_{y}(\tau)e^{-i\omega\tau}d\tau$$

$$= \int_{-\infty}^{\infty} \int_{0}^{\infty} \int_{0}^{\infty} R_{x}(\tau + t_{1} - t_{2})h(t_{1})h(t_{2})e^{-i\omega\tau}dt_{2}dt_{1}d\tau$$

$$= \int_{-\infty}^{\infty} \int_{0}^{\infty} \int_{0}^{\infty} R_{x}(\tau + t_{1} - t_{2})e^{-i\omega(\tau + t_{1} - t_{2})}e^{i\omega t_{1}}e^{-i\omega t_{2}}$$

$$h(t_{1})h(t_{2})dt_{2}dt_{1}d\tau$$

$$= \int_{-\infty}^{\infty} R_{x}(\tau_{1})e^{-i\omega\tau_{1}}d\tau_{1}\int_{0}^{\infty} h(t_{1})e^{-i\omega t_{1}}dt_{1}\int_{0}^{\infty} h(t_{2})e^{-i\omega t_{2}}dt_{2}$$

$$= |\tilde{H}(\omega)|^{2}S_{x}(f).$$
(2.84)



Figure 2.11

 $\frac{1}{10.0}\omega$

Most analysis and characterization of LTI filters takes place in the frequency domain. In particular, certain types of transfer functions are common, for example functions that decay to zero as $|\omega| \to \infty$ [23]. These are known as *low-pass filters* in that they pass signals of low frequency with little attenuation, but block high-frequency components. They are frequently used for noise reduction in the event where a signal of interest is disturbed by random fluctuations; these fluctuations are often modelled as high-frequency additive noise. For example, the signal of interest may be a low-frequency sinusoid or modulated carrier with most of its energy centred at a single frequency. Even if the noise is taken to be white, with its energy spread across a wide range of frequencies as in Figure 2.9(d), most of this energy will be outside the frequency band of the signal. Then, a low-pass filter may be used to eliminate most of noise, while passing the signal unaltered plus any small noise components in its frequency band. As such, the magnitude of the fluctuations will be greatly reduced.



Figure 2.12: Operation of filters from Figure 2.11 on input signal (a) $x(t) = \sin t + \sin 10t$ to produce output signals (b) $y(t) \approx (1/e) \sin t$, and (c) $y(t) \approx \sin 10t$.

As an idealized example of this, consider a filter with transfer function

$$\tilde{H}(\omega) = \exp[-\omega^2],$$

as shown in Figure 2.11(a). This is a function that decays quickly to zero outside the range $|\omega| > 2$. By inspection of the transfer function and considering that the filter operates in the frequency domain according to equation (2.80), one might expect that this filter would block signals, or

components of signals, that have frequencies outside the range $|\omega| \leq 2$, the *passband* of the filter. If we now take the input signal to be

$$x(t) = \sin t + \sin 10t,$$

as shown in Figure 2.12(a), then the filter's input signal has frequency components at $|\omega| = 1$ and $|\omega| = 10$ only. Indeed, we have

$$\tilde{X}(\omega) \sim \delta(\omega \pm 1) + \delta(\omega \pm 10).$$

Given that the $|\omega| = 10$ component is well outside the passband of the filter we would expect to see this to be eliminated from the filter's output, y(t). In fact, it is found that when simulating this in the time domain that the output is

$$y(t) = \frac{1}{e}\sin t + \frac{1}{e^{100}}\sin 10t$$

which is plotted in Figure 2.12(b). Thus, the high-frequency component has been attenuated by a factor of e^{100} , and effectively eliminated. The low-frequency component passes unaltered, though attenuated by a factor of e. A frequency-domain representation of the process is shown in Figure 2.13(a); the delta spike at $\omega = 10$ has essentially been eliminated, while the tone at $\omega = 1$ has simply been attenuated.

In general low-pass filters are designed to pass the signal of interest with an attenuation factor of as close to unity as possible, while eliminating as effectively as possible frequencies outside the signal's bandwidth. The ideal low-pass filter is one with a rectangular shape, with a vertical cutoff at the edge of the frequency band of interest. This type of filter is impossible to realise in practice, though it is possible to synthesize filters that approach this behaviour. Low pass filters are commonly used in phase-locked loops.

Another important type of low-pass filter is an *(ideal) integrator*, that

is a filter that has an output equal to the integral of the input signal:

$$y(t) = \int_{-\infty}^{t} x(\tau) d\tau.$$
 (2.85)

For this filter, the response to an impulse applied at time t = 0, $x(t) = \delta(t)$, will be a constant $h(t) = y(t) \equiv 1$ for t > 0. Taking a Laplace transform yields the frequency response

$$H(s) = \frac{1}{s},\tag{2.86}$$

and frequency magnitude response

$$|H(i\omega)| = \frac{1}{\omega}.\tag{2.87}$$

Clearly this is similar to the low-pass filter of Figure 2.11(a) in that the response of the filter falls off for large ω , so an integrator will similarly allow through slow variations in the input signal but eliminate high-frequency oscillations.

A second category of filter that is commonly used is a high-pass filter. As its name suggests, it eliminates slow variations in the incoming signal and allows through the rapidly-changing components. A sample transfer function, $\tilde{H}(\omega) = \exp[-(\omega - 10)^2]$ is shown in Figure 2.11(b). When this filter is applied to the input signal $x(t) = \sin t + \sin 10t$ from Figure 2.12(a), the output is approximately $y(t) = \sin 10t$ as shown in Figure 2.12(c), i.e., the low-frequency component is eliminated. The action of the high-pass filter is shown in Figure 2.13(b); the high-frequency tone passes unaltered, while the delta spike at $\omega = 1$ is all but eliminated.

It is often a requirement to recover a high-frequency sinusoid or modulated carrier in the presence of high- and low-frequency noises as well as in-band noise. For this, a *bandpass* filter may be employed. This is a filter that passes signals close to a given centre frequency only. For example, an FM radio receiver may be tuned to receive a signal at 90 Mhz, the signal being embedded in white noise. A bandpass filter that



Figure 2.13: Action of (a) low-pass and (b) high-pass filters from Figure 2.11 on signal from Figure 2.12(a) shown in the frequency domain.

passes frequencies in the range 89.5-90.5 Mhz could be employed to allow the signal through, while largely eliminating the white noise. A more common approach, however, is to mix (multiply) the 90 Mhz signal with a 90 Mhz reference sinusoid in the receiver, thus producing replicas of the input signal shifted to *baseband* (0 Mhz) and 180 Mhz. A low-pass filter can then be used to eliminate the double-frequency term as well as most of the noises.

2.3.2 Digital linear filters

Filters are very often implemented in the digital domain, on a computer or similar digital system, in which case the filter is an algorithm that runs on discrete samples of the input signal $\{x_i\}$ to produce samples of an output signal $\{y_i\}$. In some cases the signals and filter may be inherently digital, and a digital implementation is the only option. In many other cases, the digital method is the preferred option; for example, in a telecoms receiver an analogue signal may be sampled using an *analogue-to-digital converter (ADC)*, with resulting signal samples then fed into a digital filter. Further processing may then performed in order to extract the information content of the input signal. In other digital signal-processors, the output digital samples might be fed into a *digitalto-analgue converter (DAC)* to produce an output analogue signal.

Digital implementations are usually chosen to avoid the practical is-

sues encountered in the synthesis of analogue filters. A perfect integrator, for example, is impossible to realise exactly. In contrast, a digital integrator is a simple mathematical process, easily programmed on a digital computer. Digital algorithms are also far more flexible than analogue circuits; very often, changing the characteristics of a digital filter is as simple as changing a few lines of computer code!

The integrator of equation (2.85) can be written for digital signals as:

$$y_i = \sum_{n=-\infty}^i x_n. \tag{2.88}$$

Clearly, each y_i is function of all previous values of the input signal x_j . However, such an integrator would be implemented using the simple recursive relation

$$y_i = y_{i-1} + x_i. (2.89)$$

The advantage of the digital version is obvious: the filter can simply be a single variable, which is the running sum of all input signal samples.

The digital equivalent of the Dirac delta impulse, $\delta(x)$, is the Kronecker delta δ_{i0} . Taking the input signal to the integrator to be such an impulse (so that $x_0 = 1, x_i = 0 \ \forall i \neq 0$), then we have the output signal $y_j = 1 \ \forall j \geq 0$. Thus, the impulse response never decays to zero. Such filters are called *infinite impulse response (IIR)* filters, to distinguish them from *finite impulse response (FIR)* filters. These classes also exist for analogue filters but the differences are more pronounced in the digital domain. For example, the IIR property of filter (2.89) can easily be deduced from the *feedback* in the equation, i.e., each value of the output signal is dependent of the previous value of the output signal. FIR filters cannot be written in such a form.

A general digital filter can be written as a difference equation

$$\sum_{n=0}^{N} a_n y_{i-n} = \sum_{n=0}^{M} b_n x_{i-n}.$$
(2.90)

In this, N is the order of the feedback filter, and the a_n the feedback

coefficients. The value M is the feed-forward filter order and the b_n the feed-forward coefficients. In general, we have $a_0 = 1$ so that the equation can be written in the form $y_i = \dots$ as in equation (2.89).

Many of the concepts for analogue filters outlined in Section 2.3.1 carry through for digital filters also. For example, frequency domain and transfer function analysis can be performed similarly using *Z*-transforms, which are the digital, discrete-time equivalent of Laplace transforms. The *Z*-transform of a time-discrete signal $\{x_i\}$ is defined as the power series

$$X(z) = \sum_{n=0}^{\infty} x_n \, z^{-n}.$$
 (2.91)

Note that, like the Laplace transform, the Z-transform is a function of the complex, continuous variable z even though the underlying time variable is discrete. The Z-transform of the difference equation (2.90) is Y(z) = H(z)X(z), where the transfer function H(z) is, assuming $a_0 = 1$,

$$H(z) = \frac{\sum_{n=0}^{M} b_n \, z^{-n}}{1 + \sum_{n=1}^{N} a_n \, z^{-n}}.$$
(2.92)

An FIR filter has $a_n = 0$ for all $n \ge 1$ so that $H(z) = \sum_{n=0}^{M} b_n z^{-n}$. An IIR filter has at least one $a_n \ne 0$ for $n \ge 1$. A common type of FIR filter is the *moving average filter*, which is often used to smooth input signals and works by simply taking the average of the previous M input samples so that $b_i = \frac{1}{M}$ for $i = 0, \ldots, M - 1$ and $b_n = 0$ for $n \ge M$. In this case we have

$$y_i = \frac{1}{M} \sum_{n=0}^{M-1} x_{i-n},$$

and

$$H(z) = \frac{1}{M} \sum_{n=0}^{M-1} z^{-n}.$$

The digital integrator of equation (2.89) has $a_0 = 1$, $a_1 = -1$, $b_0 = 1$,

and all other a_n and b_n set to zero. Thus, its Z-tranform is

$$H(z) = \frac{1}{1 - z^{-1}}.$$
(2.93)

Another component commonly used in digital circuits is a delay element, which simply gives as an output a delayed version of the input, i.e.,

$$y_i = x_{i-D}.$$

This simply produces a delay of D time intervals. The delay element has all a_n and b_n coefficients set to zero apart from $a_0 = 1$ and $b_D = 1$. Thus the Z-transform is

$$H(z) = z^{-D}.$$
 (2.94)

Real-time digital systems are always causal, i.e., cannot produce an output instantaneously in response to a given input. For example, in practice the digital integrator described earlier might actually incur a delay of a single time interval, in which case equation (2.89) is modified as

$$y_i = y_{i-1} + x_{i-1}, (2.95)$$

which, in contrast with (2.93), has Z-transform

$$H(z) = \frac{z^{-1}}{1 - z^{-1}}.$$
(2.96)

There is further reading on the general topic of digital filters in the "Digital PLLs" chapter of [2], and for a more in-depth text on the topic, for example, see [24].

Chapter 3

PLL theory

This chapter provides the background material on phase-locked loops (PLLs) needed for an understanding of our results, which are presented from Chapter 4 onwards. The PLL introduction is split into two sections: analogue and digital loops are described separately as the mathematics used in their analysis is quite different in each case.

Analogue PLL theory is a mature area of study and literature on the topic is extensive, from the early pioneering works by Viterbi in the 1950s and 60s [1] through to the present day [25]. The underlying mathematical theory used in the analysis of analogue PLLs is also quite mature, from linear and non-linear ODEs to continuous stochastic processes, Fokker-Planck equations etc. ([13], [18], [19], [20], [21], [22]).

Digital PLLs, in contrast, are a much newer area of study, stemming from the rise of computing and other digital systems in the 1980s. Few complete works on digital PLLs exist at this time; a good reference is an earlier work on analogue PLLs by Gardner [2], which has recently been updated to include several chapters on digital theory. Most of the current literature is in the form of journal papers such as those by Teplinsky et al. [4, 5].

3.1 Analogue PLL summary

The section summarises the basics of analogue phase-locked loops, required for an understanding of the remainder of this thesis. This summary is based largely on the approach taken in the first four chapters of [1], chapters 1, 2, and 6 of [2], as well as some additional material from the later work by Gardner [3].

3.1.1 Basics



Figure 3.1: Simplified block diagram of a phase-locked loop.

A phase-locked loop consists of three main components arranged in a closed loop as shown in Figure 3.1. The output of the loop is a timecontinuous sinusoidal signal which is generated by a *voltage-controlled oscillator (VCO)*. Oscillators are widely used in analogue circuits for the generation of radio-frequency (RF) signals. A standard oscillator generates a sinusoid at a fixed frequency. A VCO, however, accepts an input control voltage that is used to adjust its output frequency. The control voltage is shown as e(t) in Figure 3.1. Typically, a VCO will have a *quiescent frequency*, ω_0 , that will be output in the case where the control voltage is zero, and the frequency is adjusted linearly about ω_0 in response to positive or negative control voltages.

In the simplest case, the aim of the loop arrangement is to match the output signal of the VCO to a reference input signal as closely as possible. The input signal is shown in Figure 3.1 as $\sqrt{2}A \sin \theta_i(t)$ entering the loop at the top left, where A represents the amplitude of the signal. n(t) as shown in the diagram represents any noise that may be added to the input signal; for now we assume there is no noise present so that $n(t) \equiv 0$. The output signal of the VCO is denoted $\sqrt{2}K_1 \sin \theta_o(t)$. The loop aims to keep $\phi(t) = \theta_i(t) - \theta_o(t)$ as close as possible to zero under all conditions, i.e., the output signal is continually adjusted such that its frequency and phase matches the input signal as closely as possible. In general, the input signal may be embedded in noise, and recovery of this input signal is the main function of the loop. This is considered later, from Chapter 4 onwards. In other cases, a more advanced loop containing a frequency multiplier or divider is used to, for example, produce an output signal that is phase-locked to the input but at a multiple of the frequency; this is not considered here.

In order to drive the output signal towards the input, a comparator is needed to produce an error signal. This is implemented in the *phase detector (PD)* in Figure 3.1. Ideally, this would output simply $x(t) = \theta_i(t) - \theta_o(t) = \phi(t)$, and this is often realised in practice in digital circuits. However, for analogue loops the phase detector is often simply a multiplier operating on the input and output sinusoids, producing a combination of sinusoids involving sums and differences of the input and output phases. It will be seen later how this is usually sufficient for our purposes.



Figure 3.2: Second-order PLL block diagram.

The linear filter block is simply an LTI filter as described in Section 2.3. It is generally the filter type and tuning of its parameters that determine the overall dynamics of the loop. The filter takes as an input the difference signal produced by the PD and gives a filtered version of the difference signal as the output, an error signal e(t) that is used to adjust the VCO. In the simplest case of the *first-order loop*, the filter may simply be passthrough or proportional gain path, e.g., e(t) = Kx(t). This would result in an instantaneous adjustment of the VCO in response to a change in the input signal. This may be appropriate in some, but not all cases. Sometimes a different response is required, e.g., a "smoother" response where the output signal is adjusted more slowly in response to any input changes, or we might also require a non-zero response where the error signal is zero, to hold to VCO at a particular frequency. There are many ways this can be achieved in the loop filter, including the use of an integrator so that the filter output contains both a proportional and integral path. In this case the PLL is described as a second-order loop. Such a loop is shown in Figure 3.2; in this case the linear filter consists of a summation of a straight-through path and an integrator. A full exploration of the basic first- and second-order loops will be undertaken in Sections 3.1.2 and 3.1.3.

Note that, throughout this thesis, the *order* of a PLL, both analogue and digital, refers to the number of integrators in the loop. The VCO is essentially an integrator per equation (3.3) below, since it produces a signal with phase that is an integral of the input error signal. Therefore, a PLL will always be of order at least 1. The number of integrators is sometimes referred to in the literature as the *type* of the PLL, e.g., in [2].

In Figure 3.1 the input signal is

$$\sqrt{2}A\sin\theta_i(t),\tag{3.1}$$

while the output is

$$\sqrt{2}K_1 \sin \theta_o(t). \tag{3.2}$$

As described above, the VCO output frequency is given by

$$\frac{d\theta_o(t)}{dt} = \omega_0 + K_2 e(t), \qquad (3.3)$$

where ω_0 is the quiescent frequency and K_2 is the VCO gain. For a multiplier-type PD, the output is

$$x(t) = AK_1 \left\{ \sin \left[\theta_i(t) - \theta_o(t) \right] + \sin \left[\theta_i(t) + \theta_o(t) \right] \right\}$$
(3.4)

The second term in equation (3.4) involves the sum of the input and output phases, and when the loop is locked or close to locked, this is a term that is at double the input frequency. For PLLs of order higher than 1, the loop filter is generally a *low-pass filter*, which eliminates high-frequency components in order to provide the necessary smoothing action. As such, it would be arranged that it effectively eliminated the double-frequency term in (3.4). Just as often, the PD itself would have a built-in filter that would eliminate this term so that only the zerofrequency term may be considered for all loop types, including those of first order. The error signal produced by the filter is, as in equation (2.78)

$$e(t) = e_0(t) + \int_0^t x(t-u)f(u)du, \qquad (3.5)$$

where f(t) is impulse response of linear filter, and e_0 is the initial output of the filter, often assumed to be zero. Combining (3.3)-(3.5), we get

$$\frac{d\theta_o(t)}{dt} = w_0 + K_2 \int_0^t f(t-u) A K_1 \sin\left[\theta_i(u) - \theta_o(u)\right] du$$
(3.6)

Now if we take the phase error

$$\phi(t) = \theta_i(t) - \theta_o(t), \qquad (3.7)$$

the loop gain

$$K = K_1 K_2, \tag{3.8}$$

and the normalized input and output phases

$$\theta_1(t) = \theta_i(t) - \omega_0 t$$

 $\theta_2(t) = \theta_o(t) - \omega_0 t$

we get, finally

$$\frac{d\phi(t)}{dt} = \frac{d\theta_1(t)}{dt} - AK \int_0^t f(t-u)\sin\phi(u)du.$$
(3.9)

This last is a non-linear integro-differential equation for the phase error, ϕ , in terms of the loop parameters A, K, and f(t), as well as the input signal phase, $\theta_1(t)$. The loop is said to be *in lock* when the phase error is at zero.

3.1.2 Linear approximation



Figure 3.3: Laplace domain representation of linearized PLL.

If the loop as described in Section 3.1.1 remains close at all times to the locked state, so that $\phi(t) \ll 1$ rad, then we can use the approximation $\sin \phi(t) \simeq \phi(t)$. In this case (3.9) becomes the linear equation

$$\frac{d\phi(t)}{dt} = \frac{d\theta_1(t)}{dt} - AK \int_0^t f(t-u)\phi(u)du.$$
(3.10)

This can be analyzed most simply by taking Laplace transforms of each

side, assuming from the outset that the transforms exist, to obtain

$$s\Phi(s) = s\Theta_1(s) - AKF(s)\Phi(s), \qquad (3.11)$$

where Φ , Θ_1 , and F, are the transforms of ϕ , θ_1 , and f respectively. Alternatively, this can be written in the following form:

$$\Phi(s) = \frac{1}{1 + AKF(s)/s} \Theta_1(s).$$
(3.12)

Noting now that $\Phi(s) = \Theta_1(s) - \Theta_2(s)$, where $\Theta_2(s)$ is the transform of $\theta_2(t)$, we can also write

$$\Theta_2(s) = \frac{AKF(s)/s}{1 + AKF(s)/s} \Theta_1(s).$$
(3.13)

These last two equations give directly the phase error (3.12) and phase of the loop's output signal (3.13), where the input phase and loop parameters are known. In particular, because the equation of operation (3.13) can be written in the form given in (2.80), it is clear that the PLL is itself a LTI filter with transfer function

$$H(s) = \frac{AKF(s)/s}{1 + AKF(s)/s},$$
(3.14)

so that the equations for the output phase and phase error can be written more concisely as

$$\Theta_2(s) = H(s)\Theta_1(s), \tag{3.15}$$

$$\Phi(s) = [1 - H(s)] \Theta_1(s). \tag{3.16}$$

For a PLL, H(s) is known as the closed-loop transfer function.

3.1.2.1 First-order loop

If we take firstly the simplest case of a loop without a filter — a standard first-order loop — then we have $F(s) \equiv 1$ and $f(t) = \delta(t)$. This gives a closed-loop transfer function of

$$H(s) = \frac{AK}{s + AK}.$$
(3.17)

If we further assume an input signal that is a sinusoid of constant frequency ω and initial phase θ_s , then we have

$$\theta_1(t) = (\omega - \omega_0)t + \theta_s.$$

This is, equivalently, in the Laplace domain

$$\Theta_1(s) = \frac{\omega - \omega_0}{s^2} + \frac{\theta_s}{s},$$

where it is assumed that the input signal starts at t = 0. Equation (3.12) now yields, for this example,

$$\Phi(s) = \frac{\omega - \omega_0}{s(s + AK)} + \frac{\theta_s}{s + AK}.$$

It is possible to take inverse Laplace transforms of this to return to the time domain, where we find the phase error is given by

$$\phi(t) = \frac{\omega - \omega_0}{AK} \left(1 - e^{-AKt} \right) + \theta_s e^{-AKt}.$$
(3.18)

Clearly, this expression for $\phi(t)$ contains two terms that decay as $t \to \infty$, which means that for this first-order loop we have a steady-state phase error $\frac{\omega-\omega_0}{AK}$. For the linear approximation to be valid requires $\phi(t)$ to be small for all t, so in particular $\frac{\omega-\omega_0}{AK}$ must be small, as must ϕ_s . The form of the steady-state phase error means that the first-order loop will not lock exactly as $t \to \infty$ unless $\omega = \omega_0$, i.e., the frequency of the input frequency is identical to the quiescent frequency of the VCO. This is difficult to realise in practice. This non-zero static phase error is a property of first-order loops in general, as we shall see in later sections.

3.1.2.2 Second-order loop

As a further example, consider a PLL where a loop filter has been inserted, where the filter consists of a perfect integrator as well as the passthrough path. Such a PLL is shown in Figure 3.2. This gives

$$F(s) = 1 + \frac{a}{s}.$$

In operator notation, this is

$$\mathscr{F}\{i(t)\} = i(t) + a \int_0^t i(t')dt'.$$

Now, inserting the expression for F(s) for this example into equation (3.14), the closed-loop transfer function for the second-order loop is obtained as

$$H(s) = \frac{AK(s+a)}{s^2 + AKs + aAK}.$$
 (3.19)

Note that for a = 0 this reduces to the transfer function for the firstorder loop (3.17) as expected. Now, for a constant-frequency sinusoid as before, the phase error in the Laplace domain can be obtained from (3.16) as

$$\Phi(s) = \frac{s^2}{s^2 + AKs + aAK} \left(\frac{\omega - \omega_0}{s^2} + \frac{\theta_s}{s}\right) = \frac{(\omega - \omega_0) + \theta_s s}{s^2 + AKs + aAK}.$$
 (3.20)

Again, this can be inverse-transformed to obtain the phase error, $\phi(t)$, in the time-domain. However, typically we are interested only in the steady-state phase error, in which case we can make use of the finalvalue theorem for Laplace transforms [26], which is

$$\lim_{t \to \infty} \phi(t) = \lim_{s \to 0} s \Phi(s).$$

For the second-order loop, with $\Phi(s)$ as in (3.20), and assuming $a, A, K \neq 0$, this is just $\lim_{t\to\infty} \phi(t) = 0$. So, unlike the first-order loop, this loop can lock with zero static phase error to a signal that differs in frequency to the VCO quiescent frequency. This is made possible by the inclusion

of the integrator in the loop; the filter can produce a non-zero output for phase error of zero, which maintains the VCO frequency at ω .

3.1.2.3 Varying input frequency

Next, we return to the example of the first-order loop, but now take the input signal to have a frequency that is varying linearly with time with rate R. That is, the input signal is of form

$$\sqrt{2}A\sin\left[\int_{0}^{t}(\omega+Rt^{'})dt^{'}+ heta_{s}
ight],$$

which yields

$$\theta_1(t) = \frac{1}{2}Rt^2 + (\omega - \omega_0)t + \theta_s.$$
(3.21)

This characteristic would be typical of signals received from a sinusoidal generator when the transmitter and receiver are accelerating relative to one another. Inserting once again into equation (3.16) with $F(s) \equiv 1$ for the first-order loop, gives the phase error in the Laplace domain

$$\Phi(s) = \frac{s}{s + AK} \left(\frac{R}{s^3} + \frac{\omega - \omega_0}{s^2} + \frac{\theta_s}{s} \right).$$

Application of the final-value theorem for Laplace transforms shows that $\phi(t)$ grows unbounded at $t \to \infty$, so the first-order loop is never able to lock to a signal with linearly-varying frequency. If we use instead a second-order loop for this input signal, the transformed phase error becomes

$$\Phi(s) = \frac{s^2}{s^2 + AKs + aAK} \left(\frac{R}{s^3} + \frac{\omega - \omega_0}{s^2} + \frac{\theta_s}{s}\right),$$

which gives

$$\lim_{t \to \infty} \phi(t) = \frac{R}{aAK}.$$
(3.22)

This says that the second-order loop is able to lock with a static phase error that is proportional to the radial acceleration, R. In general, the order of a PLL is equal to the number of integrators in the loop, including the VCO. It can be shown that a third-order loop could lock to a signal with linearly-varying frequency with zero static phase error. For an input signal with non-zero *n*-th phase derivative, a *n*-th order PLL can generally lock with (possibly) non-zero static phase error, and any higher-order loops will lock with zero phase error. Higher-order loops are more difficult to realize in practice, particularly in analogue circuits and are more prone to instabilities than those of lower-order.

3.1.3 Non-linear PLLs in the absence of noise

For this section we return to the full non-linear equation of operation for the PLL, (3.9), and take specific examples of loop types and input signals. Use of the full equation removes any restriction on the size of the phase error, $\phi(t)$, thus allows for analysis of PLLs over the full range of behaviours. This is important, for example, for examining the dynamics during acquisition, various transients, and loops that cannot lock with zero or small phase error. The nonlinearity in equation (3.9) arises from the multiplier type PD, which, when filtered appropriately, gives an error signal proportional to $\sin \phi(t)$. Note that there exist PDs with characteristics other than the sinusoidal one assumed here. However, the multiplier PD is the one most commonly used [2], so for the remainder of this thesis it is assumed the PD is of this type.

3.1.3.1 First-order loop

If we first consider the non-linear version of the simplest example considered in Section 3.1.2, the first-order loop with constant-frequency input signal, we have, for equation (3.9), $f(t) = \delta(t)$ so that $F(s) \equiv 1$ and $\theta_1(t) = (\omega - \omega_0)t + \theta_s$. Thus, (3.9) becomes

$$\dot{\phi}(t) = (\omega - \omega_0) - AK\sin\phi(t), \qquad (3.23)$$

where $\dot{\phi}(t)$ represents $d\phi(t)/dt$.

The behaviour of this system can be understood most easily by look-


Figure 3.4: Trajectory of first-order loop on phase plane.

ing at its trajectory on a phase plane plot. Phase plane analysis is a standard method for visualizing and understanding non-linear systems. A very readable introductory text on the subject is [27], where Chapter 6 deals with the phase plane. If we plot $\dot{\phi}$ against ϕ , it can be seen that the trajectory is a sinusoid, shifted in the horizontal and vertical directions. While the trajectory is above the horizontal axis $\dot{\phi}$ is positive, so ϕ will increase along the curve, and vice-versa. At the points of intercept with the axis, the derivative will be zero and the system will have reached an equilibrium point. It can be seen from the directional indicators along the curves that alternate points are stable and unstable, i.e., any trajectory in the vicinity of the first intercept to the right of the origin will converge to that point, whereas any trajectory near the next equilibrium point will diverge from that point and settle instead near the first or third point. The system will remain at the unstable equilibrium points only if the derivative is exactly zero at the outset.

It should be clear from equation (3.23) and Figure 3.4 that the points of stable equilibrium for the first-order loop are

$$\phi_n = 2n\pi + \sin^{-1}\frac{\omega - \omega_0}{AK},$$

and the points of unstable equilibrium

$$\phi_n = 2(n-1)\pi + \sin^{-1}\frac{\omega - \omega_0}{AK},$$
(3.24)

where $n \in \mathbb{Z}$. The stable equilibrium point for n = 0 is

$$\phi_0 = \sin^{-1} \frac{\omega - \omega_0}{AK}.$$
(3.25)

If this ϕ_0 is small compared to 1 radian, then we have

$$\phi_0 \approx \frac{\omega - \omega_0}{AK},$$

which was exactly the static phase error found in the linear model. This approximation in turn imposes a requirement on the input signal's frequency, as would be expected, i.e., that its frequency offset from the VCO frequency is small as compared to the loop parameters:

$$|\omega - \omega_0| \ll AK.$$

Note that, even without any approximations, for any stable or unstable equilibrium points to exist for this first-order loop, we must have

$$|\omega - \omega_0| \le AK.$$

Were this not the case, the sinusoid in Figure 3.4 would lie entirely either above or below the horizontal axis, and any trajectory would continue indefinitely in one direction along the curve.

3.1.3.2 Second-order loop

If we next take again the example of the second-order loop with constantfrequency input, so that F(s) = 1 + a/s and $\theta_1(t) = (\omega - \omega_0)t + \theta_s$ so that (3.9) in this case is

$$\frac{d^2\phi(t)}{dt^2} + AK\left(\frac{d}{dt} + a\right)\sin\phi(t) = \frac{d^2\theta_1(t)}{dt^2} = 0, \qquad (3.26)$$

which becomes

$$\frac{d^2\phi}{dt^2} + AK\cos\phi\frac{d\phi}{dt} + aAK\sin\phi = 0.$$
(3.27)

Using the normalization $\tau = AKt$, this can be simplified further as follows

$$\ddot{\phi} + \dot{\phi}\cos\phi + a'\sin\phi = 0, \qquad (3.28)$$

where $\dot{\phi} = d\phi/d\tau$ and a' = a/AK. We can eliminate the independent variable, τ , from this and consider $\dot{\phi}$ and ϕ to be independent variables to be plotted on the phase plane according to

$$\frac{d\dot{\phi}}{d\phi} = -\cos\phi - a'\frac{\sin\phi}{\dot{\phi}}.$$
(3.29)

This equation is most easily visualized on a phase-plane plot. Because it is periodic in ϕ with period 2π it is sufficient to plot it only in the range $-\pi \leq \phi \leq \pi$. Some other properties of the plot are also notable from (3.29):

- The slope of the trajectories across the $\phi = 0$ axis is -1 always.
- For large $\dot{\phi}$ the trajectories are almost sinusoidal.
- At the origin and at at all points $(\phi = n\pi, \dot{\phi} = 0)$ the last term in (3.29) is indeterminate, and at this singularity the system has either a stable equilibrium point or unstable saddle point.

An examination of the plot in Figure 3.5 bears out the observations above as well as several other interesting properties:

- Away from the ϕ -axis, the trajectories are almost sinusoidal. In the upper half-plane, $\dot{\phi}$ is positive, so ϕ will increase. That is, in the upper half-plane, trajectories move from left to right, and in the lower half-plane the opposite is the case.
- Above the ϕ -axis, any trajectory that begins below the separatrix (plotted in bold) will be pulled in to the ϕ -axis to the right of the



Figure 3.5: Trajectory of second-order loop on $\phi - \dot{\phi}$ phase plane for a' = 1.

origin and from there attracted back around to the stable equilibrium point at the origin. Clearly, the stable equilibrium point corresponds to the locked state of the PLL, since $\phi = 0$ and the point is attractive.

- In the lower half-plane a similar characteristic is evident: any trajectory that begins (on the right) between the ϕ -axis and the separatrix will be pulled into lock.
- Trajectories that begin in the upper-half plane but above the separatrix will not be pulled into the origin on this period. However, note that any such trajectory will end closer to the ϕ -axis at the end of the period than it was at the start, i.e., $\dot{\phi}(\pi) < \dot{\phi}(-\pi)$. Thus, any such trajectory will eventually start below the separatrix on some period, and from there will be pulled into lock. Such a phenomenon is called *cycle slipping* and is a well-known characteristic of the acquisition phase of a PLL: while a PLL starts trying to lock on an input signal with an arbitrary frequency, the phase error will go through several cycles of 2π until the integrator sum builds up sufficiently to match the VCO frequency to that of the

input frequency. Once frequency lock as been achieved, phase lock will be achieved without any further cycle slips.

- Symmetry again applies in the lower half-plane. The PLL will slip cycles for large negative frequency errors before eventually achieving frequency lock, and then phase lock on the same cycle
- Clearly the second-order loop can lock to a signal with any input frequency with zero static phase error. This is the same result we found earlier using the linear approximation for small φ. In this case, however, we have allowed for large φ and φ, and the graphical approach shows that, although lock will always be achieved eventually, it may take arbitrarily long for the cycle slipping to stop, depending on how large φ is, which is usually determined by the input signal frequency. Usually, an acquisition aid is required to get reasonable *pull-in* performance for a second-order loop. Once such method is to sweep the VCO frequency linearly with time. A variant of this problem is considered in the next section.
- Finally note that any point near the unstable equilibria at $\phi = \pm (2n+1)\pi$ will be taken quickly away from that point and either back toward the origin, or off into the next cycle, depending on which side of the ϕ -axis the point was on.

The characteristics of the plot shown in Figure 3.5 vary depending on the loop parameters, in particular the value of a'. Further plots are shown in [1]. However the observations above apply to all second-order loops.

If we take as a final example again the case where the input signal has constant radial acceleration, which gives, as in (3.21)

$$\theta_1(t) = \frac{1}{2}Rt^2 + (\omega - \omega_0)t + \theta_s.$$

Note that this same form for $\theta_1(t)$ would result if the input frequency had constant frequency input ω , and the VCO frequency was swept with rate -R. Then the general non-linear PLL equation of operation (3.9) becomes

$$\frac{d^2\phi}{dt^2} + AK\cos\phi\frac{d\phi}{dt} + aAK\sin\phi = R.$$
(3.30)

This is identical to (3.27), except that now we allow the right-hand side to be non-zero. Proceeding exactly as before, this yields the analog of (3.28):

$$\ddot{\phi} + \dot{\phi}\cos\phi + a'\sin\phi = R', \qquad (3.31)$$

where $R' = R/(AK)^2$. For phase-plane analysis, this can be written as

$$\frac{d\dot{\phi}}{d\phi} = -\cos\phi + \frac{R' - a'\sin\phi}{\dot{\phi}},\tag{3.32}$$

which reduces to (3.29) for R = 0. The singular points now become

$$\dot{\phi} = 0, \ \phi = \sin^{-1}\left(\frac{R'}{a'}\right) \pm 2n\pi, \ n = 0, 1, 2, \cdots,$$

and

$$\dot{\phi} = 0, \ \phi = \pi - \sin^{-1}\left(\frac{R'}{a'}\right) \pm 2n\pi$$

If R' = 0, these reduce to the stable and unstable equilibrium points of Figure 3.5. For R' = a' (R = aAK), the two sets of singularities coincide, and for R' > a' there are no singularities.

In Section 3.1.2 the linear approximation for the second-order loop with linearly-varying input frequency yielded (3.22), which showed that the static phase error was R/aAK. This approximation requires ϕ to be small, and in particular in the range $[-\pi, \pi]$. Our stable equilibrium point in this range for the non-linear case is

$$\phi = \sin^{-1}\left(\frac{R'}{a'}\right) = \sin^{-1}\left(\frac{R}{aAK}\right).$$

For the linear approximation to be valid we must have small R/aAK, and so we have

$$\sin^{-1}\left(\frac{R}{aAK}\right) \approx \frac{R}{aAK},$$

i.e., our exact solution agrees with the linear approximation.



Figure 3.6: Trajectory of second-order loop on $\phi - \dot{\phi}$ phase plane for $a' = \frac{1}{2}$ and $\frac{R'}{a'} = \frac{1}{2}$.

The phase-plane plot for (3.32) is shown in Figure 3.6. Some observations can be made regarding this system:

- From (3.32), the graph is still periodic in ϕ , so it is sufficient for plot for $\phi \in [-\pi, \pi]$.
- The singularities are of exactly the same nature as those of Figure 3.5, but they are shifted on the φ-axis. The stable singularity on the left in Figure 3.5 is shifted to the right be an amount sin⁻¹ (R'/a'), while the unstable one on the right is shifted to the left by the same amount.
- Like the case where R' = 0, any trajectory between the upper and lower separatrix will be pulled into phase lock for this set of parameters (R'/a' = 0.5). This, however is not true for larger values of R' (1/2 < R' < 1), for which some plots are shown in [1].
- Any trajectory that begins above the upper separatrix ends the $[-\pi,\pi]$ period higher than where it began, i.e., with larger ϕ .

Therefore, these trajectories, taken over several periods, will diverge from the separatrix and phase lock will never be achieved for these initial conditions.

- For the R'/a' = 1/2 case shown, all trajectories below the lower separatrix move upwards on each cycle until they are eventually above the separatrix on some cycle, from where they will be pulled into phase lock. The rate at which the trajectories move upwards in Figure 3.6 is greater than it was in Figure 3.5.
- These last two observations confirm the intuition that, where the input frequency is linearly increasing with time, it is better to have that frequency initially lagging the VCO frequency (negative $\dot{\phi}$), so that it will increase towards the VCO frequency. A frequency with initial negative frequency offset will be acquired by the PLL faster for R' > 0 than for the case of R' = 0.
- For R'/a' → 1 the upper separatrix moves towards the φ-axis, thus reducing the number of trajectories that can achieve phase lock. There are also more trajectories above the lower separatrix that get pulled above the φ-axis, above the upper separatrix, from where frequency and phase-lock are lost. This corresponds to a loop that is close to phase lock being pulled out of lock because the frequency sweep rate is too great
- For $R'/a' \ge 1$ there are no singularities and no stable points. Any trajectory will have monotonically-increasing or decreasing ϕ depending on the sign of $\dot{\phi}$.
- For the acquisition of a signal with fixed frequency ω , sweeping the VCO frequency with rate -R yields an identical system that shown in Figure 3.6. Therefore it can be seen that, if we start the VCO at a fixed frequency and sweep it linearly towards that of the input signal, the acquisition time will be reduced. This approach is often used as an acquisition aid in real systems [1]. However, this comes

at the price of reduced stability; for large values of R $(R'/a' \approx 1)$ there is the risk of the VCO frequency being swept right past that of the input signal without ever achieving lock.

3.1.4 Additive noise in PLLs



Figure 3.7: Simplified block diagram of a phase-locked loop from Figure 3.1, showing equivalent addition of noise after phase detector.

We return now to the block diagram shown in Figure 3.1 and take the case where noise is added to the input signal, that is $n(t) \neq 0$. In order to derive useful analytical results it is necessary to make some assumptions about the properties of the additive noise. These assumptions may only be approximated by the noise processes in real systems. Under these assumptions, to be outlined below, the system with noise added to the input signal in Figure 3.1 is equivalent to the system shown in Figure 3.7, where the noise is instead added after the phase detector. This means the noise enters as an additional term in equation (3.9), which means we can often use the tools of stochastic calculus to obtain useful results.

We assume firstly that the noise can be represented as a zero-mean stationary Gaussian stochastic process. The simplest analytical results are typically derived by assuming the noise process to be white with spectral density $N_0/2$, i.e., having a flat spectral density of $N_0/2$ across the entire frequency range. One could assume this to be the case here also. However, the linear filter in a PLL will typically result in the PLL passing only a certain range of frequencies, centred at the VCO quiescent frequency, ω_0 . So, even if the original additive noise is white, one could equally assume it to have a flat PSD equal to $N_0/2$ across a small frequency range around ω_0 only, that is a *narrowband Gaussian process*. The reasoning here is similar to that used in the elimination of the doublefrequency term in equation (3.4). It can be shown (see Appendix A of [1], for example) that such a process, n(t), can be represented as

$$n(t) = \sqrt{2}n_1(t)\sin\omega_0 t + \sqrt{2}n_2(t)\cos\omega_0 t + \tilde{n},$$
 (3.33)

where \tilde{n} is the mean of the process n(t) and $n_1(t)$ and $n_2(t)$ are zero-mean, stationary narrowband Gaussian processes centred at the zero frequency (DC) with PSDs identical to that of n(t) but shifted to DC. Figure 3.8 shows the PSD of a sample narrowband noise process n(t) centred around frequency ω_0 as well as the PSD of the corresponding process $n_1(t)$ shifted to DC.



Figure 3.8: Power spectral densities of sample narrowband noise process, n(t) centred at frequency ω_0 and process $n_1(t)$ with spectrum shifted to DC.

Taking all of this, the expression (3.1) for the input signal becomes,

in the noisy case,

1

$$\sqrt{2}A\sin\theta_{i}(t) + n(t) = \sqrt{2}\{A\sin[\omega_{0}t + \theta_{1}(t)] + n_{1}(t)\sin\omega_{0}t + n_{2}(t)\cos\omega_{0}t\}.$$
(3.34)

In this, we have used the normalized input phase $\theta_1(t)$ and the fact that the noise process is zero-mean to set \tilde{n} in (3.33) to zero. The output of the VCO is, as in (3.2)

$$\sqrt{2}K_1 \sin[\omega_0 t + \theta_2(t)].$$
 (3.35)

We can combine (3.34) and (3.35) to find that the multiplier output, after eliminating double-frequency terms as before, is

$$x(t) = K_1 \{ A \sin \phi(t) - n_1(t) \sin \theta_2(t) + n_2(t) \cos \theta_2(t) \}, \qquad (3.36)$$

where $\phi(t) = \theta_1(t) - \theta_2(t)$ is the phase error. Defining a new noise term $\breve{n}(t) = n_1(t) \sin \theta_2(t) + n_2(t) \cos \theta_2(t)$, the error signal input to the VCO is

$$e(t) = K_1 \int_0^t \left[A \sin \phi(u) + \breve{n}(u) \right] f(t-u) du, \qquad (3.37)$$

so we get a integro-differential equation for the phase error as before, but now with an additional noise term:

$$\frac{d\phi(t)}{dt} = \frac{d\theta_1(t)}{dt} - K \int_0^t \left[A\sin\phi(u) + \breve{n}(u)\right] f(t-u)du, \qquad (3.38)$$

where $K = K_1 K_2$.

It is shown in [1] that, under the assumptions above on $n_1(t)$ and $n_2(t)$, namely that they are Gaussian processes with flat PSDs with density $N_0/2$ within the bandwidth of the loop, $\breve{n}(t)$ is similarly a Gaussian process with flat spectrum $N_0/2$ within the loop bandwidth. This implies that for practical purposes we may treat it as AWGN in our loop equation (3.38). This last is the equivalent of the integro-differential equation (3.9) found earlier in the noise-free case, but in this case has an extra AWGN term. The PSD and variance of this loop phase error term, $\breve{n}(t)$, are identical to those of the original additive noise on input, n(t).

From (3.9) we were able to obtain ODEs for the phase error, ϕ , for specific types of PLL, e.g., equation (3.23) for the first-order loop. Now, considering the noisy signal scenario, these ODEs become Langevin equations, which can be treated using the tools of stochastic differential equations. Some examples are detailed later in Section 3.1.6. Where we use the linear approximation $\sin \phi \approx \phi$, further simplifications are possible: we can apply the principle of superposition to treat the signal and noise separately and work in the Laplace domain as we did in Section 3.1.2. The next section explores this in more depth.

3.1.5 Additive noise in the linear model



Figure 3.9: Conceptual representation of signal+noise passing through linearized PLL.

In the linear model, we use the approximation $\sin \phi \approx \phi$ and so the equation of operation for the PLL (3.38) becomes

$$\frac{d\phi(t)}{dt} = \frac{d\theta_1(t)}{dt} - K \int_0^t \left[A\phi(u) + \breve{n}(u)\right] f(t-u)du.$$
(3.39)

Rather than work directly in the time domain with (3.39), results can more easily be obtained by working in the Laplace domain, as we did in the noise-free case in Section 3.1.2. There we found that the linearized PLL is an LTI filter with transfer function

$$H(s) = \frac{AKF(s)/s}{1 + AKF(s)/s},$$

so that the output of the loop is related to the input as

$$\Theta_2(s) = H(s)\Theta_1(s).$$

This relation applies to all input signals Θ_1 . Because the system is linear it also applies to all sums of input signals. In particular, we can consider the action of the PLL on the signal and noise separately. For example, Figure 3.9 shows the action of the linear PLL on an input that is the sum of a signal of interest plus noise. Here, the response of the filter to the signal, $\Theta_1(s)$ is $\Theta_{2S}(s)$, and the response to the noise, N(s)is $\Theta_{2N}(s)$, the total output being $\Theta_2(s) = \Theta_{2S}(s) + \Theta_{2N}(s)$. Therefore, to investigate the response of the loop to an input signal plus noise, we can equivalently look at the response to noise alone and so can set $\theta_1(t) = 0$.

We saw in Section 3.1.4 that, under certain conditions on the noise, a noise term n(t) added to the input signal was equivalent to adding the noise term $\breve{n}(t)$ to the phase error $\phi(t)$ after the phase detector, which was in that case non-linear. This equally applies to the linear phase detector. In that case, assuming $\theta_1(t) = 0$ so that the PLL's input is noise only, the loop shown in Figure 3.7 becomes the one in Figure 3.10. Now If we take the additive noise, $\breve{n}(t)$, to be white, zero-mean, with one-sided spectral density N_0 so that $S_n(\omega) = N_0/2$, then the spectral density of the input signal is $N_0/2A^2$ (we are scaling the input noise by 1/A as we include a scaling factor of A inside the loop itself). The spectral density of the output error signal, $\phi(t)$, caused by the noise is found from equation (2.84) as

$$S_{\phi}(\omega) = \frac{N_0}{2A^2} |H(i\omega)|^2.$$
 (3.40)

The variance is then calculated as

$$\sigma_{\phi}^{2} = \frac{N_{0}}{2A^{2}} \int_{-\infty}^{\infty} |H(i\omega)|^{2} \frac{d\omega}{2\pi} = \frac{N_{0}}{A^{2}} \int_{0}^{\infty} |H(i\omega)|^{2} \frac{d\omega}{2\pi}.$$
 (3.41)

Defining the noise bandwidth of the loop as

$$B_L = \int_0^\infty |H(i\omega)|^2 \frac{d\omega}{2\pi},\tag{3.42}$$

we get, finally, a simple relation between the variance of the output error signal and the input noise level:

$$\sigma_{\phi}^2 = \frac{N_0 B_L}{A^2}.$$
 (3.43)

Note that the noise bandwidth is a property of the loop only and can be applied independently to any input signal.



Figure 3.10: PLL from Figure 3.7, linearized, and equivalent noise process shown as input.

Using the fact that for real, approximate noise processes, the noise variance, σ_n^2 , is proportional to the spectral density within the bandwidth of interest, N_0 . Thus, we have, from (3.43)

$$\sigma_{\phi}^2 = \frac{\kappa \sigma_n^2 B_L}{A^2},\tag{3.44}$$

where $\kappa \in \mathbb{R}$ is a constant. Thus, the output phase error variance plotted against the input noise variance would be a straight line of slope $\kappa B_L/A^2$.

3.1.6 Analysis of additive noise in non-linear model

We now return to the full non-linear equation of operation (3.38) for the PLL

$$\frac{d\phi(t)}{dt} = \frac{d\theta_1(t)}{dt} - K \int_0^t \left[A\sin\phi(u) + \breve{n}(u)\right] f(t-u)du.$$

Following the approach taken in [1], the most useful analytical results are obtained with the the simplest case: the first-order loop with sinusoidal input of frequency ω .

3.1.6.1 First-order loop

In this case the loop filter disappears so that $F(s) \equiv 1$ and $f(u) = \delta(u)$, in which case (3.38) becomes

$$\frac{d\phi(t)}{dt} = (\omega - \omega_0) - AK\sin\phi(t) - K\breve{n}(t).$$
(3.45)

This is a first-order system driven by white Gaussian input noise, \check{n} , with spectral density N_0 , and is modelled in [1] as a Markov process. In general, a process that is described by an *n*th order ODE with AWGN driving function can be modelled as an *n*-dimensional system of Markov random processes. Under the Markov assumption on ϕ in (3.45), its probability density satisfies the Fokker-Planck equation (2.73) and its exact form can be obtained by calculating the moments according to (2.72). In particular, these are

$$A_{1}(\phi) = \lim_{\Delta t \to 0} \frac{\mathscr{E}\left[\Delta \phi \mid \phi\right]}{\Delta t} = (\omega - \omega_{0}) - AK \sin \phi(t),$$

$$A_{2}(\phi) = \lim_{\Delta t \to 0} \frac{\mathscr{E}\left[(\Delta \phi)^{2} \mid \phi\right]}{\Delta t} = \lim_{\Delta t \to 0} \frac{K^{2}}{2} \int_{t}^{t+\Delta t} \int_{t}^{t+\Delta t} \mathscr{E}\left[\breve{n}(u)\breve{n}(v)\right] du dv$$

$$= \lim_{\Delta t \to 0} \frac{K^{2}N_{0}}{2\Delta t} \int_{t}^{t+\Delta t} \int_{t}^{t+\Delta t} \delta(u-v) du dv = \frac{K^{2}N_{0}}{2}.$$
(3.46)

In these, the quantity $\Delta \phi$ has been calculated directly from (3.45) as

$$\Delta \phi = \int_{t}^{t+\Delta t} \left. \frac{d\phi(u)}{du} \right|_{u=v} \, dv,$$

and the fact that the noise is white has yielded the simplified noise term in A_2 . It is also possible to show directly in this case [1] that all higher moments vanish, i.e., $A_n(\phi) = 0$ for $n \ge 3$. This is also true in the general case of a first-order ODE with white Gaussian driving function [18]. The moments yield the Fokker-Planck equation for the first-order loop:

$$\frac{\partial p(\phi,t)}{\partial t} = -\frac{\partial}{\partial \phi} \left[(\omega - \omega_0 - AK\sin\phi)p(\phi,t) \right] + \frac{K^2 N_0}{4} \frac{\partial^2 p(\phi,t)}{\partial \phi^2}.$$
 (3.47)

This is an exact PDE for the time-dependent probability distribution of the phase error, ϕ . Knowledge of this function would allow us to calculate all statistics relating the the loop's output. Note that the Fokker-Plank equation (3.47) could also have been derived directly from the Langevin equation (3.45) using the calculus of Itô or Stratonovich.



Figure 3.11: Examples of phase error densities for a first-order non-linear PLL with additive noise and $\omega = \omega_0$ at sample value of t for (a) full phase error density $p(\phi, t)$, and (b) density of phase error wrapped to a single cycle, $p_{\pi}(\phi, t)$.

For the noisy first-order PLL we would expect the phase error probability density to be centred around the expected steady-state phase error, which, in the first cycle, is $\phi_0 = \sin^{-1} \frac{\omega - \omega_0}{AK}$ from (3.25). Because of the additive noise, there will also be a non-zero probability that the trajectory will get pushed out of the first cycle, onto a higher or lower part of the trajectory shown in Figure 3.4, and eventually towards an equilibrium point in a different cycle. This phenomenon is known as *cycle slipping* and results in a probability distribution that looks like that shown in Figure 3.11(a). The probability density in the first cycle $[-\pi, \pi)$ generally decreases with increasing t (assuming the system starts in the first cycle), with the density in all other cycles correspondingly increasing. The rate of 'leakage' of probability from the first to secondary cycles depends on the level of the additive noise.

Rather than solving for the absolute phase $\phi \in (-\infty, \infty)$, we are instead more interested in the phase within a particular cycle, i.e., ignoring the effects of cycle-slipping. So, if we take the sum

$$p_{\pi}(\phi, t) = \sum_{n=-\infty}^{\infty} p(\phi + 2n\pi, t), \qquad (3.48)$$

then each term in the sum is a solution of (3.47) for the initial condition $\phi = \phi_0 + 2n\pi$, and so also is the sum, $p_{\pi}(\phi, t)$. $p_{\pi}(\phi, t)$ taken over just the interval $-\pi \leq \phi < \pi$ represents the wrapping of the full probability distribution $p(\phi, t)$ into a single cycle $[-\pi, \pi)$, an example of which is shown in Figure 3.11(b). Since, as shown in Figure 3.11(b), the full form (3.48) of $p_{\pi}(\phi, t)$ is periodic in ϕ with period 2π it is enough to solve the Fokker-Planck equation (3.47) for $p_{\pi}(\phi, t)$, with the initial condition

$$p_{\pi}(\phi, 0) = \delta(\phi - \phi_0),$$

boundary condition for all t,

$$p_{\pi}(\pi, t) = p_{\pi}(-\pi, t),$$

and normalizing condition for all t,

$$\int_{-\infty}^{\infty} p_{\pi}(\phi, t) \, d\phi = 1.$$

The solution of most interest generally is the steady-state solution, if it exists. This is

$$p_{\infty}(\phi) = \lim_{t \to \infty} p_{\pi}(\phi, t).$$

To check that such a solution exists we substitute this into the Fokker-Planck equation, which reduces the $\partial/\partial t$ term on the left-hand side to zero. That is, we obtain

$$\frac{d}{d\phi} \left[(\alpha \sin \phi - \beta) p_{\infty}(\phi) + \frac{dp_{\infty}(\phi)}{d\phi} \right] = 0, \qquad (3.49)$$

where

$$\alpha = \frac{4A}{KN_0},$$

and

$$\beta = \frac{4(\omega - \omega_0)}{K^2 N_0}$$

Equation (3.49) is readily solvable as

$$p_{\infty}(\phi) = C e^{\alpha \cos \phi + \beta \phi} \left[1 + D \int_{-\pi}^{\phi} e^{-\alpha \cos x - \beta x} \, dx \right], \qquad (3.50)$$

for $\phi \in [-\pi, \pi)$. Here, C and D are arbitrary constants of integration to be determined using the conditions

$$p_{\infty}(\pi) = p_{\infty}(-\pi), \qquad (3.51)$$

and

$$\int_{-\infty}^{\infty} p_{\infty}(\phi) \, d\phi = 1. \tag{3.52}$$

In the special case where the PLL's quiescent frequency is matched to the input signal frequency, we have $\omega = \omega_0$, and thus $\beta = 0$. In this case, the conditions above give

$$D=0,$$

and

$$C = \frac{1}{\int_{-\pi}^{\pi} \exp\left(\alpha \cos\phi\right) d\phi} = \frac{1}{2\pi I_0(\alpha)},$$

where I_0 is the zeroth-order modified Bessel function. This gives, for the special case of $\omega = \omega_0$, the solution

$$p_{\infty}(\phi) = \frac{\exp\left(\alpha\cos\phi\right)}{2\pi I_0(\alpha)}.$$
(3.53)

Clearly the solution is characterized entirely by the parameter α . For a first-order loop, from (3.17) the closed loop transfer function is AK/(s + AK), and the noise bandwidth of the loop is AK/4, using (3.42). The parameter α can be written as

$$\alpha = \frac{A^2}{N_0(AK/4)} = \frac{A^2}{N_0B_L}.$$

Since A represents the magnitude of the input signal, and N_0 the level of the input noise, the value of α is exactly the signal-to-noise ratio (SNR) within the loop bandwidth.



Figure 3.12: Steady-state phase error densities for first-order loop with $\omega = \omega_0$, and α values 1 (large dash), 3 (smaller dash), 10, 30, and 100 (solid).

The steady-state density is shown in Figure 3.12 for a number of different values of the signal-to-noise ratio. As expected for the case of $\omega = \omega_0$, the density is concentrated around the value $\phi = 0$; our work leading to Figure 3.4 predicts a stable equilibrium point at $\phi = 0$, the density in this case being a delta spike at $\phi = 0$. Indeed for large values of

the signal-to-noise ratio, such as $\alpha = 100$, the resulting density is a tall, narrow Gaussian-like curve centred at $\phi = 0$, as shown in Figure 3.12. For lower signal-to-noise ratios, the density is a wider and less peaked curve with more probability density at values away from $\phi = 0$. This corresponds to the additive input noise disturbing the operation of the loop so that the output phase error is thrown away from the noise-free locked state $\phi = 0$.



Figure 3.13: Steady-state output phase error variance from first-order loop with $\omega = \omega_0$ for (a) exact theory *(solid)*, and (b) linear model *(dashed)*.

The variance of the zero-mean output phase error process is calculated as

$$\sigma_{\phi}^{2} = \int_{-\pi}^{\pi} \phi^{2} p_{\infty}(\phi) \, d\phi = \frac{1}{2\pi I_{0}(\alpha)} \int_{-\pi}^{\pi} \phi^{2} e^{\alpha \cos \phi} \, d\phi. \tag{3.54}$$

The integral in (3.54) can be evaluated numerically, or alternatively, the exponential term may be expanded as a Fourier series, an expansion known as the Jacobi-Anger formula [11]. Thus, the variance may be written in the form:

$$\sigma_{\phi}^{2} = \frac{1}{2\pi I_{0}(\alpha)} \int_{-\pi}^{\pi} \phi^{2} \left[I_{0}(\alpha) + 2\sum_{n=1}^{\infty} I_{n}(\alpha) \cos n\phi \right] d\phi$$

$$= \frac{\pi^{2}}{3} + 4\sum_{n=1}^{\infty} \frac{(-1)^{n} I_{n}(\alpha)}{n^{2} I_{0}(\alpha)}.$$
 (3.55)

This series converges very rapidly, so only a few terms are needed for

accurate numerics. A plot of the output variance against the inverse of the SNR, $1/\alpha$, is shown in Figure 3.13. Note that the linear PLL theory in Section 3.1.6 predicted, from equation (3.43), $\sigma_{\phi}^2 = 1/\alpha$. The linear output variance, a line of slope 1, is also shown in Figure 3.13 for comparison. It can be seen from the plot that for high SNR (small noise), the linear model is a good match to the exact, non-linear theory. For moderate SNR, the variance output of the non-linear PLL is higher than that of the linear model. As the SNR approaches zero (high noise), the output noise variance, σ_{ϕ}^2 , approaches the asymptote $\pi^2/3$. This is exactly the variance that would result from the output phase error being uniformly distributed in the range $[-\pi, \pi)$; see equation (2.22), for example. In contrast, the output variance of the linear model grows unbounded.

For $\omega \neq \omega_0$ we have $\beta \neq 0$ and equation (3.50), together with conditions (3.51) and (3.52) must be solved numerically to obtain $p_{\infty}(\phi)$ and associated statistics, such as σ_{ϕ}^2 . As an example, we can take $\beta/\alpha = (\omega - \omega_0)/AK = \sin(\pi/4)$. From our work on the noise-free situation for the non-linear model, using (3.25) we would expect to find the phase error probability density to be centred around $\phi_0 = \pi/4$. Solving (3.50) and plotting $p_{\infty}(\phi)$ as before bears this out as shown in Figure 3.14.

3.1.6.2 Second-order loop

To examine the case of the second-order loop we first note that it is possible to derive, for a multi-dimensional Markov process $\mathbf{y}(t)$, where

$$\mathbf{y}(t) = [y_0(t), y_1(t), \dots, y_{n-1}(t)],$$

a vector equivalent of the Fokker-Planck equation (2.73). The derivation is sketched, for example, in [1], and yields



Figure 3.14: Steady-state phase error densities for first-order loop with $(\omega - \omega_0)/AK = \sin(\pi/4)$, and α values 1 (large dash), 3 (smaller dash), 10, 30, and 100 (solid).

$$\frac{\partial p(\mathbf{y},t)}{\partial t} = -\sum_{k=0}^{n-1} \frac{\partial}{\partial y_k} \left[A_k(\mathbf{y}) p(\mathbf{y},t) \right] + \frac{1}{2} \sum_{j=0}^{n-1} \sum_{k=0}^{n-1} \frac{\partial^2}{\partial x_j \partial x_k} \left[A_{jk}(\mathbf{y}) p(\mathbf{y},t) \right],$$
(3.56)

subject to the initial condition

$$p(\mathbf{y}, 0) = \prod_{k=0}^{n-1} \delta \left[y_k - y_k(0) \right]$$

for an initial vector of real values $[y_0, y_1, \ldots, y_{n-1}]$. In equation (3.56) the moments $A(\mathbf{y})$ are defined similarly to the scalar versions (2.72). That is, we have

$$A_{k}(\mathbf{y}) = \lim_{\Delta t \to 0} \frac{\mathscr{E}\left[\left(\Delta y_{k}\right) \mid \mathbf{y}\right]}{\Delta t},$$

$$A_{jk}(\mathbf{y}) = \lim_{\Delta t \to 0} \frac{\mathscr{E}\left[\left(\Delta y_{j} \Delta y_{k}\right) \mid \mathbf{y}\right]}{\Delta t}.$$
(3.57)

Once again, in the derivation of (3.56) it has been assumed that the process is sufficiently slowly-varying that all higher moments vanish, e.g.,

 $A_{jkl}(\mathbf{y}) = 0.$

The relevance of the vector Markov process to the second-order PLL is that, in general, an *n*-th order ODE can be written as a system of n first-order ODEs, which in this case are driven by white noise. For the 2nd order PLL, equation (3.38) becomes

$$\frac{d\phi(t)}{dt} = (\omega - \omega_0) - K \left[A\sin\phi(t) + \breve{n}(t)\right] - aK \int_0^t \left[A\sin\phi(u) + \breve{n}(u)\right] du,$$
(3.58)

where the loop filter is F(s) = 1 + a/s. This can be split as follows:

$$\frac{dy_0(t)}{dt} = y_1(t)$$

$$\frac{dy_1(t)}{dt} = -AK\sin\left[ay_0(t) + y_1(t)\right] - K\breve{n}(t),$$
(3.59)

where $\phi(t) = ay_0(t) + y_1(t)$. Then, $\mathbf{y}(t) = [y_0(t), y_1(t)]$ is a vector Markov process driven by white noise $\mathbf{\tilde{n}}(t) = [0, \mathbf{\tilde{n}}(t)]$, $\mathbf{\tilde{n}}$ having spectral density N_0 . Calculating the moments involved, the Fokker-Planck equation becomes

$$\frac{\partial p}{\partial t} = -y_1 \frac{\partial p}{\partial y_0} + \frac{\partial}{\partial y_1} \left[AK \sin\left(ay_0 + y_1\right)p \right] + \frac{K^2 N_0}{4} \frac{\partial^2 p}{\partial y_1^2}, \qquad (3.60)$$

subject to initial condition $p(y_0(t), y_1(t), 0) = \delta[\hat{y}_0 - y_0(0)]\delta[\hat{y}_1 - y_1(0)].$ Here $[\hat{y}_0, \hat{y}_1]$ is the fixed, initial value of $[y_0(t), y_1(t)].$

Equation (3.60) can be solved only by using advanced numerical methods (see [18]). However, using the substitution $z(t) = ay_0(t)$ and considering only the steady-state probability distribution $p_{\infty}(y_0, y_1) = \lim_{t\to\infty} p(y_0, y_1, t)$, equation (3.60) leads to

$$a(\phi - z) \left(\frac{\partial}{\partial \phi} + \frac{\partial}{\partial z}\right) p_{\infty}(\phi, z) = AK \frac{\partial}{\partial \phi} [\sin \phi \, p_{\infty}(\phi, z)] + \frac{K^2 N_0}{4} \frac{\partial^2 p_{\infty}(\phi, z)}{\partial \phi^2},$$
(3.61)

which, when we integrate over all z to obtain the marginal density

 $p_{\infty}(\phi) = \int_{-\infty}^{\infty} p_{\infty}(\phi, z) dz$, becomes

$$a\left[\frac{d}{d\phi}[\phi p_{\infty}(\phi)] - \frac{d}{d\phi}\int_{-\infty}^{\infty} z p_{\infty}(\phi, z) dz\right] = AK \frac{d}{d\phi}[\sin\phi p_{\infty}(\phi)] + \frac{K^2 N_0}{4} \frac{d^2 p_{\infty}(\phi)}{d\phi^2}.$$
(3.62)

The second term on the left, $\int_{-\infty}^{\infty} z p_{\infty}(\phi, z) dz$ can also be written as

$$\int_{-\infty}^{\infty} z p_{\infty}(\phi, z) \, dz = p_{\infty}(\phi) \mathscr{E}(z \mid \phi).$$

Since $z = \phi - y_1$, this last expectation in turn is $\phi - \mathscr{E}(y_1 | \phi)$, and, going back to the original equations (3.59) for y_0, y_1 , the expectation is found to be

$$\mathscr{E}[y_1(t) \mid \phi(t)] = AK \int_t^\infty \mathscr{E}[\sin \phi(u) \mid \phi(t)] \, du.$$

Inserting now back in (3.62), we get an ODE for the steady-state density of the phase error, ϕ that involves terms in ϕ only:

$$\frac{d}{d\phi} \left(\frac{4A}{KN_0} \left[\sin\phi - a \int_0^\infty \mathscr{E}[\sin\phi(t+\tau) \,|\, \phi(t)] \,d\tau \right] p_\infty(\phi) + \frac{dp_\infty(\phi)}{d\phi} \right) = 0$$
(3.63)

Like the case of the first-order loop, we can take ϕ in equation (3.63) to be in the range $[-\pi, \pi)$ by wrapping all values of the phase back into the first cycle of the PLL. The same boundary and normalization conditions on $p_{\infty}(\phi)$ apply as did in the case of the first-order equation.

Even equation (3.63) is not particularly amenable to analysis, though some useful approximate results can be found. First of all, note that as $a \rightarrow 0$ equation (3.63) reduces to the corresponding equation for the firstorder loop (3.49) with $\beta = 0$. This has solution (3.53) as before, which is to say that the second-order loop equation reduces to the first-order loop with $\omega = \omega_0$ as we would expect.

For an arbitrary value of a, for large SNR ϕ will be small and so we can use the approximation $\sin \phi \approx \phi$. This is the same approximation as was used for the linear PLL but now we use it in equation (3.63),

which was derived using the non-linear analysis. Then, for wide-sense stationary ϕ , the expectation within the integral is proportional to the autocorrelation function $R_{\phi}(\tau)$. That is

$$\int_0^\infty \mathscr{E}[\sin\phi(t+\tau) \,|\, \phi(t)] \,d\tau \approx \frac{1}{2\sigma^2} \left[\int_{-\infty}^\infty R_\phi(\tau) \,d\tau \right] \sin\phi(t), \quad (3.64)$$

where σ^2 is the variance of the process ϕ . Using the Wiener-Khintchin theorem, the integral in (3.64) is found to be related to the spectral density of the process, ϕ . Specifically, we have

$$\int_{-\infty}^{\infty} R_{\phi}(\tau) \, d\tau = S_{\phi}(0).$$

Because we've used the linear approximation, we can use the results from Section 3.1.5 to find the spectral density. For the second-order loop with F(s) = 1 + a/s, equation (3.40) gives

$$S_{\phi}(\omega) = \frac{N_0 K^2}{2} \left| \frac{i\omega + a}{-\omega^2 + AKi\omega + aAK} \right|^2,$$

which gives $S_{\phi}(0) = N_0/2A^2$ and

$$\sigma^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} S_{\phi}(\omega) \, d\omega = \frac{N_0}{4A^2} (AK + a).$$

Using these simplifications, equation (3.63) now reduces to

$$\frac{d}{d\phi} \left(\frac{4A}{KN_0} \left[\sin \phi \left(\frac{AK}{AK+a} \right) \right] p_{\infty}(\phi) + \frac{dp_{\infty}(\phi)}{d\phi} \right) = 0.$$
(3.65)

This is again has the same form as the equation for the first-order density (3.49), with $\beta = 0$ and modified SNR α :

$$\alpha' = \frac{A^2}{N_0(AK+a)/4} = \frac{A^2}{N_0B'_L}.$$

Note that α here has the same form as for the first-order loop, now inversely proportional to the bandwidth of the second-order loop, B'_L .

The solution of (3.65) has exactly the same form as found previously, namely

$$p_{\infty}(\phi) = \frac{\exp\left(\alpha'\cos\phi\right)}{2\pi I_0(\alpha)}.$$
(3.66)

Thus, the approximate solution indicates that the response of the second-order loop is identical to that of the first-order loop with $\omega = \omega_0$. This is borne out by empirical results; the effect of changing from a first- to a second-order loop is to eliminate the static phase error at the expense of some loop stability. However, when the loop does lock, the noise output from the second-order system is similar to that of the first.

3.2 Digital PLL summary

In modern communications and other electronic devices, phase-locked loops are more commonly implemented in digital form rather than in the traditional analogue version, summarized in Section 3.1. The predominance of the digital version is for reasons of cost, flexibility of design, deterministic behaviour, and the tendency towards all-digital systems, of which PLLs are a component [7].

A digital PLL (DPLL) may be implemented on dedicated hardware, in which case the circuitry replaces the equivalent analogue circuit of lumped elements, such as resistors, capacitors and oscillators. In this case, the input signal is sampled using an *analogue-to-digital converter* (ADC), with the PLL operations taking place digitally on these samples. The output of the loop is another set of samples representing the output signal, which may be fed into a *digital-to-analogue converter* (DAC) to convert back to an analogue signal if required.

Alternatively, the PLL may be run simply as an algorithm on a general-purpose computer or processor, taking an input that is a sequence of samples of the input signal and providing a corresponding sequence of output signal samples as an output. This scenario is most often used in all-digital systems, where the input signal is sourced in the previous digital block, and the DPLL's output samples are fed directly into the next processing block in the chain. No conversion from or to analogue signals are needed in this case.

In either case, all DPLL operations are computations, i.e., the operations of phase-detector, linear filter and VCO all take place as digital operations within the circuit or processor. These are generally sequences of basic operations such as additions, subtractions and multiplications. All operations take place on sequences of samples, representing the input signal, to produce another set of samples, representing the output. The input samples, output samples, and intermediate values used in computations are digital numbers, and thus quantized. In particular, the intermediate representing the "voltage" to be fed to the VCO is quantized, which means the frequency of the output signal is similarly quantized. Quantization at any point in the loop introduces an immediate non-linearity in the system. The effects of quantization may be reduced by increasing the resolution of the digital values, e.g., by storing all samples and intermediates as 64- or 128-bit values. This, however, adds complexity and cost to the circuit, and increases power requirements. A common approach when analyzing the behaviour of DPLLs is to assume that the effects of external, additive noise will always be much greater than the effects of quantization. In this case, quantization can be ignored, and the DPLL essentially treated like an analogue loop. However, as we will see later, this assumption is not always valid, and often additive and quantization noises can interact in unexpected ways.

The following sections provide a detailed introduction to DPLLs, based largely on the new material added in the latest edition of the book by Gardner [2], leading on to the work by Teplinksy et al. [4], [5]. This material provides the background to our new results, which are presented from Chapter 4 onwards. As is standard in these publications, the alldigital loop we examine is referred to in this thesis simply as a "DPLL", whereas in some texts [7] this is classified as an *(all-digital) ADPLL*. This is to distinguish it from earlier incarnations of loops that had digitization in the phase-detector only [7]; we don't make the distinction in this work.

3.2.1 Basics



Figure 3.15: Block diagram of a first-order DPLL.

The basic elements of a first-order DPLL are shown in Figure 3.15. The corresponding diagram for the analogue version of this loop is Figure 3.1, where the loop filter is absent so that the straight-through path gives $e(t) \equiv x(t)$. The DPLL may be used in an all digital circuit, in which case the input consists of a sequence of discrete-time quantized representations of a sinusoidal signal, $\sin \phi_i(t)$, for $t = 0, 1, 2, \ldots$ Alternatively, the DPLL may be part of an analogue signal receiver, in which case the input signal is fed to an ADC, which samples the signal and produces an equivalent set of discrete samples. In either case, the operation of the DPLL from the phase detector onwards is the same. The analogue-todigital conversion option is shown on the top left of Figure 3.15.

The components of the digital loop directly correspond to their analogue counterparts from Figure 3.1. The principal differences are:

- The input and output of each component are series of discrete values representing signal samples, whereas in the analogue case, the signals themselves are transmitted across the loop.
- All sample values are quantized and all operations take place on quantized values using finite-precision arithmetic. This results in

possible truncation of results, which represents an immediate nonlinearity.

On the first point, we can write the discrete time samples of a signal x(t) as the set $\{x(t_0), x(t_1), \ldots, x(t_i), \ldots\}$. However, we generally use equally spaced time instants so that the interval can be normalized and we can deal with signal samples $\{x(t) | t = 0, 1, 2, \ldots\}$. The second point above deserves further elaboration. The simplest component in the loop shown in Figure 3.15 is the scaler K. This corresponds to an analogue attenuator or amplifier, and simply multiplies each input sample by K. Thus, in the analogue case, the operation of the scaler is simply

$$e(t) = Kx(t),$$

where x(t) is the output of the phase detector and e(t) is the output of the scaler. In the digital case, if we assume $x(t_i)$ is the *b*-bit quantized phase detector output and the output of the scaler, $e(t_i)$, is also *b*-bit quantized, then the corresponding relationship is

$$e(t_i) = \frac{1}{2^b} \operatorname{Int}[2^b K x(t_i)],$$

where Int[x] denotes the integer part of x.

All other loop operations are subject to similar truncation. As a result, the equations describing the overall loop behaviour are both nonlinear and unwieldy, and not particularly amenable to analysis. A common approach in the analysis of digital systems is to assume the effects of quantization are negligible. This is usually achieved in practice by using a sufficiently high number of bits, b. If this number is high enough, the nonlinear loop noise introduced by quantization is often much smaller than the magnitudes of either the signals themselves or external additive noise. In such cases it is often valid to ignore the quantization noise. For the initial stages of this introduction to digital loops we will ignore the quantization as standard. However, we will re-introduce it later in Section 3.2.3. The digital phase detector may be one of several types, but is typically based on a multiplier, which produces, at each time instant t the output

$$\sin [\phi_i(t) - \phi_o(t)] + \sin [\phi_i(t) + \phi_o(t)], \qquad (3.67)$$

for t = 0, 1, 2, ... Similar to the analogue case, the second, doublefrequency term in (3.67) is usually virtually eliminated by either the loop filter or an internal filter in the PD, so that the PD output is

$$\sin\left[\phi_i(t) - \phi_o(t)\right] = \sin\Phi(t). \tag{3.68}$$

In cases where the phase error remains small we can use the linear approximation $\sin \Phi(t) \approx \Phi(t)$ in the loop equations as we did in the analogue case. In other loops, the phase detector may simply output $\phi(t) = \phi_i(t) - \phi_o(t)$ always, by performing appropriate operations on the inputs $\sin \phi_i(t)$ and $\sin \phi_o(t)$. However, this is less common; multiplier-type PDs are generally easiest to implement in both analogue and digital circuits (see, e.g., [2], Ch. 13).

After scaling the output by K, the error signal enters the delay component, which delays the input value by D clock cycles, where $D \ge 1$. Thus, the input to the NCO at time t is

$$K\sin\left[\Phi(t-D)\right].\tag{3.69}$$

In any digital loop, there must be at least one delay element. This is to maintain causality, i.e., the output ϕ_0 fed back into the phase detector at time t can only affect subsequent values of ϕ_0 . Otherwise, the value of $\phi_o(t)$ could not be calculated until the value of $\phi_o(t)$ was known! Also, in general a real-time digital component cannot produce an output value based on an input value at the same clock cycle. Usually the component will take in an input value at time t to produce an output at the next clock cycle t + 1, or at some later time again. This often results in delays of several clock cycles through the loop operations. For example, the scaler may produce e(t) = Kx(t-1), where x(t) is the output of the phase

detector. This per-component delay may not be present, for example, in a post-processing situation, where the PLL algorithm is run on a list of digital samples that have all been captured before the algorithm is run. In other cases, a real-time component may produce its output at a fractional clock cycle later than the input, in which case it may be well modelled as producing the output instantaneously. However, there always needs to be a delay of at least one clock cycle across the entire circuit to maintain a causal loop.

Here we model the delays in the various components by combining them all in a single delay element before the NCO. Not all DPLL models allow for the assorted loop delays to be combined in this manner, but this was the only model considered by Gardner in his most complete work on DPLLs [2]. The model covers the case where the delay is positioned as shown in Figure 3.15, or in the output of the phase detector and\or scalers K, K_1 in Figures 3.15 and 3.16. To keep the analysis tractable, this model of loop delay is the only one we use in this thesis. The delay element is over D clock cycles, where $D \ge 1$ for causality. Most of our results in this thesis are based on the simplest case of D = 1, but we do also consider the non-trivial delayed D > 1 case in Chapter 5.

A numerically-controlled oscillator (NCO) is the digital equivalent of a VCO, that is a component that produces discrete samples of the signal $\sin \phi_o(t)$. The frequency of the signal, the rate of change of phase $[\phi_o(t) - \phi_o(t-1)]$, is proportional to the NCO input value or control word, x(t), where x(t) is given by (3.69), i.e.,

$$\phi_o(t) = \phi_o(t-1) + 2\pi K_v x(t) \mod 2\pi.$$
(3.70)

Without loss of generality we may take $K_v = 1$ by including it in the scaling factor K. Note that the operation of the NCO is essentially the same as that of the digital integrator given by equation (2.89). This is to be expected as the input is proportional to the output signal's frequency, and the output is the signal phase. Now, if we take the input signal to be a sinusoid of constant frequency, ν , then we have

$$\phi_i(t) = 2\pi t\nu \mod 2\pi,\tag{3.71}$$

and so

$$\phi_i(t) = \phi_i(t-1) + 2\pi\nu \mod 2\pi.$$
(3.72)

Combining equations (3.68), (3.69), (3.70), and (3.72), we obtain finally a difference equation for the phase error Φ :

$$\Phi(t) = \Phi(t-1) + 2\pi \left[\nu - K\sin\Phi(t-D)\right] \mod 2\pi.$$
(3.73)

The equations for the input signal phase, ϕ_i , and output phase, ϕ_o , use the modulo operation to remove any differing integer number of cycles in each since time t = 0. While the DPLL is in the locked state, or close to locked, we expect ϕ_o to remain forever close to ϕ_i , and certainly to remain on the same cycle. Thus, for the purposes of this thesis, where we expect the phase difference, Φ , to remain close to zero, the modulo- 2π operation is redundant; this is needed only to handle the case of cycle slipping.

In the simplest case where D = 1, (3.73) becomes an equation for Φ where the value at time t depends only on the previous value at time t - 1. If we shrink the normalized time interval in (3.73) back towards the case where we have continuous time, the equivalent equation would be

$$\frac{d\Phi}{dt} \approx 2\pi \left[\nu - K\sin\Phi\right]. \tag{3.74}$$

It should be noted that, up to a constant, this is the same as equation (3.23) for the analogue, first-order non-linear PLL obtained in Section 3.1.3.

The second-order DPLL is shown in Figure 3.16. It is identical to the first-order loop apart from the inclusion of an additional integrator on the right-hand side of the diagram. The equation of operation for the



Figure 3.16: Block diagram of a second-order DPLL.

first-order DPLL (3.73) is modified accordingly:

$$\Phi(t) = \Phi(t-1) + 2\pi \left[\nu - K_1 \sin \Phi(t-D) - u(t-D)\right] \mod 2\pi, \quad (3.75)$$

where u(t) is the output of the integrator at time t. The integrator output is given by

$$u(t) = u(t-1) + K_1 K_2 \sin \Phi(t).$$
(3.76)

For D = 1, these become a pair of first-order equations for the loop, as can be seen by taking the continuous-time equivalents:

$$\frac{d\Phi}{dt} \approx 2\pi \left[\nu - K_1 \sin \Phi - u\right],$$

$$\frac{du}{dt} \approx K_1 K_2 \sin \Phi.$$
(3.77)

These can be combined to form a single, second-order equation for $\Phi(t)$:

$$\frac{d^2\Phi}{dt^2} \approx -2\pi K_1 \cos\Phi \frac{d\Phi}{dt} - K_1 K_2 \sin\Phi.$$
(3.78)

This equation is again a version of equation (3.27) found earlier for the analogue, non-linear second-order PLL. In general, a digital PLL with n integrators (including the NCO) will be described by a system of n first-order equations, for the phase error Φ and (n-1) other state variables.



(c) Path to steady-state.

Figure 3.17: (a) Time-domain plot of phase error for second-order DPLL with $\nu = 0.0021875$, $K_1 = 2^{-6}$, $K_2 = 2^{-5}$, and $\{\Phi(0), u(0)\} = \{0, 0\}$, (b) corresponding integrator output, (c) path to steady-state on phase plane.

Figure 3.17 shows an example of the behaviour of the second-order DPLL. The plots show the dynamics of the system as is starts from $\{\Phi, u\} = \{0, 0\}$ at t = 0. Plot (a) is a time-domain plot of the phase error, while plot (b) is a similar plot for the integrator output. It can be seen that the phase error initially grows from 0 until a maximum value is reached, while the integrator output climbs towards steady-state. As the steady-state integrator output is approached, the phase error starts to decrease again and is pulled in towards zero. This is the same behaviour as we would get with the second-order analogue PLL of Section 3.1.3, where the frequency of the input signal differs from the quiescent frequency of the VCO (the latter is normalized to zero here in the case of a DPLL). This is as would be expected since the DPLL is simply a

time-discrete version of the analogue case. The corresponding behaviour for the analogue PLL can be seen by taking a trajectory in Figure 3.5 that begins on the positive- $\dot{\phi}$ axis ($\phi = 0, \dot{\phi} > 0$); on such trajectories ϕ increases initially before being pulled into the stable equilibrium point at $\phi = 0$. The steady-state integrator output is that which is needed to keep the NCO frequency-locked to the non-zero input frequency.

Figure 3.17(c) shows the combined behaviour on the $(\phi - u)$ phase plane. We will return to similar phase-plane plots for the frequencyquantized DPLL in the following sections.

3.2.2 Linear approximation

The digital PLL can be linearized in exactly the same way as outlined in Section 3.1.2 for the analogue version. If the phase error $\Phi(t)$ remains small for all t, then we can use the approximation $\sin \Phi(t) \approx \Phi(t)$. Then all elements in the digital loop are linear and can be represented as operators in the Z-domain as shown in Figure 3.18. Here, the Z-transforms of the digital integrator and delay element are $1/(1 - z^{-1})$ and z^{-D} respectively, from (2.93) and (2.94) in Section 2.3.2.

With linearization, equation (3.70) for ϕ_o in the case of the first-order loop becomes

$$\phi_o(t) = \phi_o(t-1) + 2\pi K \left[\phi_i(t-D) - \phi_o(t-D)\right] \mod 2\pi.$$
(3.79)

This is a linear equation, and Z-transforms can be used to obtain a closed-form solution in the Z-domain

$$\Phi_o(z) = H(z)\Phi_i(z), \qquad (3.80)$$

where $\Phi_i(z)$ and $\Phi_o(z)$ are the Z-transforms of the corresponding timedomain sequences. For this first-order case we have

$$H(z) = \frac{2\pi K z^{-D}}{1 - z^{-1} + 2\pi K z^{-D}}.$$
(3.81)

The Z-transform of a sinusoidal input with constant frequency ν is

$$\Phi_i(z) = \frac{2\pi\nu z}{(z-1)^2},$$

so that, for example, the output of the first-order linear loop with minimal delay D = 1 is

$$\Phi_o(z) = \left[\frac{2\pi K z^{-1}}{1 + (2\pi K - 1)z^{-1}}\right] \left[\frac{2\pi \nu z}{(z - 1)^2}\right].$$

Inverse transforms may be used to to obtain the time-domain output, $\phi_o(t)$. When $\Phi_o(z)$ is a rational function of z (ratio of polynomials) as it is in this case, the inverse transform may be calculated using elementary methods, i.e., by re-writing $\Phi_o(z)$ as a partial fraction expansion [23] and transforming each of the simpler terms individually. However, evaluation in this case is cumbersome and is best performed using a computer package such at Matlab or Mathematica. More often, once the system has been transformed to the Z-domain, the analysis tends to remain there rather than being inverse transformed back to the time domain [23].



Figure 3.18: Z-domain representation of linearized second-order DPLL.

Similarly to the analogue case, the linearization of the DPLL allows us to use the principle of superposition to consider separately the action of the system on the sums of input signals. In particular, if the input
signal is subject to phase noise, $n_i(t)$ with Z-transform $N_i(z)$, then the output is

$$\Phi_o(z) + \alpha N_o(z) = H(z) \left[\Phi_i(z) + \alpha N_i(z) \right],$$

where $\Phi_o(z) = H(z)\Phi_i(z)$ and $N_o(z) = H(z)N_i(z)$.

This predictability of linear loops yields huge simplifications in the analysis. However, in this thesis we focus mainly on systems where this linearity is broken. Examples of such systems are illustrated in the following sections, these being central to our results which are presented from Section 4 onwards.

3.2.3 Quantization effects

In an all-digital loop, any signal sample passed between elements is represented by a digital value and thus subject to some quantization. For example, the phase detector output, $\sin \Phi$ may be passed to the scaler as a *b*-bit value

$$\frac{1}{2^b} \operatorname{Int}[2^b \sin \Phi], \qquad (3.82)$$

where $\operatorname{Int}[x]$ denotes the integer part of x. In general, this means the PD output will not exactly equal $\sin \Phi$, and the difference between $\sin \Phi$ and (3.82) represents the quantization or rounding error. The errors in the loop may be made arbitrarily small by increasing the number of bits, b. However, this increases the cost and complexity of components. In general, a value of b in the range 16-64 is used depending on the application.

The presence of the quantization in the output of each component, as in (3.82), greatly complicates the DPLL equations. This is because the rounding error from each block enters the next block as an input, so the quantization accumulates across the loop. For example, assuming all loop elements are *b*-bit quantized, the NCO output for the first-order loop (3.70) may be modified as

$$\phi_o(t) = \frac{1}{2^b} \operatorname{Int}[2^b \phi_o(t-1)] + \frac{1}{2^b} \operatorname{Int}\left[\frac{2\pi}{2^b} \operatorname{Int}[2^b \sin \Phi(t-D)]\right] \mod 2\pi.$$
(3.83)

Such expressions are unwieldy and direct use of quantization in the loop equations makes analysis impossible. The situation is complicated further by the fact many components use intermediate values internally in order to produce their outputs. The PD, for example, may digitally multiply the two signals $\sin \phi_i(t)$ and $\sin \phi_o(t)$, and then apply some filtering to eliminate the high-frequency term. This would involve performing several other additions and multiplications, each stage of which would result in an intermediate value being stored as a quantized value. Thus, in practice, the PD output would often not exactly equal the value given by (3.82).

The effect of the accumulated quantization in the loop is an additional contribution towards the phase error, $\Phi(t)$. Thus, even with no external additive noise, the output signal may never lock exactly to the input signal because of the quantization in the NCO output, (3.83) for example. This may be viewed as a form of noise or 'quantization jitter' in the phase error. However, because the effect of the loop is to drive the phase error Φ towards zero, and the quantization effects are generally small, the overall effect is not to drive the loop away from lock entirely, rather the introduction of this quantization jitter around the lock point.

A common approach to handling the complications introduced by quantization is to handle it as a form of additive noise in the loop equations. The simplest model is to treat this 'quantization noise' as AWGN of an appropriate level. This relies on the external signal or additive noise being large enough that the quantization is uncorrelated with it, so the quantization effects can be added in a linear fashion. However, this small external stimulus approximation is most often not true in practice as quantization may be the dominant effect. Furthermore, the quantization is inherently non-linear and, as we will see later, interacts with external additive noise in a decidedly non-linear way. A study model chosen by Gardner [3] as well as Teplinsky et al. [4] is to concentrate the effects of quantization into a single block before the NCO. The quantization block could be placed at any location in the loop and the analysis would be similar. However, NCO quantization seems to be the predominant model in the literature, with little having been published to date on any other form of quantization. Gardner also argues [2] that in certain circumstances, quantization in the output of the phase detector or loop filter may be combined with the quantizer at NCO input in which case it can also be treated using this same model, though possibly with a lower number of bits. This "frequency quantization" model is the one we use throughout our results in this thesis. This means that the NCO control word is a quantized value and can only take on a finite number of values, and the same is then true of the NCO's output frequency.

3.2.3.1 First-order loop

A block diagram of a first-order DPLL with frequency quantization is shown in Figure 3.19. Following [4], we assume a noise-free sinusoidal input signal of frequency ν so that its phase, $\phi_i(t)$ in Figure 3.19, is

$$\phi_i(t) = 2\pi t\nu \mod 2\pi,\tag{3.84}$$

for t = 0, 1, 2, ... The quantizer produces a *b*-bit quantized version of the loop filter output $x_f(t)$

$$x_q(t) = \frac{1}{2^b} \text{Int}[2^b x_f(t)],$$
 (3.85)

where Int[x] denotes the integer part of x. The quantized loop filter output, x_q , is used as the input control word for the NCO, whose output frequency is proportional to the input, so that its phase is

$$\phi_o(t) = \phi_o(t-1) + 2\pi x_q(t) \mod 2\pi.$$
(3.86)



Figure 3.19: Block diagram of a first-order DPLL with frequency quantization.

Since x_q is quantized, so also are the NCO phase and frequency. This implies that the NCO output signal will, in general, not be able to lock exactly to the input signal. Instead, the difference between the phases of the input and output signals, the phase error $\Phi(t)$, will incur a jitter about the locked position.

For the first-order loop in Figure 3.19, and taking the simplest case initially of the minimal delay, D = 1, the loop filter output x_f is

$$x_f(t) = K \sin \Phi(t-1),$$
 (3.87)

where K is the loop gain. Combining equations (3.84)-(3.87) we have, finally, as in [4]

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \left(\mu - \text{Int}[2^b K \sin \Phi(t)] \right) \mod 2\pi \text{ for } t = 0, 1, 2, \dots$$
(3.88)

where $\mu = 2^{b}\nu$. This is a non-linear difference equation for the phase error of a frequency-quantized first-order DPLL in the absence of external additive noise. A time-domain plot of this phase error is shown in Figure 3.20(a). From the graph, it would appear that Φ initially converges to a range of values in which it remains indefinitely, approximately [0.245, 0.265] in the example shown. This is in contrast to the analogue case where the phase error converged to a single value if the input frequency was small enough; see, for example, Figure 3.4. However, note from Figure 3.4 that in general the first-order analogue loop settles to a non-zero value of Φ , i.e., incurs a static phase error. Similarly, it can be seen that, although the loop does not settle exactly, it does converge to a range of values not necessarily centred about $\Phi = 0$. We will see later how the dynamics in this steady state are described by a map on the circle provided the input frequency ν is sufficiently small, the jagged nature of the trajectory corresponding to the steps in the VCO output frequency needed to minimise the phase error.



Figure 3.20: (a) Time-domain plot of phase error for frequency-quantized first-order DPLL with b = 8, $\mu = 0.56$, $K = 2^{-6}$, and $\Phi(0) = 0.3$, (b) similar plot for DPLL without quantization.

For now, note also the similarity of equation (3.88) to the one found earlier (3.73) where no quantization was involved; this is a simplification brought about by concentrating all quantization effects on the NCO input. The behaviour of the corresponding non-quantized map (3.73)for D = 1 is shown in Figure 3.20(b) for comparison. The similarities in behaviour between the two systems is clear: starting from the initial condition, the phase error quickly tends towards a constant, non-zero value. However, for the non-quantized case (b), the the transition occurs smoothly, without the 'sawtooth' pattern of the quantized system. Also, steady-state is reached more slowly, and to a different static phase error than in the quantized case. Because of the truncation in equation (3.88), this higher average phase error in steady-state in (a) is required to keep the NCO's average output frequency equal to that of the input signal.

The phase error trajectory may also be visualized on the $\Phi(t+1) - \Phi(t)$

phase plane, as was done in [4]. This is shown in Figure 3.21, where equation (3.88) is graphed as the bold line segments. The map (3.88) can be iterated by stepping between the bold line segments and the identity line $\Phi(t+1) = \Phi(t)$. Steady-state is reached when the graph of $\Phi(t+1)$ lies on both sides of the identity line.

We shall investigate the quantized first-order system more fully in Section 3.2.5.



Figure 3.21: Plots of phase error trajectory on $\Phi(t) - \Phi(t+1)$ plane for frequency-quantized first-order DPLL with b = 8, $\mu = 0.56$, $K = 2^{-6}$, and $\Phi(0) = 0$.

3.2.3.2 Second-order loop

For the frequency-quantized second-order DPLL in the absence of external additive noise with D = 1, equations (3.75) and (3.76) found earlier are similarly modified as:

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \left(\mu - \operatorname{Int}[2^b (K_1 \sin \Phi(t) + u(t))] \right) \mod 2\pi,$$

$$u(t+1) = u(t) + K_1 K_2 \sin \Phi(t+1),$$
(3.89)

where u(t) is again the integrator output.

The steady-state dynamics of this particular system were studied in great detail by Teplinsky et al. in [4], [5]. In these papers, the precise nature of the periodic and non-periodic limit cycles of the map were described for a range of loop parameters, bounds found on the phase error and integrator output etc. Because the first-order loop is simply a special case of second-order loop ($K_2 = 0$) a study of the second-order DPLL also reveals details of the corresponding first-order loop behaviour.



Figure 3.22: (a) Time-domain plot of phase error for frequency-quantized second-order DPLL with b = 8, $\mu = 0.56$, $K_1 = 2^{-6}$, $K_2 = 2^{-5}$, and $\{\Phi(0), u(0)\} = \{0, 0\}$, (b) corresponding integrator output, (c) path to steady-state region on phase plane, (d) steady-state phase error and integrator output on phase plane.

For now the second-order loop behaviour can be summarized by the various plots in Figure 3.22. In this, it can be seen from graph (a) that, although the phase error, Φ first jumps from its zero value to a positive value, it eventually tends towards a corridor around $\Phi = 0$. This is in contrast to the first-order loop, which remained forever in a corridor centred at a non-zero Φ . However, for the second-order loop, both its eventual limit cycle and the path taken towards that limit, follow a similar 'sawtooth' pattern as the first-order case. The force that pulls the phase

error back towards zero for the second-order system is provided by the integrator, whose output climbs gradually towards a limit cycle that is reached as Φ nears zero. This is shown in plot (b). This "restoring force" is the component that is missing from the first-order system in order to cancel the static phase error.

The similarities and differences between these plots and those in Figure 3.17 for the quantization-free case should be noted. While the macro behaviour is identical there are some significant differences in the detail. For example, the maximum phase error excursion for the quantized case is over twice as large as the corresponding value for the DPLL without quantization. This is because, with quantization, the phase error initially increases by an amount $2\pi\nu$ until the quantizer output becomes 1. This can be seen by an examination of equations (3.89). However from equation (3.75), without the quantizer in place, the increment in the phase error starts to get smaller as soon as the integrator output becomes non-zero. The result is a much smaller maximum excursion in Φ and a smaller steady-state value of the restoring force, u.

Using numerical simulations only, Gardner made the following observations in [2] which are applicable to the limit cycles of both first- and second-order loops:

- 1. The pattern of the limit cycles (of both phase error and integrator output u(t), for the second-order loop) are dependent on $\operatorname{Frac}[\mu]$, where $\operatorname{Frac}[x]$ denotes the fractional part of x.
- For the second-order loop, the limit cycles are independent of Int[μ].
 For the first-order system, the limit cycles are also independent of Int[μ] apart from the particular non-zero value of Φ the phase error settles to (static phase error). However, for this loop we must have Int[μ] ≤ 2^bK − 1 for the loop to achieve lock.
- 3. Apart from amplitude and any static phase error, the limit cycle waveforms are also unaffected by any changes to b, K, K_1, K_2 if the value of μ is maintained.

- If Frac[μ] is rational and is equal to p/q in its lowest form, then the limit cycles are periodic with period q.
- 5. If $\operatorname{Frac}[\mu]$ is irrational then the limit cycles are not periodic, but are instead said to be *quasi-periodic*; the values do not repeat exactly but are instead densely distributed within a finite range. In the general case where the input frequency is independent of the PLL clock, μ will be irrational with probability 1.
- 6. For loop parameters of practical interest the NCO output signal takes only two output frequencies, i.e., $Int[\mu]$ and $1 + Int[\mu]$. The loop adjusts the frequency between the two values such that the average output frequency is exactly μ . The loop parameters of interest have K, K_1 , K_2 sufficiently small as to allow stable lock and non-integer input frequency μ .
- 7. For most loop parameters the range (peak-to-peak excursion) of the phase error Φ is approximately $2\pi D/2^b$.
- 8. There is no straightforward formula or approximation for the output phase error variance in the case where additive input noise is combined with the internal quantization jitter. However, for input noise levels greater than $\sigma_N^2 = (1/2^b K_1)^2$, the additive noise dominates and the digital PLL behaves like an analogue PLL.

We re-visit some of these findings in our results Chapters 4, 5, and 6.

Of particular interest for the second-order case is the combined $\Phi - u$ behaviour in steady-state, shown in plot (d) of Figure 3.22. The dynamics in this region were studied in great detail in [4], [5]. In these papers, it was shown that, from its initial state, the system transitions through a corridor as shown in plot (c) to a "trapping region" from which it never emerges. Within this trapping region, it finally settles to a region that is invariant under the map. The dynamics within the region are a combination of circle maps in Φ and u, the exact details being dependent on the system parameters, μ , b, K_1 , K_2 , and the initial conditions. An example of the invariant region dynamics are shown in plot (d). We will go into more detail on the circle map dynamics in the following sections.

Of most interest to engineers is the range, or magnitude, of the phase errors in this trapping region. However, the precise nature of the limit cycles are also of interest since periodic behaviour often introduces unwanted spectral effects. An understanding of the dynamics in this regime is also crucial for analysis of how external additive noise might interact with these quantization effects. We will see this later from Chapter 4 onwards.

3.2.4 Circle rotation map preliminaries

In this section we provide the definition of the circle rotation map as well as outlining some of its basic properties. Maps of the circle are well studied in the literature – further reading on circle rotation maps and similar dynamical systems is available, for example, in [28] – and what is presented here is just a summary of the existing material using our notation.

The *circle rotation map* on $[0, 2\pi)$ is defined as:

$$x(t+1) = \{x(t) + 2\pi\eta\} \mod 2\pi, \text{ for } t = 0, 1, 2, \dots$$
(3.90)

The circle rotation map is periodic with period q if η is rational and equal to p/q in its lowest form. Otherwise, the mapping is quasi-periodic and is dense in $[0, 2\pi)$. The parameter η is called the *rotation number* of the map.

In this thesis we frequently deal with circle maps that are placed arbitrarily on the real line. We can define a circle rotation map on [M, M+S)as

$$x(t+1) = \begin{cases} x(t) + \alpha & \text{for } M \le x(t) < M + S - \alpha \\ x(t) + \alpha - S & \text{for } M + S - \alpha \le x(t) < M + S, \\ & \text{for } t = 0, 1, 2, \dots \end{cases}$$

This second form of the circle map is equivalent to the first except shifted from $[0, 2\pi)$ to [M, M+S) and is periodic if the rotation number, $r = \alpha/S$, is rational, and quasi-periodic otherwise.

Lemma 1. For a circle rotation map with parameters (M, S, α) as before, where $r = \alpha/S$ is rational and equal to p/q in its lowest form, the arithmetic mean value of the map is

$$(M+\gamma) + \frac{(q-1)S}{2q},$$

and the variance is

$$\frac{(q^2 - 1)S^2}{12q^2}$$

where $\gamma \in [0, \frac{1}{q})$ is the minimum value of the map on [M, M + S).

Proof. Note that the "mean" and "variance" here are those obtained using time-averages, i.e., are not the expectations of Section 2.1.4 since the circle map output is not a random variable. The proof is straightforward once it is noted that the set of values taken on by the map over a single period is $\{M + \gamma, M + \gamma + \frac{S}{q}, M + \gamma + \frac{2S}{q}, \ldots, M + \gamma + \frac{(q-1)S}{q}\}$, though not necessarily in this order. For example, the mean is then calculated as

$$\frac{1}{q} \sum_{n=0}^{q-1} (M + \gamma + \frac{nS}{q})$$

= $(M + \gamma) + \frac{S}{q^2} \sum_{n=1}^{q-1} n$
= $(M + \gamma) + \frac{(q-1)S}{2q},$

where the last equality uses a result for the sum for the first q-1 natural numbers as found in Section 2.1.4. The variance is calculated in a similar fashion.

Note that in Lemma 1 that as $q \to \infty, \gamma \to 0$, the mean $\to M + \frac{S}{2}$ and the variance $\to \frac{S^2}{12}$ as would be expected for a uniformly distributed quasi-periodic mapping on [M, M + S).

3.2.5 Circle rotation map and DPLL

This subsection continues our review of the existing literature on DPLLs and circle rotation maps in the absence of additive input noise. It was noted in [4] that the output of the noise-free first and second-order DPLLs obeyed a circle map once it reached its steady-state limit cycle. Here we look at the basic DPLL equations and show that this is indeed the case.

3.2.5.1 First-order loop

The following straightforward result was assumed without proof in [4]. For completeness, we provide the details here.

Lemma 2. For a sufficiently small input frequency and loop gain, the steady-state first-order DPLL phase error obeys a circle rotation map.

Proof. Firstly we make the following assumptions regarding the input frequency and loop gain:

$$\operatorname{Int}[\mu] \le 2^b K - 1, \tag{3.91}$$

$$K < \frac{1}{2\pi}.\tag{3.92}$$

Equation (3.91) is a hard limit on the input signal's frequency and reflects the fact that, in general, a first-order PLL can only lock to a finite range of frequencies. Equation (3.92) is a sufficient but not a necessary condition for the DPLL to remain in steady-state and is used here to simplify the analysis. In practice, all DPLLs we consider in our numerical examples satisfy this constraint. Suppose that $\Phi(t)$ is such that $\text{Int}[2^bK\sin\Phi(t)] =$ $\text{Int}[\mu]$. This is possible by (3.91). Then, from (3.88), we have

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \operatorname{Frac}[\mu], \qquad (3.93)$$

where $\operatorname{Frac}[x]$ is the fractional part of x. Because $\Phi(t+1) > \Phi(t)$, we

will now have the following:

$$\operatorname{Int}[2^{b}K\sin\Phi(t+1)] = \begin{cases} \operatorname{Int}[\mu] & \text{if } \Phi(t+1) < \arcsin\frac{\operatorname{Int}[\mu]+1}{2^{b}K} \\ \operatorname{Int}[\mu]+1 & \text{if } \Phi(t+1) \ge \arcsin\frac{\operatorname{Int}[\mu]+1}{2^{b}K} \end{cases}$$

This is, the quantizer value for $\Phi(t+1)$ either remains unchanged from the previous value, or increases by 1. That only two states are possible is guaranteed by (3.92), which ensures that $\text{Int}[2^bK\sin\Phi]$ never changes by more than 1 for changes of Φ less than $\frac{2\pi}{2^b}$ as in equation (3.93). In the case where the quantizer increments at $\Phi(t+1)$, we then have, from (3.88)

$$\Phi(t+2) = \Phi(t+1) + \frac{2\pi}{2^b} \operatorname{Frac}[\mu] - \frac{2\pi}{2^b}.$$
(3.94)

In general, equations (3.93) and (3.94) will hold for some t, as, in steadystate, Φ cannot continue to increase indefinitely as in (3.93) alone. That is to say, in this steady-state regime, Φ obeys the mapping

$$\Phi(t+1) = \begin{cases} \Phi(t) + \frac{2\pi}{2^{b}} \operatorname{Frac}[\mu] & \text{for } \Phi(t) < \arcsin\frac{\operatorname{Int}[\mu] + 1}{2^{b}K} \\ \Phi(t) + \frac{2\pi}{2^{b}} \operatorname{Frac}[\mu] - \frac{2\pi}{2^{b}} & \text{for } \Phi(t) \ge \arcsin\frac{\operatorname{Int}[\mu] + 1}{2^{b}K}, \\ & \text{for } t = 0, 1, 2, \dots \end{cases}$$

This is a circle rotation map with $M = \arcsin \frac{\operatorname{Int}[\mu]+1}{2^b K} + \frac{2\pi}{2^b} (\operatorname{Frac}[\mu]-1),$ $S = \frac{2\pi}{2^b}, \text{ and } \alpha = \frac{2\pi}{2^b} \operatorname{Frac}[\mu], \text{ so that the rotation number } r = \frac{\alpha}{S} = \operatorname{Frac}[\mu].$

It should be noted from Lemma 2 above that the nature of the motion (periodic, quasi-periodic) depends only on $r = \frac{\alpha}{S} = \text{Frac}[\mu] = \text{Frac}[2^b\nu]$. Therefore, if the level of quantization, b, is changed and the input frequency altered to compensate such that μ remains constant, only the scale (S) and base (M) of the map changes. Lemma 1 may be used to calculate the variance of the map when μ is known, or to obtain bounds on the variance when it is not.

It should also be noted here that the fact that the DPLL obeys a circle map with size $S = 2\pi/2^b$ agrees with Gardner's observation (7) as noted in Section 3.2.3, that the peak-to-peak excursion of the phase

error is approximately $2\pi D/2^b$.

3.2.5.2 Second-order loop

We return again in this section to the equations (3.89) for the second order loop:

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \left(\mu - \operatorname{Int}[2^b \left(K_1 \sin \Phi(t) + u(t) \right)] \right) \mod 2\pi,$$

$$u(t+1) = u(t) + K_1 K_2 \sin \Phi(t+1).$$

Following the approach of Teplinksy et al. [4], we first make some assumptions.

Assumption 1. $b \ge 8$

This simply says that the number of bits in the NCO quantization is not too small. Typical hardware implementations or software algorithms work with 32- or 64-bit values.

Assumption 2. $K_2 < \frac{2}{q+1}$

This is a restriction on the integrator gain, implying that it must be smaller the 'less rational' μ is.

Assumption 3. $K_1 < \frac{1}{2\pi [\max\{\mu, 1-\mu\} + K_2]}$

Dependent on the input frequency μ , the gain of the proportional path is not too large. At a minimum, this assumption requires $K_1 < 1/\pi$.

With these assumptions in place, from equations (3.89) we see that u(n) continues to increase as long as $\sin \Phi(t)$ is positive and decreases when it is negative, and the rate of change of $\Phi(t)$ alters when the trajectory of the map (in Φ -u state space) crosses one of the 'switching curves' given by

$$S_i : u = 2^{-b}i - K_1 \sin \Phi, \qquad (3.95)$$

for $i \in \mathbb{Z}$.

The rate of change of $\Phi(t)$ will change sign will change sign when the trajectory crosses the curve with index $i = \text{Int}[\mu] + 1$. Throughout this thesis, we assume that μ is not an integer as this case requires special treatment as noted in [4]. This is an unlikely scenario in practice as it would require an input signal that was already frequency-locked exactly to the DPLL reference clock. For the cases of practical interest that we consider here we generally have $\text{Int}[\mu] = 0$ so this change of sign occurs across S_1 : to the left of this curve $\Phi(t)$ increases, and to its right it decreases. The S_i are examples of critical curves, as described in many books on nonlinear dynamics, e.g., [29].

Assuming a trajectory that begins at the origin, we see from equations (3.89) that $\Phi(t)$ will continue to increase by an amount $(2\pi/2^b)$ Frac $[\mu]$ until it crosses the switching curve S_1 . Meanwhile, u(t) will increase by an amount $\sin \Phi(t+1) \approx \Phi(t+1)$ for small Φ . This initial transient is shown in Figure 3.22(c). Once the trajectory crosses S_1 , Φ decreases by an amount $(2\pi/2^b)(1-\text{Frac}[\mu])$ and u continues to increase by approximately Φ while Φ remains positive.

The overall effect of the motions in Φ and u is, from the system's initial state, to carry the trajectory in a jagged path towards a 'trapping region' centred around $\Phi = 0$ and a non-zero u. This was shown previously in Figure 3.22(c) and is reproduced in more detail in Figure 3.23(a). The trapping region is marked "T" in Figure 3.23(a). Once the trajectory enters the trapping region, it never emerges, i.e., the region is invariant under the map.

It is easy to show that the trajectory is always bounded to the left and right by curves, also shown in Figure 3.23(a):

$$L_1 : u = 2^{-b} - K_1 \sin\left[\Phi + \frac{2\pi}{2^b}(1-\mu)\right] + K_1 K_2 \sin\Phi,$$

$$R_1 : u = 2^{-b} - K_1 \sin\left[\Phi + \frac{2\pi}{2^b}(-\mu)\right] + K_1 K_2 \sin\Phi.$$

The trapping region is reached when the switching curve, S_1 is close

to the axis $\Phi \equiv 0$. Then, in addition to the trajectory being bounded by the curves L_1 and R_1 , it is also contained within the lines

$$u = \frac{1}{2^{b}} - K_{1} \sin \frac{2\pi}{2^{b}} (1 - \mu),$$

$$u = \frac{1}{2^{b}} + K_{1} \sin \frac{2\pi\mu}{2^{b}},$$

$$\Phi = -\frac{2\pi}{2^{b}},$$

$$\Phi = \frac{2\pi}{2^{b}}.$$

It should be noted here again that, near $\Phi = 0$, the width of the $L_1 - R_1$ corridor is approximately $2\pi/2^b$, which agrees with Gardner's earlier observation (7), noted in Section 3.2.3, on the size of the peak-to-peak excursion of the phase error.

Inside the trapping region the trajectory tends towards a motion that is contained entirely within a smaller invariant region. Within the trapping region the exact nature of the motion depends very much on the system parameters. The integrator output, u increments or decrements depending on what side of the switching curve the system is on. Φ will obey a circle rotation map, similar to the first-order case; apart from the location of the switching curve, the motion in Φ is independent of u. Φ will increment by an amount $(2\pi/2^b)$ Frac $[\mu]$ to the left of the curve and decrement by $(2\pi/2^b)(1 - \operatorname{Frac}[\mu])$ to its right. For cases where the rotation number $\operatorname{Frac}[\mu]$ is rational, and equal to p/q in its lowest form, this circle map is periodic. After q iterations, the map will repeat. For systems that start at the origin, the circle map will be centred at $\Phi = 0$, i.e., one of the points on the map will be at $\Phi = 0$ and the (q-1) points will be equally distributed each to the left and right of the axis at intervals of $(2\pi/2^b)(1/q)$. From the equation (3.89) for u, and for small Φ we have $\sin \Phi \approx \Phi$, so that in this case u also obeys a circle rotation map. The maps for Φ and u are periodic, both with period q. Thus the overall motion is periodic with period q as shown in Figure 3.23(b) for the case where q = 25. Here, the invariant region is simply this period-25 limit cycle.



Figure 3.23: Plot of phase error for frequency-quantized second-order DPLL with b = 8, $\mu = 0.56$, $K_1 = 2^{-6}$, $K_2 = 2^{-5}$, and $\{\Phi(0), u(0)\} = \{0, 0\}$, showing on phase plane (a) path to steady-state region, (b) steady-state phase error and integrator output.

Nearby trajectories do not however settle to similar periodic orbits. For example, Figure 3.24(a) shows the invariant region for the same map when the trajectory starts from $\Phi = 10^{-4}$, u = 0 rather than the origin. While Φ in general follows the same period-25 cycle, the same periodic motion in u has been lost and instead varies across a much wider range of values than before. Compared to the periodic map, an additional value of Φ is added to the far left of the plot. This corresponds to a position on the map that is reached after a point with $\Phi = (2\pi/2^b)(-2/q)$ gets pushed to the positive side of the switching curve due to the rotation in u.

In general, where the circle map in Φ is *q*-periodic, the *q*th iterate of the map is a circle rotation map in *u*. The rotation number of this map depends on the initial conditions. For example, for the period-25 map that started at the origin, the rotation number was 0; the motions of Φ and *u* repeated together so that the 25th iterate of the map gave a fixed point in *u*. For the map in Figure 3.24(a) that starts at $\{\Phi, u\} =$ $\{10^{-4}, 0\}$, the rotation number is irrational. Therefore, while the motion in Φ is periodic, the motion in *u* is quasi-periodic. Thus the combined trajectory is quasi-periodic also. For the map shown in Figure 3.24(b), $\mu = 0.4$ so that the motion in Φ is a period-5 circle map. However, the



Figure 3.24: Similar plots to Figure 3.23(b) except with (a) $\mu = 0.56$, $\{\Phi(0), u(0)\} = \{10^{-4}, 0\}$, (b) $\mu = 0.4$, $\{\Phi(0), u(0)\} = \{2\pi/(2^b.10), 0.002\}$.

overall motion is 10-periodic because the 5th iterate of the map is a circle map in u with rotation number 1/2. That is, when Φ has completed two full 5-cycles, u has returned to its original value.

It was shown in [4] that all limit cycles for reasonable system parameters and initial conditions are contained within an absorbing set D within the trapping region. This applies to all rational input frequencies (i.e., $Int[\mu] = p/q$) whether the overall motion was periodic or quasi-periodic. The set D is invariant under the 2-dimensional map. The invariant set is shown in Figure 3.25 and is a rather complicated construction, derived from first principles in [4]. It consists of q + 1 narrow strips within the trapping region that contain all trajectories. The q + 1 strips correspond to the values in the rotation map for Φ , which is generally q-periodic with an additional point reached for a rotation in u as in Figure 3.24(a).

Figure 3.25 shows the invariant region for $\mu = 2/5$ and so in this case consists of 6 strips. The steady-state trajectories for three maps all with $\mu = 2/5$ are overlaid. The first, at the bottom of the plot, is for the trajectory that starts at the origin. This is the simplest trajectory for $\mu = 0.4$ where both Φ and u repeat together with the same period-5 motion. This is similar to the 25-periodic limit cycle shown earlier in Figure 3.23(b). The middle trajectory is the same as the one shown in Figure 3.24(b), i.e., has { $\Phi(0), u(0)$ } = { $2\pi/(2^b.10), 0.002$ }. Here, again



Figure 3.25: Plot showing invariant region from [4] for b = 8, $\mu = 0.4$, $K_1 = 2^{-6}$, $K_2 = 2^{-5}$. Overlaid are plots for (from bottom) $\{\Phi(0), u(0)\} =$ (a) $\{0, 0\}$, (b) $\{2\pi/(2^b.10), 0.002\}$, (c) $\{10^{-3}, 0\}$.

the motion in Φ is 5-periodic and the 5th iterate of the map is a rotation in u with rotation number 1/2. Thus, the overall motion is 10-periodic. The uppermost trajectory fills out a range of values for u for each value taken on by Φ and results from the system starting from $\{\Phi, u\} = \{10^{-3}, 0\}$ so that the 5th iterate of the map results in an irrational rotation in u. This is similar to the quasi-periodic map shown earlier for q = 25 in Figure 3.24(a). In all three cases in the plot here, the trajectories are contained within the invariant region.

Teplinsky et al. [5] also looked at the case where the input signal frequency, μ , is not a low-denominator rational; that is, μ is either a large-denominator rational (so that Assumption 2 above does not hold) or an irrational number. In both cases, there is still an absorbing set, D, though it has a more complex structure than what is shown for the low-denominator case in Figure 3.25. For the large-denominator rational frequency, the dynamics for Φ are again governed by a periodic circle map. However, the motion in u is no longer a circle map, but is instead a piecewise discontinuous circle rotation. The combination of the two leads to behaviour in the phase-plane that is not yet fully understood. For irrational input frequencies, Φ is governed by a quasi-periodic circle map while the motion in u is more complex again.

3.3 Summary

In this chapter we have introduced the phase-locked loop, starting in Section 3.1 with a summary of the existing literature on the classical analogue system. We showed how the basic equation of operation can be derived, and then linearized in the case where the phase-error remains small. Using the linear approximation, the steady-state behaviour of the PLL can easily be found in the cases of both first-order and second-order loops, where the input signal has either a constant or linearly-varying frequency. Analysis of the full non-linear system is significantly more complex; the phase-plane approach used by Viterbi [1] was recapped in Section 3.1.3 and we noted both the similarities and differences between the linear and non-linear systems.

In Section 3.1.4 we saw how, under some mild conditions, noise added to the input signal is equivalent to noise of the same variance added to the output of the phase detector. In the case of a linear PLL, the entire loop can be viewed as a simple linear filter, and the output response of the loop to input noise can easily found using the standard frequency-domain analysis for linear systems. For non-linear loops, once again the situation is not so straightforward; the tools of stochastic calculus are immediately employed to obain a Fokker-Planck equation for the steady-state PDF of the phase error. Analytical results can only be obtained in the simplest of scenarios, e.g., the first-order loop with constant-frequency input, though some useful approximate results may also be found for the second-order system.

Building on this knowledge of the analogue loop, the digital PLL was introduced in Section 3.2. Once again, the basic equations of operation for the first- and second-order loops were derived and similarities to the equivalent equations for the analogue case were illustrated. In Section 3.2.3 we introduced the concept of quantization, as well as the study model used by Gardner [3] and Teplinsky et al. [4] where this quantization is applied to the NCO input only. We summarized the previous findings on the steady-state limit cycles of the first- and second-order loops with NCO quantization. Finally, in Section 3.2.5 we showed how the steadystate behaviour of both loops obeys a map on the circle. For the firstorder system, the phase error follows a circle rotation map, while for the second-order loop, both the phase-error and integrator output obey circle maps with rotation numbers that depend on the initial conditions. For different combinations of rotation numbers, various types of limit cycle can be observed, as seen in [4], [5].

Our new material begins in the next chapter, where we consider the first-order digital loop with noise added to the input signal. Just as in the analogue case, we will see that this external noise is equivalent to the same noise added after the phase detector. Using simulations, we will show how digital loops with various levels of quantization respond to additive noise, recreating some of the earlier work in [3]. Similarities to the analogue case will be apparent, but also some differences where the additive noise and quantization interact in a manner that has not hitherto been understood. To gain an insight, we will formulate the DPLL as a unique form of noisy circle rotation map. The map can be treated as a Markov chain, and its Chapman-Kolmogorov equations yield a number of new results. In later chapters we will also look at the first-order loop with delay, as well as the second-order DPLL, and see how each system is affected by external noise using a number of analytical approaches as well as numerical simulations.

Chapter 4

Results for the first-order DPLL without loop delay

This section marks the beginning of our new material and results. In this chapter we present our results for the frequency-quantized, nonlinear first-order DPLL with additive input noise. The following chapters continue on to the case where we have a delay D > 1 in the first-order loop, and finally to the second-order DPLL.

Building on the existing literature, summarized in the preceding chapters, here we use the assumption that, under reasonable conditions, the steady-state behaviour of the non-linear, frequency-quantized DPLL of Section 3.2.3 is well approximated by the circle rotation map as shown in Section 3.2.4. In Section 4.1 it is shown how the equation for the first-order DPLL with additive noise can be written as a modified circle map with probabilistic jumps at each time interval. This unique way of writing the noisy DPLL equations forms the basis for our subsequent analysis from Section 4.3 onwards.

In Section 4.2 the main numerical results for the DPLL are shown, agreeing with earlier results from [3]. In addition, simulation results from the corresponding circle maps are shown, and it is seen that they agree well with the DPLL for a large range of additive noise levels.

Our analysis of the noisy circle map equations starts in Section 4.3. First we show that the circle map is a Markov chain and that its PDF



Figure 4.1: Block diagram of a first-order DPLL with additive noise.

obeys a Chapman-Kolmogorov equation. We proceed to derive several properties of the steady-state PDF. We then continue in Sections 4.4.1 and 4.4.2 with some approximate, asymptotic solutions for the steadystate using results from the previous section. We see that these solutions agree with the earlier numerical results within the valid parameter ranges for each approximation.

4.1 Circle map and the first-order DPLL with additive input noise

We are interested in the behaviour of the DPLL when noise is added to the input signal, since this is what is encountered by designers, and, as mentioned in [3], there is no clear understanding of this problem in the existing literature. The situation is shown in Figure 4.1. As described in [3], it is equivalent to model the additive input noise noise as being added after the phase detector in Figure 3.19. This was also the approach we used for the analogue PLL in Section 3.1.4 and is a standard model for PLL analysis, though it is strictly accurate only when the noise is Gaussian with the bulk of its power at frequencies greater than the loop bandwidth. We use this approximation to simplify both the analysis and simulation of the problem. Additionally, the input noise used in the simulation results presented here is uncorrelated and from a uniform distribution. We use this distribution as it allows us to derive some particular results in later sections, such as Theorem 7. However, results obtained using the more usual Gaussian-distributed input noise were virtually indistinguishable from those shown below.

Using the approximation of adding the noise after the phase detector, the noise term enters equation (3.88) as an additional delayed input to the quantizer Int[...]. Therefore, the numerical results in the following section involve simulating the difference equation (4.2) below, allowing $\Phi(t)$ reach steady state, when at each time step the noise sample N(t) is drawn randomly from a uniform density with bounds determined by the input noise variance.

An alternative way of writing the circle rotation map in equation (3.90) is as follows:

$$x(t+1) = x(t) + \alpha + Q(x(t))S, \text{ for } t = 0, 1, 2, \dots, \text{ where}$$
$$Q(x) = \begin{cases} -1 & \text{if } x \ge M + S - \alpha, \\ 0 & \text{if } x < M + S + \alpha. \end{cases}$$
(4.1)

Here, Q(x) represents the quantizer and provides the corresponding function to the modulo operation in (3.90).

In the case where noise is added to the DPLL input signal, equation (3.88) governing the behaviour of the output phase error is modified as follows:

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \{ \mu - \operatorname{Int}[2^b K(\sin \Phi(t) + N(t))] \}.$$
(4.2)

Here N(t) is the noise sample added after the phase-detector at time-step t, and, again, is statistically equivalent to adding it to the input signal. A similar equation exists for the analogue case, equation (3.45) in Section 3.1.6, for example.

Remark 1. For large noise, i.e., $N(t) \gg \sin \Phi(t)$ and $N(t) \gg \frac{\nu}{K}$, (4.2) can be approximated by

$$\Phi(t+1) = \Phi(t) - 2\pi K N(t).$$
(4.3)

In this case $\phi(t)$ obeys a random walk and thus is non-stationary so no steady-state variance exists. This thesis considers only cases where steady-state solutions occur, these being the cases of engineering interest.

The addition of the random noise sample in (4.2) will sometimes cause the quantizer (Int[...]) to output a value other than what it would have done in the noise-free case. The DPLL can therefore be modelled by a noisy circle rotation map as follows:

$$x(t+1) = x(t) + \alpha + Q(x(t))S$$
, for $t = 0, 1, 2, \dots$
 $x(0) = x_0$,

where

$$Q(x) = \begin{cases} -1 & \text{with probability } F_N(x), \\ 0 & \text{with probability } 1 - F_N(x). \end{cases}$$
(4.4)

In the above, the function $F_N(x)$ represents the cumulative distribution



(a) Function $F_N(x)$ in case where additive noise is uniformly distributed in [-L, L].

(b) Initial distribution p(x, 0) in case of quasi-periodic circle rotation map.

Figure 4.2

function (CDF) of the noise N(t) as follows

$$F_N(x) = P_N(x-T)$$
, where $T = (M+S-\alpha)$, and $P_N(x) = \int_{-\infty}^x p_N(y) dy$,
(4.5)

where $p_N(y)$ is the PDF of N(t).

The function $F_N(x)$ is shown in Figure 4.2(a) for the case where N(t)is uniformly distributed in [-L, L]. The noise variance, σ_N^2 , is $\frac{L^2}{3}$ in this case. Note that in the range [T - L, T) the probability of getting a quantizer value of -1 is in the range [0, 0.5), whereas in the noise-free case it would have been 0: the quantizer would output zero always. Similarly, above x = T, the probability of a quantizer value of -1 increases from 0.5 to 1. At x = T, the probability of a jump in either direction is 0.5. On the other hand, in the case where the additive noise is zero the PDF of N(t) is the delta spike $\delta(0)$, and so $F_N(x)$ is a step function at x = T. In this situation, definition (4.4) for the noisy circle map reduces to the original (4.1).

We are interested in finding the statistics of x(t), given those of N(t), and its variance in particular. Therefore it makes sense to define the time-dependent PDF of x(t) as p(x, t), where $x \in \mathbb{R}$, t = 0, 1, 2, ...

Equation (4.4) above is similar to that studied recently in the literature for a first-order digital bang-bang PLL with zero loop delay and non-zero loop detuning, in [30], [31]. Combining equations (1)-(3) of [31], and setting D = 0, yields an equation similar to (4.4) in this paper, though x(t) in the latter corresponds to the auxiliary variable Δt_k^* of [31], rather than actual output timing jitter, Δt_k . Our quantizer here also differs slightly, taking on values 0 and 1, rather than the ±1 of the sgn[.] function of [31]. This latter paper, however, focusses primarily on the case of non-zero delay. The same system with zero delay was considered earlier by Da Dalt in [30]. In this thesis our main results involve approximations of the output jitter variance, rather the phase-detector gain of [30].

The behaviour of the noisy first-order DPLL (4.2) may also be visualized on the $\Phi(t+1) - \Phi(t)$ phase plane, similar to the noise-free case that was plotted earlier in Figure 3.21. This is shown in Figure 4.3. As before, steady-state is reached when the graph of (4.2) without the noise term lies on each side of the identity line. However, with noise noise now added, there is a region of uncertainty at the transition point where additional points now lie above and below the identity line beyond the bounds of the original map. This is indicated by the shaded region in Figure 4.3. The behaviour within this region is described by the noisy circle map (4.4) and will be investigated further in Section 4.3.



Figure 4.3: Plot on the $\Phi(t) - \Phi(t+1)$ plane of frequency-quantized first-order DPLL with b = 8, $\mu = 0.56$, $K = 2^{-6}$, and $\Phi(0) = 0$, with additive noise variance $\sigma_N^2 = 10^{-5}$.

4.2 Numerical results for circle map and the first-order DPLL with noise

The main numerical results of the first-order DPLL are shown in Figure 4.4(a). The response of the variance of the output jitter, σ_O^2 , to input noise variance, σ_N^2 , is displayed for various DPLL quantization levels. As mentioned in the previous sub-section, the noise is uncorrelated and drawn from a uniform-density distribution with variance σ_N^2 ; the variance response obtained using this uniform-density input is virtually indistinguishable from that using the more standard Gaussian.



(a) DPLL variance response.

(b) DPLL with circle maps.

Figure 4.4: Variance of output jitter, σ_O^2 , as a function of input noise variance, σ_N^2 , for (a) DPLL with $\mu = 0.56$, $K = 2^{-6}$, and b = 14 (solid), b = 12 (large dash), b = 10 (med dash), b = 8 (small dash), and for (b) DPLLs with corresponding circle rotation maps.

Three distinct regions can clearly be identified. The first is on the left of the graphs where the additive noise has little or no effect, and the output variance can easily be calculated or bounded using Lemma 1. Unsurprisingly, the DPLLs with lower quantization values have a higher noise-free variance, but are also less sensitive to the input noise, i.e., it requires a higher level of input noise to affect the output. The PDF of the output jitter in this region is dependent on whether the corresponding circle rotation map is periodic or quasi-periodic, which in turn is determined by $Frac[\mu]$.

The second, intermediate, region is where the additive noise starts to have an effect, but there is also still a significant dependency on the quantization level, demonstrating that the additive noise and inherent jitter are interacting in some way. The region encompasses several orders of magnitude on both the input and output variance axes, and there are points where, for example, the output variance has risen to over an order of magnitude above the noise-free level, but quantization still has a significant effect. Given this, an understanding of the behaviour in this region should be of interest to PLL designers. In addition it was noted that the PDF of the output jitter becomes closer to Gaussian as this intermediate region is traversed, as can be seen, for example, in the histogram plots of Figure 4.5. The final region is where the output variance is independent of quantization, and the graph of the output variance is a line of slope 1 on log-log axes. This, as expected, and as noted in [3], is exactly the response of a linear, analogue PLL. It can also be seen from an examination of equation (3.44): since the variance of the output jitter is proportional to that of the input noise, this gives a line of slope 1 on a log-log plot. The PDF of the output in this region is close to Gaussian.

It should be noted that the parameter μ (and ν) has little effect on these plots. The exact details of the transition from the noise-free to the intermediate region does depend on μ , as the DPLL obeys a circle map in the absence of noise, the dynamics of which are determined entirely by μ . However the overall behaviour for increasing σ_N is essentially the same for all input frequencies.

While equation (4.2) describes the noisy DPLL, the noisy circle map is governed by equation (4.4); the correspondence in parameters between the two equations is given by the result of Lemma 2. In Figure 4.4(b) the responses of the corresponding circle rotation maps to the same input noise levels are shown, the noise being added to the circle such that it corresponds exactly to the DPLL additive noise. In general the circle map plots track the DPLL plots until the noise levels become large, whereupon they start deviating from the DPLL plots, each circle map plot continuing in a straight line with slope of approximately $\frac{1}{2}$. That the line has slope $\frac{1}{2}$ indicates that the output variance of the circle map is proportional to the square root of the input noise variance, when the noise is added according to equation (4.4). This will be investigated later in Section 4.4.1.

The measured output variance of the circle maps fluctuates and becomes more unstable as the additive noise level is increased and the system's output becomes non-stationary. This is of little importance as this behaviour is evident only in the region where each DPLL is responding like an analogue PLL. It should be clear from the plot that the regime where the circle map variance begins increasing with slope $\frac{1}{2}$ corresponds to the intermediate region in the DPLL response, so the behaviour of the





Figure 4.5: Histograms of steady-state circle map output for M = 0, S = 1, $\alpha = 0.1387$, and input noise variance (a) 0.0001, (b) 0.001, (c) 0.01, (d) 0.1, (e) 1.0, (f) 2.0.

We also examine in some detail in Figure 4.5 the PDF of the circle map output when subjected to the additive DPLL noise. When no noise is added, and in the general case where $r = \frac{\alpha}{S}$ is irrational, the output has a uniform distribution on [M, M + S]. As noise is added, and its variance increased, the region within [M, M + S] of uniform probability shrinks as the range of non-zero probability increases and 'tails' form, as shown in Figure 4.5. As the input noise level is increased further, the region of constant probability vanishes entirely, and the PDF appears Gaussian. Indeed, the measured kurtosis becomes close to 3, and the Gaussian approximates the measured PDF very closely as the noise variance increases. The PDF of the circle map output is analyzed in more detail later in Section 4.4.

Another feature noticeable from the sequence in Figure 4.5 is that the mean of the output, μ_O , decreases as the input noise level is increased. This is better illustrated in Figure 4.6. In general, as the input noise-level is increased, the mean increases or decreases according to the noiseindependent circle map parameter sgn(r-1). This is investigated further in Theorem 7 on page page 144.



Figure 4.6: Mean of circle map output, μ_O , for $M = 0, S = 1, \alpha = 0.1387$ plotted against input noise variance, σ_N^2 .

4.3 Analysis of circle map with additive noise

Starting with equation (4.4) we may derive an equation for the timedependent PDF p(x,t) of x(t). Firstly, note that we can write

$$x(0) = x_0,$$

and

$$x(t+1) = \begin{cases} h_1[x(t)] & \text{with probability } F_N(x) \\ h_2[x(t)] & \text{with probability } 1 - F_N(x). \end{cases}$$
(4.6)

Here we have

$$h_1(z) = z + \alpha - S,$$
$$h_2(z) = z + \alpha.$$

Clearly, x(t) defined this way describes a Markov chain as defined in Section 2.2.5, since the next state, x(t + 1), depends only on x(t). In particular, we can use the Chapman-Kolmogorov equation (2.68) for this particular process, x(t), to obtain the next result.

Theorem 3. For the noisy circle rotation map as defined in (4.4), the PDF, p(x, t), of x(t) obeys the following equation

$$p(x,t) = p(x-\alpha,t-1)[1-F_N(x-\alpha)] + p(x-\alpha+S,t-1)F_N(x-\alpha+S).$$
(4.7)

Proof. Since x(t) is a Markov chain we can write its Chapman-Kolmogorov equation as follows:

$$p(x,t) = \int_{-\infty}^{\infty} q(x|z,t-1)p(z,t-1)dz,$$
(4.8)

where q(x|z, t-1) is the transition PDF of p(x, t) conditioned on x(t-1) = z. From (4.6) we can write

$$q(x|z,t-1) = F_N(z)\delta(x-h_1(z)) + (1-F_N(z))\delta(x-h_2(z)), \quad (4.9)$$

and substituting for h_1 and h_2 we obtain

$$q(x|z,t-1) = F_N(z)\delta(x-z-\alpha+S) + (1-F_N(z))\delta(x-z-\alpha).$$
(4.10)

Finally, inserting (4.10) into (4.8), and integrating, the result follows. \Box

This result may be used to simulate the behaviour of the PDF of x by evolving it over time from an initial condition p(x, 0). For example, in the case of an irrational $\operatorname{Frac}[\mu]$, the noise-free mapping is quasi-periodic, and the ensemble will initially be uniformly distributed in [M, M+S) as shown in Figure 4.2(b).

Figure 4.7 shows the evolution of the initial PDF p(x, 0) to a steadystate PDF over 500 time steps. In this case, the circle map parameters and noise level used are the same as in Figure 4.5(d), and the noise CDF $F_N(x)$ is as shown in Figure 4.2(a). For all such cases, the circle map PDF quickly reaches steady-state and for larger levels of input noise the PDF is approximately Gaussian, as measured by its kurtosis.



Figure 4.7: Initial PDF, and PDF after 500 time steps, of output for circle map with M = 0, S = 1, $\alpha = 0.1387$ and input noise variance, $\sigma_N^2 = 0.1$.

If the PDF in (4.7) reaches a steady-state (as it is observed to do in numerical simulations) then this state $p_{\infty}(x)$ is described by the solution of the non-local equation

$$p_{\infty}(x) = p_{\infty}(x-\alpha)[1 - F_N(x-\alpha)] + p_{\infty}(x+S-\alpha)F_N(x+S-\alpha) \quad (4.11)$$

or, equivalently

$$p_{\infty}(x)F_N(x) = p_{\infty}(x+\alpha-S) - p_{\infty}(x-S)[1-F_N(x-S)]. \quad (4.12)$$

Since $p_{\infty}(x)$ is a PDF, it obeys $p_{\infty}(x) \geq 0 \quad \forall x$, and the normalization condition $\int_{-\infty}^{\infty} p_{\infty}(x) dx = 1$, which requires the limiting behaviour $\lim_{x \to \pm \infty} p_{\infty}(x) = 0$.

We next derive a symmetry result for the solution, $p_{\infty}(x)$, of (4.11). In the case of additive noise N(t) of mean 0, with PDF symmetric about y = 0, the noise CDF, $F_N(x)$, has the following symmetry property as shown in Figure 4.2(a)

$$F_N(T+x) + F_N(T-x) = 1.$$
(4.13)

Let us now make the following transformations:

$$\alpha \longmapsto (S - \alpha)$$
$$F_N(x) \longmapsto f(x + (S - 2\alpha))$$

These are the transformations that result from interchanging α and $(S - \alpha)$ in the physical circle map, the latter transformation following because $F_N(x)$, which was centred at $x = T = (M + S - \alpha)$ (see Figure 4.2(a)) is now centred at $(M + \alpha)$. With these changes equation (4.11) for the new solution q_{∞} becomes the symmetry equation:

$$q_{\infty}(x) = q_{\infty}(x - (S - \alpha))[1 - F_N(x - \alpha)] + q_{\infty}(x + \alpha)F_N(x + (S - \alpha)). \quad (4.14)$$

The next result relates the solution $q_{\infty}(x)$ of the symmetry equation to the original solution, $p_{\infty}(x)$ of equation (4.11).

Theorem 4. For $p_{\infty}(x)$ that is a solution of (4.11), $F_N(x)$ that satisfies (4.13), $q_{\infty}(x) = p_{\infty}(2M + S - x)$ is a solution of symmetry equation (4.14).

Proof. Firstly note that equation (4.13) for the noise symmetry can be written as

$$F_N(M + S - \alpha + x) + F_N(M + S - \alpha - x) = 1.$$
(4.15)

Using a simple transformation of variable in each case, this yields the equivalent pair of relations

$$F_N(x - \alpha) + F_N(2M + 2S - \alpha - x) = 1, \qquad (4.16)$$

$$F_N(x+S-\alpha) + F_N(2M+S-\alpha-x) = 1.$$
 (4.17)

Now setting $q_{\infty}(x) = p_{\infty}(2M + S - x)$, we have $p_{\infty}(x) = q_{\infty}(2M + S - x)$,

and, substituting in (4.11), we get

$$q_{\infty}(2M + S - x) = q_{\infty}(2M + S + \alpha - x)[1 - F_N(x - \alpha)] + q_{\infty}(2M + \alpha - x)F_N(x + S - \alpha).$$
(4.18)

Using relations (4.16) and (4.17), this is

$$q_{\infty}(2M + S - x) = q_{\infty}(2M + S + \alpha - x)F_{N}(2M + 2S - \alpha - x) + q_{\infty}(2M + \alpha - x)[1 - F_{N}(2M + S - \alpha - x)].$$
(4.19)

Changing variables reduces the equation to

$$q_{\infty}(x) = q_{\infty}(x+\alpha)F_N(x+(S-\alpha)) + q_{\infty}(x-(S-\alpha))[1-F_N(x-\alpha)].$$
(4.20)

This is precisely equation (4.14). Finally, we note that the boundary and normalization conditions satisfied by $p_{\infty}(x)$ of (4.11), $\lim_{x\to\pm\infty} p_{\infty}(x) = 0$ and $\int p_{\infty}(x) dx = 1$, are also satisfied by $Q(x) = p_{\infty}(2M + S - x)$.

This result means that we can restrict our attention in the remainder of this section to the case $\alpha \leq \frac{S}{2}$, since the symmetry result immediately gives us the solution for the corresponding case where $\alpha \geq \frac{S}{2}$. The next result gives a further property of solution p_{∞} for a particular noise CDF $F_N(x)$.

Theorem 5. For the steady-state distribution $p_{\infty}(x)$ given in (4.12), $F_N(x)$ that is identically 0 in $(-\infty, T - L]$ and identically 1 in $[T + L, \infty)$ (such as the uniform distribution in Figure 4.2(a)), and initial distribution P(x, 0) that is zero outside the range (M - L, M + S + L) (e.g., that given in Figure 4.2(b)), $p_{\infty}(x)$ is also identically zero outside the region (M - L, M + S + L).

Proof. For $x \leq T - L + \alpha - S = M - L$ we have $F_N(x - \alpha + S) = 0$ and $F_N(x - \alpha) = 0$ from Figure 4.2(a). Then (4.7) becomes $p(x,t) = p(x - \alpha, t - 1)$ and hence $p(x,t) = p(x - t\alpha, 0)$. Since $x \leq M - L$, $x - t\alpha \leq M - L$ also, so $p(x - t\alpha, 0) = 0$ from Figure 4.2(b). Hence p(x, t) = 0.

Similarly, for $x \ge T + L + \alpha = M + S + L$ we have $F_N(x - \alpha) = 1$ and $F_N(x - \alpha + S) = 1$, so (4.7) becomes $p(x, t) = p(x + S - \alpha, t - 1)$ and hence $p(x, t) = p(x + t(S - \alpha), 0)$. Since $x \ge M + S + L$ and $(S - \alpha) \ge 0$, $x + t(S - \alpha) \ge M + S + L$ also, so $p(x + t(S - \alpha), 0) = 0$ from Figure 4.2(b). Hence p(x, t) = 0.

Therefore $\forall t, p(x,t) = 0$ for $x \leq M - L$ and $x \geq M + S + L$. The same is true of $p_{\infty}(x)$ if this steady-state solution exists. Clearly, the only region where $p_{\infty}(x)$ can be non-zero is in (M - L, M + S + L), a region of size 2L + S containing $T = M + S - \alpha$.

The following result serves to verify that (4.12) holds for the noise-free circle map, i.e., that the uniform PDF is a solution of (4.12).

Lemma 6. For $F_N(x)$ as in Figure 4.2(a) in the noise-free case, i.e., L = 0, the uniform PDF on [M, M + S], as given in Figure 4.2(b), is a solution for the steady-state $p_{\infty}(x)$ in (4.12).

Proof. Firstly, note that the result of Theorem 5 can be derived directly from (4.12) by considering separately the cases $x \leq T - L$ and $x \geq T + L + S$. In the first case, both $F_N(x)$ and $F_N(x - S)$ are identically 0, so (4.12) gives

$$p_{\infty}(x+\alpha-S) = p_{\infty}(x-S). \tag{4.21}$$

Equivalently, this is, for $x \leq M - L$

$$p_{\infty}(x) = p_{\infty}(x - \alpha). \tag{4.22}$$

Since $p_{\infty}(x) \to 0$ as $x \to -\infty$, must have $p_{\infty}(x) = 0$ for $x \leq M - L$. Similarly, considering the case $x \geq T + L + S$ yields $p_{\infty}(x) = 0$ for $x \geq M + S + L$. Hence (4.12) is consistent with (4.7) in that it gives $p_{\infty}(x) = 0$ outside of the range (M - L, M + S + L).

For the noise-free case, L = 0 (the linear ramp of Figure 4.2(a) becomes a step function), and this range reduces to (M, M+S). Clearly, the
uniform distribution on [M, M+S] satisfies this, so all that remains to be shown is that a constant $p_{\infty}(x) = \frac{1}{S}$ is a solution of (4.12) on (M, M+S). This may be done by again considering two cases, $x \in (T, T + \alpha)$ and $x \in [T + \alpha, T + S)$. The first case gives $F_N(x) \equiv 1$, $F_N(x - S) \equiv 0$, and, from above, $p_{\infty}(x - S) \equiv 0$. Equation (4.12) then becomes

$$p_{\infty}(x) = p_{\infty}(x + \alpha - S). \tag{4.23}$$

So, for $x \in (M + S - \alpha, M + S)$ we obtain

$$p_{\infty}(x) = p_{\infty}(x - (S - \alpha)).$$
 (4.24)

Similarly, considering the range $[T + \alpha, T + S)$, gives for $x \in [M + \alpha, M + S)$,

$$p_{\infty}(x) = p_{\infty}(x - \alpha). \tag{4.25}$$

Equations (4.24) and (4.25) define two pairs of regions within (M, M + S) where $p_{\infty}(x)$ must repeat. Specifically these are $(M, M + \alpha)$ and $(M + S - \alpha, M + S)$, shown as the smaller dotted regions in Figure 4.8, and $(M, M + S - \alpha)$ and $(M + \alpha, M + S)$, the larger hatched regions in the same figure. Clearly the uniform PDF on (M, M + S) satisfies these requirements.



Figure 4.8: Solution domains for steady-state PDF in noise-free case, as discussed in Lemma 6.

Remark 2. In the case of a noise-free circle map with $r = \frac{\alpha}{S}$ rational, and equal to $\frac{p}{q}$ in its lowest form, the steady-state PDF is the train of

discrete delta functions as in Lemma 1:

$$p_{\infty}(x) = \frac{1}{q} \sum_{n=0}^{q-1} \delta\left(x - \left(M + \gamma + \frac{nS}{q}\right)\right).$$

By substituting $\alpha = rS$ it is clear that this PDF satisfies relations (4.24) and (4.25).

4.4 Approximate solutions for steady-state circle rotation map

We present in the following section various analytical approximations of the solution for the circle map PDF introduced in Section 4.3. In particular, our next results, Theorem 7 and Remark 3, provide the most useful approximations to the behaviour of the circle map output variance as shown in Figure 4.4(b).

4.4.1 Large-*L* approximation

Theorem 7. For the steady-state distribution $p_{\infty}(x)$ given in (4.12), $F_N(x)$ as in Figure 4.2(a), $p_{\infty}(x)$ is approximately Gaussian with a variance of $(r - r^2)SL + \frac{S^2}{8}$, where $r = \frac{\alpha}{S}$.

Proof. Using the change of variable y = x - z, (4.12) becomes

$$p_{\infty}(y+z)F_N(y+z) = p_{\infty}(y+z+\alpha-S) - p_{\infty}(y+z-S)[1-F_N(y+z-S)].$$
(4.26)

Using a Taylor expansion $p_{\infty}(y+z) = p_{\infty}(y) + zp'_{\infty}(y) + \frac{z^2}{2}p''_{\infty}(y) + \dots$, similarly for $p_{\infty}(y-z)$, $F_N(x+z)$, and $F_N(x-z)$, neglecting higher-order terms, and choosing $z = \frac{S}{2}$, we get

$$K_1 p_{\infty}''(y) + (SF_N(y) - \alpha) p_{\infty}'(y) + SF_N'(y) p_{\infty}(y) = 0, \qquad (4.27)$$

where $K_1 = \frac{S\alpha}{2} - \frac{\alpha^2}{2} + \frac{S^3}{16L}$. Equivalently, this is

$$K_1 p_{\infty}''(y) + \frac{d}{dy} [(SF_N(y) - \alpha)p_{\infty}(y)] = 0.$$
(4.28)

Next, if we integrate, and demanding $p_{\infty}(y)$ and $p'_{\infty}(y) \to 0$ as $y \to \pm \infty$, we obtain

$$K_1 p'_{\infty}(y) + (SF_N(y) - \alpha)p_{\infty}(y) = 0,$$
 (4.29)

which we can rearrange as

$$\frac{p'_{\infty}(y)}{p_{\infty}(y)} = \frac{(\alpha - SF_N(y))}{K_1}.$$
(4.30)

Next, if we substitute for linear $F_N(y) = \frac{1}{2}(\frac{y-T+L}{L})$ as in Figure 4.2(a), we get

$$\frac{p'_{\infty}(y)}{p_{\infty}(y)} = \frac{K_2 - Sy}{2LK_1},\tag{4.31}$$

where $K_2 = 2L\alpha + S(T-L)$. Integrating again, we get, for some constant \tilde{C}

$$p_{\infty}(y) = \tilde{C} \exp\left[\frac{2K_2y - Sy^2}{4LK_1}\right].$$
(4.32)

Finally, if we complete the square and rearrange, and let $C = \tilde{C} \exp\left[\frac{K_2^2}{4LK_1}\right]$, we obtain

$$p_{\infty}(y) = C \exp\left[\frac{-S}{4LK_1}\left(y - \frac{K_2}{S}\right)^2\right].$$
(4.33)

This is a Gaussian of mean $\mu_O = \frac{K_2}{S}$ and variance $\sigma_O^2 = \frac{2LK_1}{S}$. Substituting for K_1 in the expression for the variance, we obtain

$$\sigma_O^2 = \frac{2L}{S} \left(\frac{S\alpha}{2} - \frac{\alpha^2}{2} + \frac{S^3}{16L} \right).$$
(4.34)

Rearranging, we get the result

$$\sigma_O^2 = \left(\frac{\alpha}{S} - \frac{\alpha^2}{S^2}\right)SL + \frac{S^2}{8}.$$
(4.35)



Figure 4.9: Plots of DPLL output jitter variance with b = 14, b = 8, and corresponding circle maps from Figure 4.4 with noise-free limits $\frac{S^2}{12}$ from Lemma 1 and large-*L* asymptotes (*dashed*) found in Theorem 7.

Remark 3. The output variance can be written in terms of the input variance, σ_N , by substituting for L in (4.35), i.e.,

$$\sigma_O^2 = (r - r^2) S \left(3\sigma_N^2 \right)^{\frac{1}{2}} + \frac{S^2}{8}.$$
 (4.36)

As the variance of the additive noise becomes large, the output variance increases proportional to the square root of the input, i.e., $\sigma_O^2 \approx (r - r^2)S(3\sigma_N^2)^{\frac{1}{2}}$. This agrees with the observation from the numerical simulations that, on the log-log plot of Figure 4.4(b), the output variance was asymptotically linear in the input, with slope $\frac{1}{2}$. Also, in the noise-free limit, the output variance in (4.36) is $\frac{S^2}{8}$. This does not correspond to the actual variance of the noise-free circle map as found in Lemma 1, $\frac{S^2}{12}$, in the general case where r is irrational. This indicates that approximations were made in the workings of Theorem 7 that are not valid for small L; see Section 4.4.2 below for further discussion of a case where L is small. Figure 4.9 shows two of the earlier DPLL and circle map plots, along with the corresponding asymptote (4.36) and noise-free limit $\frac{S^2}{12}$, as found analytically. Note our asymptotic formula (4.36) accurately characterises the jitter in the intermediate region between the (well-understood) limits of quantization jitter and analogue PLL behaviour. As discussed in Section 4.2, this regime is of significant interest to PLL designers, and has previously only been examined via numerical simulation.

Remark 4. The output mean was found to be $\frac{K_2}{S}$ during the workings of Theorem 7. This is T + (2r - 1)L. This implies that the output mean changes linearly with L and increases or decreases depending on whether $r > \frac{1}{2}$ or $r < \frac{1}{2}$. This corresponds with the observation noted earlier in Section 4.2 and illustrated in Figure 4.6. Indeed, similar to the asymptotic variance above, the mean found here approximates the mean obtained from direct simulation of the circle map very well for large L. This is illustrated in Figure 4.10, where it is clear that the asymptotic mean is a good match for the direct simulation, even for small L. Also shown is the mean for the DPLL corresponding to the same circle map. The DPLL mean matches that of the circle map for low levels of input noise, but diverges for larger L, just as the variance diverged in Figure 4.9.

Remark 5. There are two approximations made in the proof of Theorem 7. The first is in the Taylor series truncation of the unknown function $p_{\infty}(y)$ in obtaining (4.27). [Note: the truncation of the expansion for $F_N(y)$ is exact in the case where it is linear]. The second is the assumption in (4.31) that $F_N(y)$ is a linear function with slope $\frac{1}{2L}$ for all y, whereas in fact this is true only in the range [T - L, T + L], as in Figure 4.2(a). This approximation may be accurate in the large-noise limit, but will fail in the case where L is small.

4.4.2 Small- α approximation

Working directly from (4.11), and again assuming $F_N(x)$ that is identically 0 in $(-\infty, T-L]$ and identically 1 in $[T+L, \infty)$ as in the uniform



Figure 4.10: Plots of output mean for circle map corresponding to DPLL with b = 10, and large-*L* asymptote for circle map mean *(dashed)* noted in Remark 4. Also shown is mean for DPLL which matches circle map for low noise but diverges for large σ_N^2 .

CDF of Figure 4.2(a), we can write piecewise equations for $p_{\infty}(x)$ for various combinations of the system parameters, S, α , and L. Specifically, these combinations are

Here, Case 5 above corresponds to the large L limit considered in Theorem 7. Case 3 is equivalent to Case 2, replacing α with $(S - \alpha)$ and using the result of Theorem 4. Case 2 represents the situation where α is small, and L is also small, but not vanishingly small. The full set of equations for Case 2 are:

Case 2: $L > \frac{\alpha}{2}, L < \frac{S-\alpha}{2}$

$$p_{\infty}(x) = \begin{cases} p_{\infty}(x+S-\alpha)F_N(x+S-\alpha) & M-L < x \le M-L+\alpha \\ p_{\infty}(x-\alpha) + p_{\infty}(x+\alpha_1)F_N(x+\alpha_1) & M-L+\alpha < x \le M+L \\ p_{\infty}(x-\alpha) + p_{\infty}(x+S-\alpha) & M+L < x \le M+L+\alpha \\ p_{\infty}(x-\alpha) & M+L+\alpha < x \le M+S-L \\ p_{\infty}(x-\alpha)[1-F_N(x-\alpha)] & M+S-L < x \le M+S+L \end{cases}$$

Here, for brevity, we set $\alpha_1 = (S - \alpha)$. For this case, and for all cases by Theorem 5, we have $p_{\infty}(x) \equiv 0$ for $x \leq M - L$ and for $x \geq M + S + L$.

We next consider a small- α expansion of the equations in Case 2. Note that this corresponds to the case for the original DPLL where the fractional part of the input frequency, $\operatorname{Frac}[\mu]$, is small. In particular, for $\alpha \ll 2L$ and $\alpha \ll (S - 2L)$ the regions of size α are neglected and we obtain:

$$p_{\infty}(x) = \begin{cases} p_{\infty}(x-\alpha) + p_{\infty}(x+\alpha_1)F_N(x+\alpha_1) & M-L < x \le M+L \\ p_{\infty}(x-\alpha) & M+L < x \le M+S-L \\ p_{\infty}(x-\alpha)[1-F_N(x-\alpha)] & M+S-L < x \le M+S+L \end{cases}$$

Re-writing, and ignoring an offset of size α in the endpoints of each region, the equation set becomes:

 $p_{\infty}(x+\alpha) =$

$$\begin{cases} p_{\infty}(x) + p_{\infty}(x+S)F_{N}(x+S) & M - L < x \le M + L \\ p_{\infty}(x) & M + L < x \le M + S - L \\ p_{\infty}(x)[1 - F_{N}(x)] & M + S - L < x \le M + S + L \end{cases}$$

Substituting for linear F_N as in Figure 4.2(a), and Taylor expanding to

first order for small α we can obtain the explicit piecewise solution:

$$p_{\infty}(x) = \begin{cases} p_3(x) & M - L < x \le M + L \\ p_1(x) & M + L < x \le M + S - L \\ p_2(x) & M + S - L < x \le M + S + L \end{cases}$$

where

$$p_3(x) = C_3 - C_2 \exp\left[\frac{-(S+x-\alpha)(2L-2M-S+x+\alpha)}{4L\alpha}\right], \quad (4.37)$$

$$p_1(x) = C_1, (4.38)$$

$$p_2(x) = C_2 \exp\left[\frac{2(M+S-L-\alpha)x-x^2}{4L\alpha}\right].$$
 (4.39)

We can determine two of the arbitrary constants in this solution set immediately by matching p_2 and p_3 to the constant region p_1 , i.e., by solving

$$\begin{cases} p_3(M+L) = p_1(M+L) \\ p_2(M+S-L) = p_1(M+S-L) \end{cases}$$

This gives

$$\begin{cases} C_2 = C_1 \exp\left[-\frac{(M+S-L)(M+S-L-2\alpha)}{4L\alpha}\right]\\ C_3 = C_1 \left(1 + \exp\left[-\frac{L}{\alpha} + \frac{\alpha}{4L}\right]\right) \end{cases}$$

At this point we may proceed to integrate (4.37)-(4.39) numerically and find C_1 by normalizing the PDF. Equations (4.37)-(4.39) together then give a complete piecewise solution for the PDF in this small- α case.

We can check the accuracy of the solution by numerically integrating (4.37)-(4.39) to find the normalized piecewise solution and comparing to histograms of the noisy circle map ouput. Figure 4.11 shows the piecewise PDF as compared to the histogram obtained from direct simulation of the same circle map as in Figure 4.5. Clearly, the piecewise approximation is a good match for the numerical result.

Using numerical integration, the piecewise solution can also be used to calculate the moments, from which the mean and variance may be found. Figure 4.12 shows plots of the approximated mean and variance of the



(a) Piecewise PDF found from (4.37)-(4.39) for M = 0, S = 1, $\alpha = 0.04162$, L = 0.31. p_1 and p_3 regions are dashed.



(b) PDF compared to histogram of directly-simulated steady-state circle map output.

Figure 4.11

output circle map determined in this way, and from direct simulation, e.g., as in Figure 4.4. Both are plotted against input noise parameter, L.



Figure 4.12: (a) Mean and (b) variance of directly-simulated steady-state circle map *(solid)* vs numerical solution of piecewise PDF *(dashed)* for $M = 0, S = 1, \alpha = 0.04162$.

Instead of numerically integrating as we did in obtaining Figure 4.12, we may continue analytically by defining the *n*th moment of the PDF $p_{\infty}(x)$ as

$$I_{\infty}(n) = I_3(n) + I_1(n) + I_2(n), \qquad (4.40)$$

where

$$I_3(n) = \int_{M-L}^{M+L} t^n p_3(t) dt, \qquad (4.41)$$

$$I_1(n) = \int_{M+L}^{M+S-L} t^n p_1(t) dt, \qquad (4.42)$$

$$I_2(n) = \int_{M+S-L}^{M+S+L} t^n p_2(t) dt.$$
(4.43)

It is clear that we can find C_1 using the normalization condition $I_{\infty}(0) = 1$.

Substituting for p_1 , p_2 , and p_3 in (4.41)-(4.43) from (4.37)-(4.39) we obtain

$$I_{3}(n) \approx C_{3} \frac{(M+L)^{n+1} - (M-L)^{n+1}}{n+1} - \frac{1}{C_{2}\sqrt{\pi L\alpha}(M-L)^{n}e^{\frac{-1}{2L}(M+S-L)}e^{\frac{(M+S-L)^{2}}{4L\alpha}}},$$
(4.44)

$$I_1(n) = C_1 \frac{(M+S-L)^{n+1} - (M+L)^{n+1}}{n+1},$$
(4.45)

$$I_2(n) \approx C_2 \sqrt{\pi L \alpha} (M + S - L)^n e^{\frac{-1}{2L}(M + S - L)} e^{\frac{(M + S - L)^2}{4L \alpha}}.$$
 (4.46)

Here, (4.45) has been obtained by direct evaluation of (4.42), whereas (4.44) and (4.46) are leading-order terms in the asymptotic expansions of the integrals for $\alpha \to 0$ [32]. Substitution for C_2 from above simplifies these to

$$I_3(n) \approx C_3 \frac{(M+L)^{n+1} - (M-L)^{n+1}}{n+1} - C_1 \sqrt{\pi L \alpha} (M-L)^n, \quad (4.47)$$

$$I_1(n) = C_1 \frac{(M+S-L)^{n+1} - (M+L)^{n+1}}{n+1},$$
(4.48)

$$I_2(n) \approx C_1 \sqrt{\pi L \alpha} (M + S - L)^n.$$
(4.49)

Normalization now gives the value of C_1 explicitly:

$$C_1 = \frac{e^{L/\alpha}}{2e^{\frac{\alpha}{4L}}L + e^{L/\alpha}S}.$$
 (4.50)

By evaluating $I_{\infty}(1)$ and $I_{\infty}(2)$ we finally arrive at expressions for the mean and variance of the PDF obtained by small- α asymptotic expan-

sion. That for the mean is

$$\frac{4MLe^{\frac{\alpha}{4L}} + \left(2(M - L + \sqrt{\pi\alpha L}) + S\right)Se^{L/\alpha}}{4Le^{\frac{\alpha}{4L}} + 2Se^{L/\alpha}}.$$
(4.51)

We may check the accuracy of this expression by plotting it against the mean found from direct simulation of the circle map. The result, plotted against input noise parameter L, is shown in Figure 4.13. It can be seen that the approximation is quite accurate, displaying characteristics similar to that of the mean found by numerical solution as shown in Figure 4.12.



Figure 4.13: Mean of directly-simulated steady-state circle map *(solid)* vs numerical solution of piecewise PDF using asymptotic integral expansion *(dashed)* for $M = 0, S = 1, \alpha = 0.04162$.

The expression for the variance, found by evaluating $I_{\infty}(2)$ is not accurate, however, in contrast to the case where we found the PDF numerically rather than asymptotically expanding the integrals (Figure 4.12(b)). This indicates that use of the leading-order asymptotic integral expansions for the left- and right-hand sections of the PDF is not sufficiently accurate for moments of order higher than 1. A noteworthy feature of the plots of Figures 4.12 and 4.13 is that for L > 0.1 there is an almost constant offset between the mean and variance obtained using our asymptotic expansion and that obtained using direct numerical simulation. The reason for this is unclear, but is evidently an effect of neglecting the regions of size α in the piecewise PDF, or the use of the first-order Taylor expansions in obtaining equations (4.37)-(4.39).

4.5 Summary

In Section 4.1 we formulated the steady-state behaviour of the noisy, non-linear DPLL as a modified circle map with probabilistic jumps at each time step. This is the first of our new results, and allows us, in later sections, to perform analysis previously not undertaken on the DPLL behaviour in this noisy regime. In Section 4.2 we recap on the earlier numerical results from [3] and overlay the corresponding circle map responses and show how they agree for a range of input noise levels.

Continuing with our analysis, we show in Section 4.3 that the modified circle map is a Markov chain and we write the Chapman-Kolmogorov equation for its time-dependent PDF. We show numerically that the PDF reaches a steady-state and find a non-local equation for PDF in this state. From the equation we derive some interesting properties of the PDF, and continue in Section 4.4 to find two separate approximate solutions, one in the large-noise limit, and the other in the small- α limit with small but not vanishingly small noise. In each case we illustrate how the asymptotic solutions for the mean and variance, as well as the PDF, agree well with the direct simulation results with their range of validity.

Chapter 5

Results for the first-order DPLL with loop delay

This chapter continues the presentation of our results. Here, we consider the same frequency-quantized first-order DPLL we had in the previous chapter, but where the delay D in the loop is greater than 1. This was the case previously investigated by Gardner in [3] where a 'dip' in the phase error variance was noted as the variance of the additive noise was increased.

In the implementation of the DPLL as a digital electronic circuit, the delay is inherent in the loop, i.e., the output value from the NCO cannot be fed back into the NCO via the phase detector and loop filter until at least the next DPLL clock cycle. Often loop design and processing constraints may result in a greater lag across the circuit components. Therefore, the case where D > 1 is of interest from an engineering perspective.

The cumulative delays in the loop are modelled as a single delay element before the NCO as shown in Figure 4.1. As outlined in Section 3.2.1, this combination of all loop delays into a single block is not valid for all DPLL models. However, it does apply to a number of different situations, including where the delay is in the output of the phase-detector or loop filters for the first- and second-order loops [2]. The single delay element is the only model we consider in this thesis. In Section 5.1 of this chapter, we introduce the equations governing the operation of the first-order DPLL with delay and recap on the results from numerical simulation, as in [3]. We also see from simulations that the corresponding delayed circle maps exhibit the same behaviour in the regions of interest, in particular the curious 'variance dip' seen for the full DPLL. In Section 5.2 we take a more detailed look at the time-domain behaviour of the delayed circle map and see how the additive noise causes a reduction in the range of output phase errors. Finally, in Section 5.3 we use a small- α approximation to perform a detailed analysis of the time-domain behaviour for a sample delay value D = 6. We find an approximate PDF for the output phase error for this case, and also show how the approach extends to general $D \neq 6$. We see that the mean and variance obtained from this PDF compares well to the plots found from direct simulation of the circle map.

5.1 Preliminaries

Equation (3.87) from Section 3.2.3 is now generalized to the following:

$$x_f(t) = K\sin(\Phi(t-D)). \tag{5.1}$$

Proceeding as in Section 3.2.3, we obtain the full delayed-difference equation of operation for the DPLL:

$$\Phi(t) = \Phi(t-1) + \frac{2\pi}{2^b} \left\{ \mu - \text{Int} \left[2^b K \left(\sin \Phi(t-D) + N(t-D) \right) \right] \right\},$$

for $t = 0, 1, 2, \dots$
(5.2)

Again, as in Section 4.1, we can define a corresponding delayed circle map as

$$x(t) = x(t-1) + \alpha + Q(x(t-D))S$$
, for $t = 0, 1, 2, ...$

$$x(0) = x_0,$$

where, as before, we have

$$Q(x) = \begin{cases} -1 & \text{with probability } F_N(x), \\ 0 & \text{with probability } 1 - F_N(x). \end{cases}$$
(5.3)



(a) DPLL variance response.

(b) DPLL with circle maps.

Figure 5.1: Variance of output jitter, σ_O^2 , as a function of input noise variance, σ_N^2 , for (a) DPLL with $\mu = 0.56$, $K = 2^{-6}$, b = 8 for the nodelay (D = 1) case *(solid)*, D = 2 *(large dash)*, D = 4 *(med dash)*, D = 8*(small dash)*, and (b) for DPLLs with corresponding circle rotation maps.

In (5.3), the CDF, $F_N(x)$, is dependent on the distribution of the additive input noise exactly as before, described in Equation (4.5).

Numerical simulations, similar to those conducted in Section 4.2, once again confirm that the delayed circle map is a good approximation for the delayed DPLL. As in the Section 4.2 simulations, the noise here is uncorrelated and uniformly distributed with zero mean to match our analysis later. However, the response obtained with Gaussian noise was virtually indistinguishable from our plots here.

Figure 5.1 shows the equivalent variance responses to those shown earlier for the D = 1 case in Figure 4.4. The plot in Figure 5.1(b) shows that, for the DPLL with b = 8 and various delay values, the output variance plots for the corresponding circle maps track those of the DPLLs very closely for all but very large additive noise levels. In these latter regions, the DPLL variance increases with slope 1, in the same manner as for an analogue PLL.

A striking feature of Figure 5.1 is that within the intermediate region, between the flat, quantizer-only output jitter on the left and the analogue limit on the right, the DPLLs with higher delay values exhibit a significant decrease in output jitter variance. This is particularly evident in the case where D = 8. This was also seen previously in the numerical results of Figure 2 in [3]; the phenomenon was left unexplained in this work. It is certainly not intuitive that certain additive input noise would actually cause a decrease in the digital noise inherent in the loop! As is clear from Figure 5.1(b), the delayed circle map exhibits the same effect, and appears to capture the DPLL behaviour very well in this region. The remainder of this section focusses on the dynamics of the delayed circle map in the low and intermediate additive noise level regions.



Figure 5.2: Variance of output jitter, σ_O^2 , as a function of input noise variance, σ_N^2 , for delayed circle map M = 0, S = 1, D = 6, and α values 0.056 (solid), 0.11, 0.21, 0.56 (dashed).

The output variance response for a particular delayed circle map (D = 6) is shown in Figure 5.2. It is evident that the size of the dip in variance depends greatly on α , the dip being larger for small α . For $\frac{\alpha}{S} = 0.056$ the addition of input noise reduces the output variance by a factor of almost 5 in the case where D = 6. In order to investigate this variance dip, the analysis and results that follow all assume that α is small, i.e., $\alpha \ll S$. Note that this corresponds to the case for the original DPLL where the

fractional part of the input frequency, $Frac[\mu]$, is small.

Additive input noise is generally considered an unwanted nuisance that can at best be mitigated. In this context the dip in the output variance is unusual phenomenon since it is a system response that would generally be considered desirable, i.e., a reduction in output noise, caused by the addition of noise at the input. In this sense the effect looks similar to *dithering* [33], a technique used in signal processing where noise is intentionally added in order to remove patterns in quantization noise. However, dithering removes harmonics by whitening the spectrum of quantization noise, rather than actually reducing its magnitude. The phenomenon noted here also bears a resemblence to stochastic resonance [34]; here we see a local minimum in output noise for a certain level of input noise, while a system exhibiting stochastic resonance has a critical level of input noise where the output SNR has a local maximum [35]. However, in this case we do not have a signal of interest that we are trying to detect in a noisy environment, and nor is there any apparent 'cooperation' between an input signal and additive noise. Instead, we have an output noise in the form of limit cycles that simply seem to get disrupted, as we will see in the following section.

5.2 Analysis

Figure 5.3(a) shows the steady-state behaviour of the circle map with delay in the small- α case where the output is unaffected by additive noise. The 'sawtooth' output pattern is similar to the case of the circle map without delay (D = 1). The main change in behaviour is that x(t) will continue to increase in steps of α as long as x(t - D + 1) < T. Similarly, x(t) will decrease in much larger steps of $(S - \alpha)$ as long as x(t - D + 1) > T. This is illustrated in Figure 5.3(b); for the non-delayed case, the map would have started to decrease once the output exceeded $T = (M + S - \alpha)$, i.e., at point C. However, in the case shown, where D = 6, the map continues upwards in α -steps until 6 sequential output values have exceeded T, i.e., point H is reached. At that point, the map



(a) Multiple iterations of limit cycle. (b) Limit cycle detail.

Figure 5.3: Time-domain plots of delayed circle map output for M = 0, S = 1, D = 6, $\alpha = 0.056$ with no additive noise.

x(t) jumps downwards in 6 steps of $(S - \alpha)$ until x(t - D + 1) > T, leading to the pattern shown in Figure 5.3(a).

In this case, therefore, the range of the output jitter is bounded by $[M+(1-D)S+(2D-2)\alpha, M+S+(D-1)\alpha]$. This is a region of size $DS+(1-D)\alpha \approx D(S-\alpha) \approx DS$ for small α . This range is approximately D times greater than the range of the non-delayed map, which was S. This agrees with Gardner's observation (7), noted in Section 3.2.3, that the peak-to-peak jitter for the DPLL was approximately $2\pi D/2^b$; in Section 3.2.5 we saw that the DPLL with D = 1 obeys a circle rotation map with $S = 2\pi/2^b$.

If $\frac{\alpha}{D(S-\alpha)}$ is irrational (i.e., $\frac{\alpha}{S}$ irrational as before), the jitter is roughly uniformly distributed in a region of size $DS - (D-1)\alpha$ with variance $\frac{(DS-(D-1)\alpha)^2}{12}$. This result relies on the assumption that α is small, so that the majority of the map output points are on the upwards portion of trajectory, i.e., we can ignore the downward $(S - \alpha)$ steps for the purposes of calculating the PDF.

With the addition of noise to the quantizer, some points on our noisy delayed circle map (5.3) are now subject to having their quantizer values modified. In the simulations shown in Figure 5.4 the noise is uncorrelated and uniformly distributed in [-L, L] so that its variance is $L^2/3$. In this case, the noise CDF $F_N(x)$ is the the linear ramp as shown in Figure 4.2,

and only points that fall within the noise ramp are subject to changes in quantizer values.

Figure 5.4 shows how the dynamics of the map are changed by the addition of the quantizer noise. The change occurs – circled in Figure 5.4(a) – when one of the points within the noise-affected region [T-L, T+L has its quantizer changed from the noise-free value. In this case, points A, B, C, and D are within the noisy region, C being marginally above the line x = T. Without noise, the quantizer value at B would have been 0 so that the trajectory would have continued to increment from G to H as in Figure 5.3. However, in this case noise has changed the quantizer value at B to -1 so that the trajectory decrements by $(S - \alpha)$ instead. From there it decrements again, as shown in Figure 5.4(b), indicating that the quantizer value at C is -1, the same as its noise-free value. However, at the next step, the trajectory increments, because the noise has changed the quantizer value at D (circled) to 0 from its normal value of -1. After this, the map decrements a further 3 times, corresponding to points E, F, and G earlier in the map. Thus, the map has decremented 5 times in total, rather than the usual 6. The reason for this is the change of the quantizer value at D from -1 to 0, which has reduced the total number of downward steps.

In general, when a point on the upward part of the trajectory has its quantizer value changed from 0 to -1 (such as point B in Figure 5.4), the effect will be to reduce the range of the map from the upper end of the range, e.g., point H is removed in Figure 5.4(c). However, it does not lead to an increase in the number of downward steps since, while the trajectory starts decreasing earlier than usual, there will also be fewer points above the line x = T yielding further downward steps in the map. The result of this change is to reduce both the upper and lower endpoints of the map by α , which we are assuming to be small. On the other hand, if a point on the trajectory has its quantizer changed from -1 to 0, as in point D in Figure 5.4, the result is a reduction in the number of (much larger) downward steps of size $(S-\alpha)$. This causes a significant reduction in the range of the map for that particular iteration, by increasing the



Figure 5.4: Time-domain plots of delayed circle map output for M = 0, S = 1, D = 6, $\alpha = 0.056$ with input noise variance $\sigma_N^2 = 0.04$.

lower endpoint, as shown in Figure 5.4(a).

Since the overall effect of the noise is to either leave the range of the map almost unchanged, or to reduce it by a number of $(S - \alpha)$ -steps from the lower endpoint, the result is both an increase in the mean and a decrease in the variance of the output jitter within this particular noise regime. This accounts for the dips seen in Figures 5.1 and for the larger dips for small α values in 5.2. Note that the analysis here relies on the asymmetry brought about by the small- α approximation. It means that our upward α -steps are much smaller than the downward $(S - \alpha)$ steps, and we also assume that a single downward step takes the map below the switching region [T - L, T + L] in all cases, e.g., in Figure 5.4(b). This requires

$$D\alpha \ll S - L. \tag{5.4}$$

The histograms from direct simulations of the delayed circle map are shown in Figure 5.5 and confirm the analysis above. For low noise as in Figure 5.5(a), the distribution is essentially uniform on the range $[M + (1 - D)S + (2D - 2)\alpha, M + S + (D - 1)\alpha]$ apart from 'spikes' corresponding to the larger $(S - \alpha)$ downward steps of each iteration. As the noise level is increased, as shown in (b)-(d), the distribution becomes more concentrated at the upper end of the range. This corresponds to the analysis performed above. The entire distribution remains contained within the original range. As the noise level is increased further (e)-(f), the distribution becomes more symmetric, and approximately Gaussian. The mean also begins to shift significantly as the noise level is increased, i.e., the entire distribution moves downwards. This corresponds to the similar result in Theorem 7 for large-L in the case of the circle map without delay, where the mean was shifted according to $sgn(\frac{\alpha}{S} - 0.5)$.



Figure 5.5: Histograms of steady-state circle map output for M = 0, S = 1, D = 6, $\alpha = 0.056$, and input noise variance (a) 0.001, (b) 0.01, (c) 0.05, (d) 1.0, (e) 10, (f) 100, similar to those of the D = 1 case in Figure 4.5. The predicted uniform distribution for the noise-free case is also shown (shaded) for comparison.

5.3 Small- α results

We can derive some quantitative results by looking at a specific case in more detail, in particular the case where D = 6 and where α is small, i.e., $\alpha \ll S$ and $D\alpha \ll S - L$. In addition to these assumptions, we also assume the magnitude of the additive input noise is such that the exact number of points on the circle map trajectory closest to the switching curve x = T that fall within the linear ramp of the noise PDF, n_L , is 4. For this, we necessarily have

$$3\alpha < 2L < 4\alpha. \tag{5.5}$$

This is the situation shown in Figure 5.4(c). Note that the particular values of D and n_L chosen here are to ease the presentation of the analysis only. The same approach extends seamlessly for general D and n_L ; we use other values of n_L to obtain the PDFs in Figure 5.6 below. Also, in general n_L is a function of L, the additive noise magnitude, as well as the precise position of the points on the upward portion of map trajectory. Here, however, we ignore the latter dependency, as between iterations the upward-trajectory points will vary in position by at most α , which is small. Therefore, each integer value of n_L defines a an input noise level $L \in \left[\frac{(n_L-1)\alpha}{2}, \frac{n_L\alpha}{2}\right]$ for which that value of n_L is valid.

In Figure 5.3(b), where there is no noise, the values of the quantizer at each of the points A-D are:

$Q(x_A)$	$Q(x_B)$	$Q(x_C)$	$Q(x_D)$
0	0	-1	-1

With these noise-free quantizer values we simply get the trajectory shown in Figure 5.3(a). However, when noise is added with $n_L = 4$, the quantizer outputs become random variables with probability distribution dependent on the noise PDF $F_N(x)$. However, the quantizer values can only be -1 or 0, leading to a total of 16 possible combinations of quantizer values across the four points. With the noise-free quantizer values above, the range of the next iteration of the delayed circle map is $6(S - \alpha)$, as noted in Section 5.2 and shown in Figure 5.3(a). For each of the other cases, the ranges of the maps, R_i , can be found be calculating the trajectory from the point A onwards. These may be enumerated as shown in Table 5.1.

Next, if we define δ as $x_A - (T - L)$, where we have

$$0 \le \delta \le 2L - 3\alpha \le \alpha, \tag{5.6}$$

we can find the probabilities of the occurrence of the quantizer values from the linear noise PDF, $F_N(x)$, as follows:

$$\begin{cases}
P(Q(x_A) = -1) = \delta/2L \\
P(Q(x_B) = -1) = (\delta + \alpha)/2L \\
P(Q(x_C) = 0) = (2L - \delta - 2\alpha)/2L \\
P(Q(x_D) = 0) = (2L - \delta - 3\alpha)/2L
\end{cases}$$

i	$Q(x_A)$	$Q(x_B)$	$Q(x_C)$	$Q(x_D)$	R_i	Y_i
0	0	0	0	0	$6S - 6\alpha$	$M + S + 7\alpha$
1	0	0	0	-1	$6S - 6\alpha$	$M + S + 6\alpha$
2	0	0	-1	0	$5S - 6\alpha$	$M + S + 5\alpha$
3	0	0	-1	-1	$6S - 6\alpha$	$M + S + 5\alpha$
4	0	-1	0	0	$4S - 6\alpha$	$M + S + 4\alpha$
5	0	-1	0	-1	$5S - 6\alpha$	$M + S + 4\alpha$
6	0	-1	-1	0	$5S - 6\alpha$	$M + S + 4\alpha$
7	0	-1	-1	-1	$6S - 6\alpha$	$M + S + 4\alpha$
8	-1	0	0	0	$3S - 6\alpha$	$M + S + 3\alpha$
9	-1	0	0	-1	$4S - 6\alpha$	$M + S + 3\alpha$
10	-1	0	-1	0	$4S - 6\alpha$	$M + S + 3\alpha$
11	-1	0	-1	-1	$5S - 6\alpha$	$M + S + 3\alpha$
12	-1	-1	0	0	$4S - 6\alpha$	$M + S + 3\alpha$
13	-1	-1	0	-1	$5S - 6\alpha$	$M + S + 3\alpha$
14	-1	-1	-1	0	$5S - 6\alpha$	$M + S + 3\alpha$
15	-1	-1	-1	-1	$6S - 6\alpha$	$M + S + 3\alpha$

Table 5.1: Per-iteration ranges and upper bounds for each quantizer value in case of circle map with D = 6 and $n_L = 4$.

From this, we can get the expected value of the range, R, as follows:

$$\mathscr{E}[R] = P(Q(x_A) = 0)P(Q(x_B) = 0)P(Q(x_C) = -1)P(Q(x_D) = -1)R_1 + \dots + P(Q(x_A) = 0)P(Q(x_B) = -1)P(Q(x_C) = 0)P(Q(x_D) = 0)R_{16}$$
(5.7)

We can extend this analysis by noting that in Table 5.1 the range of the iteration in each case, R_i , is $(D-l_0)S-D\alpha$, where l_0 is the number of zeros following the leading -1 in the set $\{Q(x_A), Q(x_B), Q(x_C), Q(x_D)\}$, the quantizer values in that particular case. As enumerated, this is the number of zeros following the leading 1 in *i*, when *i* is represented in binary form. Indeed, this follows in general, for cases both were the number of points within the linear ramp, n_L , does not equal 4, and for different delay values, *D*. In general, we require $\frac{n_L}{2} \leq D$ and $S \gg$ $(n_L + D)\alpha$. In cases where $(D - l_0) < 1$ we set R_i to $(S - \alpha)$ to reflect that fact that there will always be at least one downward $(S - \alpha)$ -step per iteration of the map.

The probability of each case, as well as the expected range for the resulting map iteration, were used in the calculation of $\mathscr{E}[R]$ above. This could be used to calculate a variance, assuming a uniform output distribution within the expected range. However, looking at the histograms in Figure 5.5, this does not seem to be a valid assumption. It was found that much better numerical results were obtained if, rather than simply using the magnitude of the range for each case, the expected lower and upper endpoints of the range were used instead to derive an overall PDF for the map. The PDF can then be used to calculate the moments in order to find the mean and variance of the distribution.

In the noise-free case, the upper bound on the map output is M+S+ $(D-1)\alpha$, and the maximum value actually attained be the map will be within α of this limit. In the case where noise is added with $n_L = 4$, the upper bound depends on which of the points A-D has the first quantizer



Figure 5.6: PDFs found numerically from D = 6 analysis for delayed circle map M = 0, S = 1, $\alpha = 0.056$ for n_L values (a) 2, (b) 5, (c) 7, (d) 9. The uniform distribution expected for the noise-free, $n_L = 0$, case is also shown for each.

output of -1. In the noise-free case this is point C but this will vary when noise is added. In Table 5.1, the upper bound is determined by the position of the leading -1 in the set $\{Q(x_A), Q(x_B), Q(x_C), Q(x_D)\}$. Again, this is exactly the position of the leading 1 in *i* when represented in binary form, which is easily found as *floor* $[\log_2(i)]$. These upper bounds are shown as Y_i in Table 5.1. This approach once again generalizes for arbitrary D and n_L .

Knowledge of the probabilities, upper endpoints and ranges of the map for each of the $2^4 = 16$ cases is sufficient to numerically calculate the expected PDF of the map. This is because we can assume, for small α , that the map is uniformly distributed between the upper and lower endpoints. It is reasonable to use the Y_i values as approximations to the upper endpoints as these incur an error of at most α , which is small. In

general, the probabilities are functions of the unknown δ , which is less than α . However, it was found that the numerical results that follow were not sensitive to varying values of δ so a random value of δ was chosen in the range $[0, \alpha]$. The results may be improved slightly by averaging over δ in this range.

Therefore, the approximate PDF may be simulated simply as

$$P(x) = \sum_{i=0}^{15} P_i \mathbf{U}(Y_i - R_i, Y_i, x),$$

where P_i is the probability of each of the 16 cases and $\mathbf{U}(a, b, x)$ is the uniform distribution on [a, b]. The simulated PDFs for fixed delayed circle map parameters and various values of n_L are shown in Figure 5.6.



Figure 5.7: PDFs from graphics (c) and (d) in Figure 5.6 compared with histogram of direct circle map simulation for M = 0, S = 1, $\alpha = 0.056$ and (a) L = 0.169, (b) L = 0.225.

These PDFs can be directly compared with histograms of the direct map simulation as in Figure 5.5, by choosing a value of L in in the simulation within the valid range for the particular value of n_L , i.e., $\left[\frac{(n_L-1)\alpha}{2}, \frac{n_L\alpha}{2}\right]$. In the simulations here we chose $L = (n_L - 1)\alpha/2$ to obtain the best fit for the data, based on a visual inspection of the plots. The match between the PDFs and the histograms are shown in Figure 5.7 for $n_L = 7$ and $n_L = 9$. The plots indicate that our analysis captures the actual map dynamics very well within the noise regime of interest. The main difference between the simulation output and our PDFs are the six 'spikes' in the histogram at the left-most corner of each of the probability 'steps'. These are simply an artefact of the downward portion of the map trajectory, which we have ignored in our analysis.

The PDFs of Figure 5.6 can be used to derive the mean and variance of the map for each value of n_L . The variance is plotted in Figure 5.8 and compared with the variance curve from the direct simulation of the delayed circle map. Again, the variance obtained for each n_L is placed on the graph at a position corresponding to $L = (n_L - 1)\alpha/2$ to obtain the best fit to the data. It is clear from the graph that the analysis provides a good fit to the numerical data for smaller values of n_L . However, for larger noise levels the approximation breaks down somewhat as condition (5.4) is violated and the switching region, [T - L, T + L], is entered multiple times in a single iteration.



Figure 5.8: Variance of output jitter, σ_O^2 , plotted against input noise variance, σ_N^2 , for direct simulation of the circle map circle map with M = 0, S = 1, D = 6, and $\alpha = 0.056$. Numerical results from D = 6 analysis also shown (boxes) where variance is calculated from histograms as in Figure 5.6.

The mean can similarly be found from the PDF and is shown in Figure 5.9. Figure 5.9(a) shows a large-scale plot of the mean for the D = 6 case along with the corresponding case without delay (D = 1), the latter being similar to what was plotted previously in Figure 4.6. For the delayed map, in general, as the input noise level increases the mean decreases in a similar way to the non-delayed map. In general, the mean

will increase or decrease according to $sgn\left[\frac{\alpha}{S}-1\right]$, corresponding to an earlier result found in Figure 4.6 and Theorem 7.



Figure 5.9: (a) Mean of output jitter, μ_O , plotted against input noise variance, σ_N^2 , for direct simulation of the delayed circle map circle map with M = 0, S = 1, D = 6, $\alpha = 0.056$, compared with corresponding non-delayed case D = 1 (dashed). (b) Comparison with numerical results from D = 6 analysis (boxes).

However, for the delayed case, there is a regime where a small amount of additive noise causes the mean to increase rather than decrease, i.e., the region $\sigma_N^2 < 0.1$, shown in Figure 5.9(b). This corresponds to the region where the analysis performed in this section is a good match to the direct simulation. Indeed, the analysis does predict an increase in the mean, as the additive noise causes a 'loss' in downward $(S - \alpha)$ steps, in turn leading to a reduction in the range of the circle map from the lower end point. The result is that the lower end of the distribution gets pulled upwards. This is clearly evident in the PDFs of Figure 5.6also; the distribution gets 'pulled' towards the upper end of the range for increasing noise levels within this initial regime. It is clear, once again, from Figure 5.9(b), that the predicted mean from our analysis of this section is a good match for the simulation results for small σ_N^2 but starts breaking down once we approach $\sigma_N^2 = 0.1$. For additive noise levels greater than this, the mean begins decreasing again (and variance increasing) until the analogue characteristics take over, as in the righthand side of Figure 5.1(b). The behaviour here is identical to that in the non-delayed case, which was investigated in Theorem 7.

5.4 Summary

We began this chapter by recapping the previous results found by Gardner in [3] on the variance response of the first-order DPLL with delay. In particular, our numerical simulations confirm the variance 'dip' previously observed and otherwise match well to Gardner's results. For our new results, we extend to this D > 1 case our previous analysis where we formulated the steady-state DPLL behaviour as a unique form of noisy circle map. Here, the delayed noisy DPLL can be modelled as a delayed noisy circle map, just as for the non-delayed case in Section 4.1. Numerical simulations of the circle map with delay show similar results to those for the non-delayed map in Section 4.2, i.e., the variance response of the circle map tracks that of the DPLL for low levels of noise through an 'intermediate' region that includes the variance dip for this delayed case. Only for large additive noise levels, in the analogue PLL limit, do the responses of the DPLL and circle map diverge. We also show histograms of the noise-affected circle map output and we see that, in the region of the variance dip, the distribution gets pulled towards the upper end of the range, causing a reduction in the output variance and an increase in the mean.

From our simulations we also see a new result: the size of the variance dip is dependent on α , the dip being greater for smaller α . To gain an insight into this behaviour we perform an analysis, looking at the steady-state of the delayed circle map in detail and seeing how the additive noise disrupts the limit cycles. We see how the noise has the effect of asymmetrically reducing the range of the limit cycles, from the lower endpoint only so that the variance is decreases and the mean increases with increasing additive noise. The magnitude of the asymmetry is greater for smaller α , explaining why the variance dip is greater for small α .

Finally, we derive some quantitative results for a particular case, D = 6. We see that, for a sample input noise level, where exactly 4 points of the circle map are within the noise-affected region, we can obtain an

approximate PDF for the circle map output by considering each of the 2^4 possible iterations individually. We show how the analysis extends to arbitrary D and number of noise-affected points. For D = 6 we obtain PDFs for a range of noise-affected points, which allows us to graph the approximate output mean and variance as functions of the input noise. Comparing with the direct simulation runs on the circle map, we see the approximate mean and variances are good matches with the direct results for small noise levels, but become less accurate as the noise increases due to a breakdown in the assumptions on which our analysis was based. Beyond this noise level lies the region where additive noise dominates over quantization effects and the variance increases as an asymptote with slope 1/2, just as it did in the non-delayed case; this behaviour is well understood.

Chapter 6

Results for the second-order DPLL

This chapter completes the presentation of our results. In Section 6.1 we return to the second-order DPLL that we introduced in Section 3.2.5 and show how the equations of operation are modified in the case of additive input noise. In Section 6.2 we present some numerical results, similar to those for the first-order loop in Section 4.2, showing the input-output variance response of the second-order loop and comparing it to the first-order case. Finally, in Section 6.3 we look again at the limit cycles we had in Section 3.2.5 and, using numerical simulations, show how they are affected by additive noise.

6.1 Preliminaries

We return again in this section to the equations (3.89) for the secondorder DPLL with minimal loop delay D = 1:

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \left(\mu - \operatorname{Int}[2^b (K_1 \sin \Phi(t) + u(t))] \right) \mod 2\pi,$$

$$u(t+1) = u(t) + K_1 K_2 \sin \Phi(t+1).$$
(6.1)

We saw in Section 3.2.5 and in Figure 3.22 how, on a phase-plane plot,

the trajectory moves from the initial conditions (assumed to be close to $\{\Phi, u\} = \{0, 0\}$) into a corridor bounded by the lines:

$$L_1 : u = 2^{-b} - K_1 \sin\left[\Phi + \frac{2\pi}{2^b}(1-\mu)\right] + K_1 K_2 \sin\Phi,$$

$$R_1 : u = 2^{-b} - K_1 \sin\left[\Phi + \frac{2\pi}{2^b}(-\mu)\right] + K_1 K_2 \sin\Phi.$$

Between L_1 and R_1 lies the switching curve S_1 . This curve corresponds to a change of direction in the motion of Φ : to the left of the curve, Φ increases, and to its right it decreases, the overall effect being to keep the trajectory between L_1 and R_1 . From the initial condition where u is close to zero, we can see from equation (6.1) that u will continue to increase as long as Φ is positive. When the L_1 - R_1 corridor intersects with the $\Phi = 0$ axis, u reaches a steady state, incrementing or decrementing depending on the sign of Φ . It should be noted that, close to $\Phi = 0$, the width of the L_1 - R_1 corridor is approximately $2\pi/2^b$.



Figure 6.1: Phase-plane plot for frequency-quantized second-order DPLL with b = 8, $\mu = 0.56$, $K_1 = 2^{-6}$, $K_2 = 2^{-5}$, $\{\Phi(0), u(0)\} = \{0, 0\}$, and without additive noise.

Figure 6.1 shows a particular case of the steady-state motion in Φ and u, that is the situation where μ is rational and the system starts from the origin. As seen in Section 3.2.5, for a low-denominator rational μ , the combined motions in Φ and u result in a steady-state trajectory that is a circle rotation map for both Φ and u. The map for Φ has period q, where $\operatorname{Frac}[\mu] = p/q$ in its lowest form. The qth iteration of the map is also a circle map in u, with rotation number that in general depends on the initial conditions. The combined period-25 circle map for Φ and u that we had in Section 3.2.5 is shown again here; the rotation in Φ has period 25, and the 25th iterate in the map gives a fixed point in u. We will see in Section 6.3 below how this limit cycle is disrupted by additive input noise.

In the same way as the equation for the first-order loop was modified to obtain (4.2), the equations for the noisy second-order loop become

$$\Phi(t+1) = \Phi(t) + \frac{2\pi}{2^b} \left(\mu - \operatorname{Int}[2^b \left(K_1 \left\{ \sin \Phi(t) + N(t) \right\} + u(t) \right)] \right) \mod 2\pi$$
$$u(t+1) = u(t) + K_1 K_2 \left\{ \sin \Phi(t+1) + N(t+1) \right\}.$$
(6.2)

Here again N(t) is the noise added to input signal $\sin(2\pi t\nu)$, which can equivalently be modelled here as being added after the phase detector. We would like to be able to investigate the behaviour of this system for various levels of the input noise N(t). In particular, we would like to be able to characterize the variance of the output phase error, Φ , as well as understand how the steady-state limit cycles are disrupted by the input noise. We investigate both aspects in the following sections.

6.2 Numerical results

The pair of equations (6.2) describing the second-order system are easily simulated for various levels of input noise, N(t). To compare to the corresponding results for the first-order loop, the particular noise we choose is uncorrelated at each time step and drawn from a uniform density. We chose the uniform density for the first-order case as it allowed us to derive some particular results, such as Theorem 7. For the numerical simulations, the results obtained using the uniform distribution and the more standard Gaussian (normal) distribution are virtually indistinguishable. Therefore we persist with the uniform density here. Simulation results for the variance are shown in Figure 6.2. The overall response of the second-order loop is remarkably similar to that of the first-order loop. However, this is quite intuitive, given that the steady-state behaviour of Φ for both in the noise-free case is essentially the same as remarked in Section 3.2.5.



Figure 6.2: Comparison of output variance, σ_O^2 , of first-order *(solid)* and second-order *(dashed)* DPLLs as a function of input noise variance. DPLLs have b = 8, $\mu = 0.56$, $K_1 = 2^{-6}$, and $K_2 = 2^{-5}$ for the second-order loop.

For low levels of additive noise, the output of the second-order system is identical to that of the first, namely a flat response with value equal to the variance of the original circle map in Φ . For a rational input frequency, for example, this is $((q^2 - 1)/12q^2)(2\pi/2^b)^2$ using Lemmas 1 and 2, and in this case we have q = 25.

Through the intermediate region, where the output variance first starts to rise in response to an increase in input noise, the output of the second-order loop continues to follow that of the first. This indicates that the analysis performed on the first-order system in Section 4 likely applies in an approximate way to the second-order loop also.

Finally, for large input noise levels, both first- and second-order loops have a similar response, with output variance increasing linearly with slope 1 as would be expected for linear, analogue systems, and as noted in [3]. It is clear from the plots that there is little or no penalty incurred by using a second-order loop as opposed to a first. This is important for cases where a frequency offset between the PLL and the source causes an undesired static phase error. This can be readily eliminated by use of a second-order loop.

As previously mentioned, the fact that the low-noise response of both loops is identical should come as no surprise given the analysis of the second order loop in Section 3.2.5. Additionally, the slope-1 large-noise asymptote is to be expected in both cases as the quantization effects are swamped by the input noise and both loops behave as standard linear, analogue systems. The only slight surprise is how close the responses of both systems are through the intermediate region. However, we saw previously in Section 3.1.6, by looking at the Fokker-Planck equations, that the response of the second-order analogue PLL was approximately the same as the first for moderate levels of input noise; see equation (3.66), for example. Therefore it may not be entirely surprising that there are similarities in the digital case also. We look at the dynamics in this regime in more detail in the following section.

6.3 Analysis

6.3.1 Second-order limit cycles

It is instructive to examine the steady-state phase-plane plots of the second-order loop at the point where the additive noise first starts to disrupt the limit cycles. Some of these are shown in Figure 6.3 for the low-denominator rational input frequency we had in Section 3.2.5, and recapped in Section 6.1 above.

Looking first at Figure 6.3(a), where the variance of the additive noise N(t) is 10^{-10} , it is clear that the same period-25 motion is present in Φ , unaffected by the low level of noise. However, for each value in the cycle for Φ there is a greater range in the values of u as compared to the noise-free case in Figure 6.1, where the motion in u was also periodic.


Figure 6.3: Similar phase-plane plots as in Figure 6.1 with additive noise variances, σ_N^2 equal to (a) 10^{-10} , (b) 10^{-9} , (c) 10^{-7} , and (d) 10^{-4} .

This effect is quite similar to that seen in the noise-free limit cycle shown earlier in Figure 3.24(a) where the system started from $\Phi(0) = 10^{-4}$. Hence, both a non-zero starting point for Φ or a small amount of additive noise have a similar effect: the rotation in Φ is unaffected while the perturbation is enough to break the periodic cycle in u and instead lead to a circle map with irrational rotation number. The variance of the output phase jitter is simply that of the original circle map, corresponding to the flat part of the graph in Figure 6.2.

We saw in Section 4.1 how the noisy first-order DPLL corresponds to a circle map with a 'noisy quantization'; see, for example, equations (4.1) and (4.4). In that case, the quantizer that originally changed its output at $x = M + S + \alpha$ became noisy; values of $x < M + S + \alpha$ gave values of Q(x) = -1 rather than the original 0, and similarly for points $x \ge M + S + \alpha$. The point $x = M + S + \alpha$ corresponds to our switching curve, S_1 , in the second-order case.

In Figure 6.3(a), the noise in the steady-state points (visible in the u direction only) is not yet great enough to move any of the points across the switching curve to result in a new value for Φ . Note that in equation (6.2) for Φ , the noise enters the quantizer in two separate terms: explicitly as N(t) and also via the noisy u(t). The quantizer value will be modified either when the noise in u (visible in Figure 6.3(a)) pushes the trajectory across the switching curve, or when the trajectory is close enough to the switching curve so that the N(t) term results in a change of quantizer value.

In Figure 6.3(b) the noise level is high enough to cause sporadic changes in quantizer values. The section of the time series plotted shows several iterations of the original period-25 orbit (with noise in the u direction) before the trajectory jumps to another similar periodic orbit at higher values of u. It would seem that this orbit is more stable in that the trajectory remains there for all observable time.



Figure 6.4: Larger-scale plot of Figure 6.3(b).

A larger-scale plot of the transition between orbits is shown in Figure 6.4. It can be seen that the last point on the trajectory before it jumps to the higher orbit is very close to the switching curve. This was the point on the original period-25 orbit that was closest to the switching curve, slightly to its right. The result was that, in the absence of noise, the next

point on the trajectory would have been to the left of the curve, for a lower value of u. However, in this case, the noise has caused the quantizer to return the value associated with the left-hand side of the switching curve, sending the trajectory off in the opposite direction. This results in an new value for Φ , which is $(2\pi/2^b)(1/q)$ to the right of the original period-25 map. This is very similar to what happened in the noise-free case in Figure 3.24(a) where the irrational rotation in u caused a point to move across the switching curve, resulting in a modified quantizer value.

A similar period-25 orbit to the noise-free case then resumes for higher values of u. Because the switching curve is now further from the $\Phi = 0$ axis, no further orbit jumps occur; the point on the new orbit corresponding to the original jump is now safely to the right of the switching curve. The effect of the noise in this case is to cause a jump between two stable periodic orbits. Using numerical simulation, it can be shown that, in the noise-free case, a number of stable period-25 orbits exist depending on the initial value of u ($\Phi(0) = 0$ in all cases) and that these orbits differ in values of u only, i.e., are shifted on the vertical axis. Additive noise has triggered a jump between two of these orbits, which is an interesting phenomenon. At this noise level, with input variance equal to 10^{-9} , the output jitter variance is still on the flat part of the graph in Figure 6.2, which is to be expected since, apart from the single jump-point, the orbit in Φ is still the original circle map.

The input noise is increased by a further two orders of magnitude in Figure 6.3(c). The first thing to note is that, although the simplicity of the dynamics looks to be lost because of the large variation in u, the motion in Φ is still bounded mostly by L_1 and R_1 . Thus, from input noise variances of 10^{-10} to 10^{-7} there is no appreciable increase in output jitter. This corresponds well to the graph of Figure 6.2. If we remove the transitions between points on the trajectory, we can see that the system's steady-state output still has a relatively simple structure as shown in Figure 6.5(a). It is clear that Φ essentially follows the same period-25 map with only the occasional excursion outside the L_1 - R_1 corridor. The motion in u is more complex; the pattern of the lighter and darker regions indicate more jumps between orbits. The clustering of the *u*-values corresponds to time spent within each orbit. The fact that more jumps occur is to be expected as more points on the period-25 horizontal circle map become susceptible to being pushed close to or across the switching curve.



Figure 6.5: Similar phase-plane plots as in Figure 6.1 with additive noise variances, σ_N^2 equal to (a) 10^{-7} , and (b) 10^{-4} .

Finally, if the input noise is increased by, for example, a further three orders of magnitude, the map eventually gets pushed outside the L_1 - R_1 region. This is shown in Figure 6.3(d) and Figure 6.5(b). While the input noise is now beginning to dominate, it is clear that the underlying discrete dynamics have not been entirely lost. The motion in Φ is still generally confined to a grid of points with a spacing of $(2\pi/2^b)(1/q)$. Many more points are now added both to the left and right of the original circle map as multiple points on that map are now subject to the switching curve noise. At this level of input noise (variance 10^{-4}), the increase in output jitter is now noticeable as can also be seen in Figure 6.2.

Above this noise level it becomes more difficult to discern discrete patterns in the general noise. This is to be expected as from input noise levels greater than 10^{-3} we enter the region where the analogue behaviour takes over.

6.3.2 Comparison to first-order case

The similarities between the behaviour of the first- and second-order loops can be seen by an examination of Figure 6.6. Plot (a) shows the steady-state behaviour of the same second-order map as in Figure 6.5, except with an additive noise level of $\sigma_N^2 = 10^{-5}$. Again, the transitions between points have been omitted so that the distribution of the points in steady-state can be seen. Similar to the case where we had $\sigma_N^2 = 10^{-4}$ in Figure 6.5(b), even though the distribution in the noisy *u* direction seems to lack pattern, much of the structure of the original period-25 circle map in Φ remains. Some additional points outside the original L_1 - R_1 region have appeared, but are still on the grid with spacing $(2\pi/2^b)(1/q)$.

Also illustrated on this plot is the region where the behaviour of the map is affected by noise. This is a corridor of width approximately 2L, centred about the switching curve, S_1 , where $L = \sqrt{3\sigma_N^2}$. This shaded region on the plot corresponds to the points on the phase plane where noise can cause the quantizer in the equation (6.2) for Φ to output a different value to the noise-free case, assuming N(t) is uniformly distributed with variance σ_N^2 . The concept here is identical to that we used earlier to modify the equation for the simple circle map (4.1) to account for the linear noise ramp (4.4).



Figure 6.6: (a) Similar phase-plane plot as in Figure 6.1 with b = 8, $\mu = 0.56$, $K_1 = 2^{-6}$, $K_2 = 2^{-5}$, $\{\Phi(0), u(0)\} = \{0, 0\}$, and with additive noise variance σ_N^2 equal to 10^{-5} , with noise-affected region shown. (b) Equivalent plot for first-order case.

Without noise the behaviour of the map (6.2) was simple: to the left of S_1 the map moved right, by an amount $\frac{2\pi}{2^b} \operatorname{Frac}[\mu]$ in Φ , and to its right it moved left by an amount $\frac{2\pi}{2^b} (1 - \operatorname{Frac}[\mu])$. With the noise added, there is a region of uncertainty, just as we had in the first-order case, where the motion is altered; if the map is within the noise-affected region to the left of S_1 , it may occasionally move left again by $\frac{2\pi}{2^b} (1 - \operatorname{Frac}[\mu])$ rather than moving right as per usual. Similarly, if the map is to the right of S_1 , it may move right again. The result is that the range of the circle map in Φ is extended to each side of L_1 and R_1 . The points furthest to the left and right are least densely distributed as these correspond to jumps from the edges of the noise corridor, where the probability of an 'erroneous' jump is least (c.f. noise ramp in Figure 4.2(a)). As a firstorder approximation, the size of the extension is L on each side. The roughly equal distribution to each side is a result of the fact that our μ in this case (14/25) is close to 0.5. This is in contrast to our earlier case of a first-order delayed loop (Figure 5.5), where the noise-induced range extension was highly asymmetric about the original circle map.

For comparison, a representation of the equivalent first-order DPLL is shown in Figure 6.6(b). Here, all parameters relevant to the first-order system are identical to those of the second, and the additive noise level is again $\sigma_N^2 = 10^{-5}$. Although there is just the single state variable, Φ , the steady-state distribution is displayed on a notional phase plane for comparison with the second-order case, simply by extending the plot in the vertical direction. Shown on the plot are the lines corresponding to the circle map bounds we had in Section 3.2.4 for the first-order DPLL and circle map:

$$\Phi = L_1 = M = \arcsin \frac{\operatorname{Int}[\mu] + 1}{2^b K} + \frac{2\pi}{2^b} (\operatorname{Frac}[\mu] - 1),$$

$$\Phi = R_1 = (M + S) = L_1 + \frac{2\pi}{2^b},$$

$$\Phi = S_1 = T = (M + S - \alpha) = \arcsin \frac{\operatorname{Int}[\mu] + 1}{2^b K}.$$
(6.3)

This is to say, our previous lower and upper bounds on the circle map

for the first-order loop, M and (M + S), correspond to our L_1 and R_1 curves in the second-order case. Also, our previously unnamed point at the centre of the noise ramp for the circle map, $T = (M + S - \alpha)$, corresponds to the switching curve of the second-order loop. Just as in the second-order case, the noise-affected region is also shown for the first-order loop; here, it is a region of width exactly 2L, corresponding to the linear noise ramp of Figure 4.2(a). The steady-state distribution of phase error of the first-order loop is shown along a horizontal axis in the centre of Figure 6.6(b). The distribution is similar in many ways to the second-order case of Figure 6.6(a): the points are centred about S_1 , in the absence of noise are fully contained within the circle map bounds, L_1 - R_1 , and in the presence of noise are sent outside this range by an amount approximately equal to L on each side, the distribution thinning further out from the original L_1 - R_1 region.

The obvious difference between the two cases is that, for the first-order loop, whether or not the map is within the noisy region is determined by Φ alone, while for the second-order loop it also depends on u. For the first-order case, the probability of a jump in either direction depends on whether Φ is currently within the region $[S_1 - L, S_1 + L] = [T - L, T + L],$ and where within that region it is. For the second-order loop, it is easy to show that, near the steady-state region, the curves L_1 , S_1 , R_1 , and the bounds on the noisy region, are all approximately linear with slope $K_1(K_2-1)$. Because the bounds on the noise region are no longer vertical lines, it means that they may span several values of Φ within the steadystate region, as can be seen in Figure 6.6(a). Here, the noise region boundaries span approximately 4 points of the period-25 circle map in Φ , so that now, whether a particular point on the map is within the noise region, or what its probability of an 'erroneous' jump is, depends on what the value of u is at that time. However, because the width of the noise region is almost constant, the number of points on the period-25 map in Φ that are within the noise region does not vary with u. The overall effect of this is that the number of points thrown outside the L_1 - R_1 corridor is similar to that in the first-order case, with a distribution that is similar.

6.3.3 Future work

For large-denominator or irrational input frequencies, μ , the overall output jitter variance follows a curve almost indistinguishable from that in Figure 6.2. However, the details of how the steady-state dynamics are affected by low- and medium-level additive noise differs significantly from what we have seen here for the low-denominator rational frequency. In general, the increase in range of both the jitter, Φ , and integrator output, u, occurs more gradually as the additive noise increases. This is to be expected as the horizontal circle map is now densely distributed between L_1 and R_1 and so points on this map will get moved across the switching curve for arbitrarily small levels of input noise. Therefore, we do not see the orbit jumps we had for the period-25 case. However, these differences in dynamics are not significant enough in magnitude to be seen on a logarithmic variance plot.

We believe there is significant scope for further research into the case of additive noise in discrete second-order loops, for low- and highdenominator rational, irrational, and integer input frequencies.

6.4 Summary

We have first presented a recap of the background material in Section 3.2.5 on the second-order loop in the case of no external additive noise. Proceeding to our new results, we have shown how, on a macro level, the response of the second-order loop to additive input noise is remarkably similar to that of the first-order system; the output variance of the second-order loop generally tracks that of the first, with only minor deviations. Finally, in Section 6.3 we use numerical simulations of the second-order DPLL to look in detail at how the period-25 limit cycle is disrupted by additive input noise. We see that, for a large range of noise levels through our 'intermediate' region, while the pattern of the motion in u is lost quite easily, the structure of the orbit for Φ remains quite simple and largely confined to the original grid. As the noise level

increases, additional grid points are added to each end of the map. Only for large noise levels, as we approach the analogue limit, does the simple, quantized structure break down completely. Finally, we compare the steady-state behaviour of first- and second-order loops within the intermediate region. For sample loop parameters of interest and a particular input noise level, we look at the behaviour of the phase error, Φ in each. By visualising the first-order system on a notional 2-dimensional phase plane, we see how the similarities in the two systems come about.

Chapter 7

Conclusions and future work

In this thesis we devoted two chapters to an in-depth review of the mathematical theories required for an understanding of our later work. In Chapter 2 we covered basic probability theory, random variables, distribution and density functions, moments and expectations. We extended the presentation to include multivariate distributions and stochastic processes. In our summary of the latter, a vast area of study in its own right, we focused particularly on the concepts of autocorrelation and power spectral density, as well as Markov processes and the Chapman-Kolmogorov equation. In our final section of base material, we provided a quick introduction to analogue and digital filters, which are central to the PLLs and DPLLs of the following chapter.

In Chapter 3 we gave an extensive overview of PLL theory, treating the traditional analogue and more modern digital theories separately. The background section on the analogue PLL relies heavily on the seminal work by Viterbi [1], covering both linearized and non-linear PLLs in the absence of noise and also with additive input noise. In particular, the case of additive input noise to the non-linear loop uses much of the theory of stochastic processes covered in the previous chapter. For the summary of the digital loop, we referenced the work of Gardner [2], one of the noted experts in this field. We looked at the difference equations describing each of the loop components and derived the full loop equations for first- and second-order loops in the case of frequency quantization. We saw how, in steady state and under certain conditions, the loop obeys a circle rotation map. This map was reviewed and some of its properties examined, in particular its noise-free variance and we showed how this can be used to predict the phase-jitter variance of the DPLL output due to quantization alone.

Chapter 4 marks the beginning of the presentation of the novel aspects of our work. In Chapter 4 itself we focused on the original case considered by Gardner [3], namely the first-order, frequency-quantized DPLL without loop delay in the presence of additive input noise. It was shown that under certain conditions, the steady-state dynamics of the first-order DPLL are identical to the circle rotation map. In addition it was illustrated how additive input noise in the case of the first-order DPLL corresponds to a unique form of quantization noise in the case of the circle map.

An equation for the time-dependent PDF of the noisy circle map's phase jitter was obtained, and, using this equation, bounds and asymptotic approximations for the variance of the phase jitter were found. A non-local equation was found for the steady-state PDF, and various aspects of its solution were examined using asymptotic methods. In particular, the PDF itself, and associated mean and variance of the noisy circle map output were found, the result being to accurately describe the behaviour in the large-noise, small- α limit of the circle map, and thus also characterize the intermediate regime response of the first-order DPLL. The asymptotic results of Section 4.4 — in particular equation (4.36) for the large-L limit and equations (4.37)-(4.39) for the small- α case are expected to be of significant benefit to DPLL designers, as they constitute the first analytical results describing the jitter in this important operational regime where additive and quantization noise interact nontrivially.

In Chapter 5 we looked at the first-order DPLL with loop delay, previously introduced by Gardner in [3]. This was also examined and analyzed as a delayed circle map. A unique — and previously unexplained — feature of the delayed case, a dip in output variance as a response to additive input noise similar to a stochastic resonance, was highlighted. A full analysis from first principles was performed for the case where the dip is most prominent, i.e., where α , representing the fractional part of the DPLL input frequency, is small. From this small- α approximation, the resulting PDF, mean and variance found to accurately match numerical simulations of the DPLL equations.

Finally, in Chapter 6, the second-order DPLL with noisy input was examined. While the dynamics of the map in steady state are considerably more complicated than for the first-order loop, the magnitude response of the output phase jitter to additive input noise was seen to be very similar to that of the first-order DPLL. The complexity of the dynamics under noisy conditions was shown on the phase plane and should serve as motivation for further study to students of non-linear dynamics.

Bibliography

- A. J. Viterbi, *Principles of Coherent Communication*, New York: McGraw-Hill, 1966.
- [2] F. M. Gardner, *Phaselock techniques*, 3rd ed., New York: Wiley, 2005.
- [3] F. M. Gardner, "Frequency granularity in digital phase-locked loops", IEEE Trans. Commun., vol. 44, pp. 749-758, June 1996.
- [4] A. Teplinsky, O. Feely and A. Rogers, "Phase-jitter dynamics of digital phaselocked loops", IEEE Trans. Circuits and Systems, Part I: Fundamental Theory and Applications, 46, pp. 545-558, May 1999.
- [5] A. Teplinsky and O. Feely, "Phase-jitter dynamics of digital phaselocked loops: Part II", IEEE Trans. Circuits and Systems, Part I: Fundamental Theory and Applications., 47, pp. 458-473, April 2000.
- [6] C. Ó Tuama, J. P. Gleeson, "Combined Effects of Frequency Quantization and Additive Input Noise in a First-order Digital PLL", Electronic Notes in Theoretical Computer Science, vol. 225, pp. 255-268, January 2009.
- [7] R. E. Best, *Phase-locked Loops: Design, Simulation and Applications*, 6th ed., New York: McGraw-Hill, 2007.
- [8] R. N. McDonough and A. D. Whalen, *Detection of Signals in Noise*, 2nd ed., San Diego: Academic Press, 1995.
- [9] T. M. Apostol, Calculus, Vol. 2, 2nd ed., New York: Wiley, 1969.

- [10] W. Rudin, Real and complex analysis, 3rd ed., New York: McGraw-Hill, 1987.
- [11] M. Abramowitz and I. A. Stegun, Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables, vol. 55, New York: Dover, 1964.
- [12] P. A. M. Dirac, The principles of quantum mechanics, 4th ed., Clarendon, 1958.
- [13] C. Gardiner, Handbook of Stochastic Methods, 3rd ed., Berlin: Springer-Verlag, 2004.
- [14] B. P. Lathi and Z. Ding, Modern Digital and Analog Communication Systems, 4th ed., New York: Oxford University Press, 2009.
- [15] R. G. Bartle, The elements of integration and Lebesgue measure, New York: Wiley, 1995.
- [16] M. E. Newman and M. Girvan, "Finding and evaluating community structure in networks", Phys. Rev. E, vol. 69.2: 026113, 2004.
- [17] A. N. Kolmogorov, "On analytical methods in the theory of probability", Math. Ann 104, pp 415-458, 1931.
- [18] H. Risken, The Fokker-Planck Equation, 2nd ed., Berlin: Springer-Verlag, 1989.
- [19] N. G. Van Kampen, Stochastic Processes in Physics and Chemistry, 3rd ed., Amsterdam: Elsevier, 2007.
- [20] L. Arnold, Stochastic Differential Equations: Theory and applications, 1st ed., New York: Wiley, 1974.
- [21] B. Øksendal, Stochastic Differential Equations, 6th ed., Berlin: Springer-Verlag, 2003.
- [22] W. Horsthemke and R. Lefever, Noise-induced Transitions, Berlin: Springer-Verlag, 1984.

- [23] A. Williams and F. Taylor, *Electronic Filter Design Handbook*, 4th ed., New York: McGraw-Hill, 2006.
- [24] A. Antoniou, Digital Filters: Analysis, Design, and Applications, 2nd ed., New York: McGraw-Hill, 2000.
- [25] W. Egan, *Phase-Lock Basics*, New York: Wiley, 1998.
- [26] J. G. Truxal and R. Brown, Automatic feedback control system synthesis, New York: McGraw-Hill, 1955.
- [27] S. Strogatz, Nonlinear dynamics and chaos: with applications to physics, biology, chemistry and engineering, Reading: Perseus, 1994.
- [28] R. L. Devaney, An Introduction to Chaotic Dynamical Systems, 2nd ed., Boulder: Westview, 2003.
- [29] C. Mira et al., Chaotic Dynamics in Two-Dimensional Noninvertible Maps, Singapore: World Scientific, 1998.
- [30] N. Da Dalt, "Markov chains-based derivation of the phase detector gain in bang-bang PLLs", IEEE Trans. Circuits and Systems, Part II: Exp. Briefs, vol. 53, no. 11, pp. 15, November 2006.
- [31] B. Chun and M. P. Kennedy, "Statistical Properties of First-Order Bang-Bang PLL With Nonzero Loop Delay", IEEE Trans. Circuits and Systems, Part II: Exp. Briefs, vol. 55, no. 10, pp. 1016-1020, October 2008.
- [32] C. M. Bender and S. A. Orszag, Advanced Mathematical Methods for Scientists and Engineers, New York: Springer-Verlag, 1999, ch. 6, pp. 263-267.
- [33] K. C. Pohlmann, Principles of Digital Audio, 6th ed., New York: McGraw-Hill, 2010.
- [34] R. Benzi et al., "The mechanism of stochastic resonance", J. Phys. A 14, L453, 1981.

[35] A. R. Bulsara and L. Gammaitoni. "Tuning in to noise", Physics Today, vol. 49, p 39, 1996.