INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI®

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

On-Chip Characterization of Charge-Pump Phase-Locked Loops

Benoît R. Veillette,

(B.Eng 1993, M.Eng 1995)

Department of Electrical Engineering

McGill University, Montréal

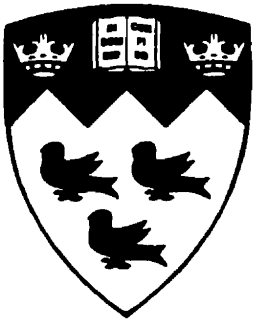

March 1998

A Thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Arriver, c'est rien

t'occupe pas de ceux qui

se lamentationnent:

on arrive pas...

on arrive jamais à arriver...

Ça compte pas.

Partir! C'est ça qui compte.

D'ailleurs pour croire qu'on peut arriver

faut déjà être complètement parti!

Je m'égalomane à moi-même, Marc Favreau (Sol)

Abstract

An all-digital technique for the measurement of the jitter transfer function of charge-pump phase-locked loops is introduced. Input jitter may be generated using one of three methods. They allow trade-offs between test clock frequency, hardware requirements and accuracy. All the methods rely on delta-sigma modulation to shape the unavoidable jitter quantization noise at high frequencies. This jitter noise is filtered by the lowpass characteristic of the device-under-test and has a minimal impact on the test results. For response measurement, the phase of the output signal is compared against a jitter threshold. As the stimulus generation and output analysis circuits are digital and do not require calibration, this jitter transfer function measurement scheme is highly amenable to built-in self-test. The technique can also be used to adaptively tune a PLL after fabrication. The validity of the scheme was verified experimentally with off-the-shelf components.

Résumé

Une technique digitale pour mesurer la réponse de fréquence à la gigue de circuits de verrouillage de phase utilisant des pompes à charge est présentée. Le signal de phase servant de stimulus peut être généré à l'aide de trois méthodes. Elles permettent des compromis entre la vitesse d'horloge, le matériel requis et la précision. Toutes utilisent la modulation delta-sigma pour façonner l'inévitable erreur de quantization de telle sorte qu'elle apparaisse surtout dans les hautes fréquences. Cette erreur est filtrée en majeure partie puisque le circuit vérifié à une réponse de fréquence passe-bas. L'impact sur les résultats du test est donc réduit. Pour mesurer la réponse du circuit, la phase du signal de sortie est comparée à un seuil de détection. Puisque la génération du signal d'entrée et l'analyse du signal de sortie sont faits dans le domaine numérique et ainsi ne nécessitent pas de calibration, cette technique de mesure de la réponse de fréquence à la phase peut-être employée pour l'auto-vérification intégrée. Cette technique peut aussi être utilisée pour ajuster un circuit de verrouillage de phase après fabrication. La validité de la méthode est démontrée à l'aide d'expériences faisant appel à des composantes commerciales.

Acknowledgements

I would first like to thank my supervisor Gordon W. Roberts. Not only did he provide direction and advice for my research work but he also gave me the chance to travel to numerous conferences where I had the opportunity to compare my work with others. I am convinced the last five years under his guidance will be very precious as I ready myself to enter the *real* world.

I must thank the Natural Sciences and Engineering Research Council (NSERC) and the Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (Fonds FCAR) for supporting me financially, making sure I only worry about electrical engineering. The Canadian Microelectronics Corporation (CMC) is also acknowledge for supplying a large share of the equipment used for this work. Jacek Slaboszewicz and Thomas Obenaus deserve praises for maintaining a suitable computing work environment with limited resources. I also wish to acknowledge the moral support from fellow graduate students of the MACS laboratory. I am further grateful to Brian Gerson of PMC Sierra for suggesting alternative methods for testing phase-locked loops.

Je m'en voudrait de ne pas mentionner le support de mes amis de Québec et Montréal. Leur respect pour mon mode de vie étudiant m'a permis de questionner les dogmes scientifiques plutôt que de remettre en cause mon choix de carrière. Leur présence dans les bons moments comme dans les moments plus difficiles fut un ingrédient essentiel de ce travail. La recherche scientifique, pour être fructueuse, doit être accompagnée de soccer, de WiF, de discussions dans un café et de pic-nics dans une voiture sous la pluie.

Finalement, j'aimerais remercier sincèrement mes parents. Sans leur support financier et leurs encouragements constants, je n'aurais pu me rendre aussi loin. Mais je leur suis particulièrement reconnaissant pour la curiosité et l'éthique de travail qu'ils ont su m'inculquer.

Claim To Originality

The core of the original work on the measurement of phase-locked loop characteristics is located in Chapters 4 and 5. Appendix A, Section A.4 introduces a new circuit required for the above purpose while Appendix B demonstrates an application.

- Chapter 4 introduces three methods for stimulating phase-locked loops for test purposes. Unlike prior art, these methods are almost entirely digital. Digital phase modulation using delta-sigma modulation, presented in Section 4.1.3, allows the creation of bandlimited phase jitter using a test clock frequency which is a small multiple of the phase-locked loop operating frequency. The use of delta-sigma modulation in the time-domain for digital-to-analog conversion is a first. The second method, presented in Section 4.2, has been presented before for communications purposes. Nevertheless, its use as a stimulus method has never been proposed before. Also, in this section appears the first report of the non-linearity inherent to this method. Then loop injection using a charge pump, another novel method, is described in Section 4.3. While it allows the use of a lower test clock frequency, it requires an additional charge pump in the phase-locked loop. Finally, an analysis of the effects of delta-sigma modulation on internal nodes of bandlimited circuits addresses an important concern of designers.

- Chapter 5 explains a method for evaluating jitter at the output. The jitter threshold circuits of Section 5.4 are adapted from a standard phase-locked loop test. Using this and one of the three stimulus method, the first comprehensive method for measuring the jitter transfer function of charge-pump phase-locked loops on-chip is introduced. Engineering trade-offs between the various implementations are discussed. Furthermore, the accuracy of the measurements is quantified.

- In Appendix A, Section A.4, a new type of delta-sigma modulator which has been labelled slope-limited is introduced. This device is used in a digital phase modulation scheme but could find its way in other applications.

- Appendix B presents a method for generating bandpass type analog signals using a phase-locked loop tested as described in Chapters 4 and 5. This application could prove useful in implementing BIST of wireless communications products.

Contents

List of Tables

List of Figures

Chapter 1 : Introduction

Every manufactured product, from cars to computers, must be verified for correct operation before being shipped to buyers. Customers expect working products and therefore a bad part finding its way to the retail market will cost dearly both in repair costs and in damaged reputation. A well known illustration of this fact is in the auto market where Japanese car makers benefit from a good reputation for reliability and may thus set higher retail prices than their American counterparts [1]. In that respect, integrated circuits (ICs) are no different than any other manufactured good. However, in contrast to larger products, a sound fabrication process does not guarantee defect-free ICs at the end of the manufacturing line. Indeed, the extremely small size of the features, on the order of the diameter of a human hair, makes the fabrication process very sensitive to the environment. At critical times, a dust particle or minute vibrations are sufficient to ruin a device. As these conditions may only be controlled to a certain extent in the factory clean rooms, the electronic industry is accustomed to the idea that a portion of any production run will be defective. However, a method must exist to screen bad parts before they are shipped to costumers. This manufacturing step, called IC testing, is an important operation of the fabrication process and can account for a significant portion of the final cost of an IC.

1.1 Economics of Integrated Circuits Testing

Contributions from a large number of factors, some of which are very difficult to quantize enter into the test cost of an IC [2]. These factors may be divided into two groups with respect to their recurring nature. In the first group are the non-recurring test costs which generally occur before full-scale production begins. They are charged once and are constant regardless of the number of devices eventually fabricated. Non-recurring engineering costs for test include the time spent by designers to make an IC easier to test, an activity which

is labelled design for test (DFT). Also a significant responsibility for engineers is the planning of the test procedures after the design has been signed off. To these manpower expenses must be added the cost to operate the resources used in these tasks such as computers for simulations and testers for test debugging. Finally, a factor which is rather difficult to evaluate is the time-to-market delay incurred because of the development of the tests. While it is well known that delays with respect to competing products will result in lost revenue, it is hard to come up with a model to quantify this statement. On the other side are the costs associated with individual manufactured unit. The first contributor is the extra silicon area required for the DFT circuits. However, the largest portion of the test costs is generally the actual test operation. The price per chip is the test time of a device multiplied by the operating cost per unit time of the tester. Figure 1.1 illustrates the different contributions to the test cost of an IC.

Testers have now become the largest capital spending items of semiconductor manufacturers. This equipment manipulates wafers or packaged ICs, applies the test vectors and mon-

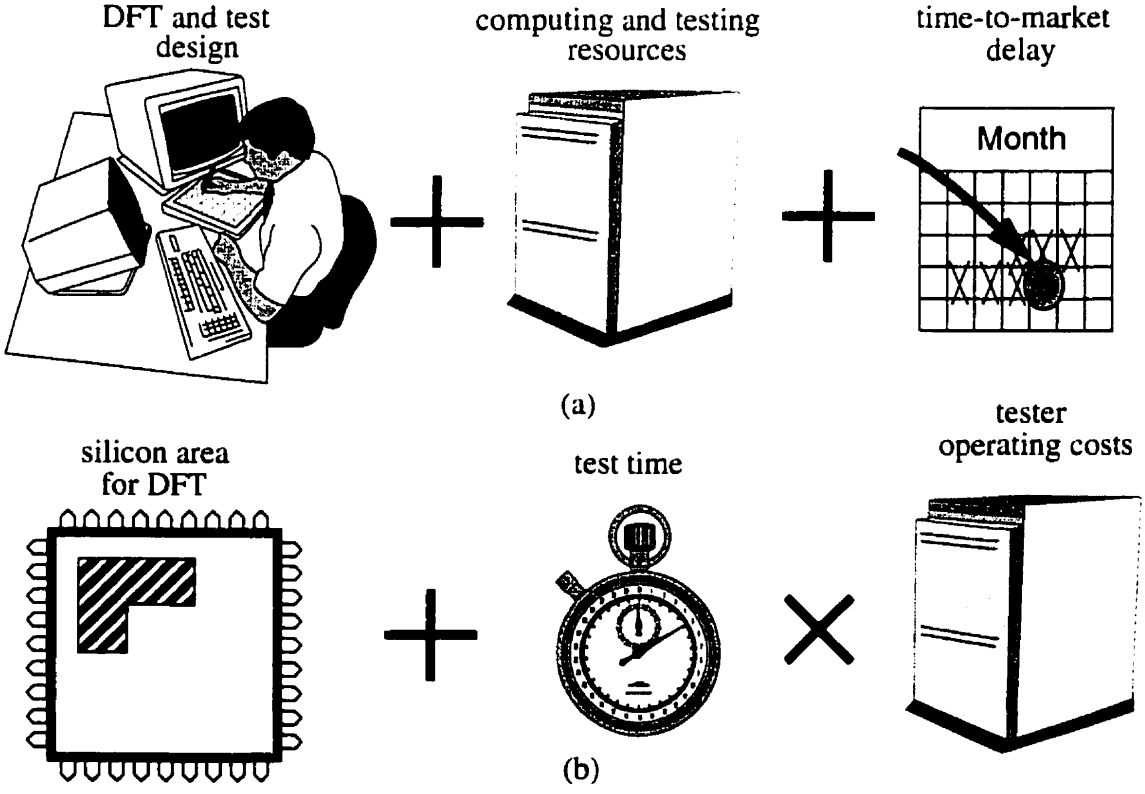

Figure 1.1: Cost of test: (a) Non-recurring costs. (b) Per manufactured unit.

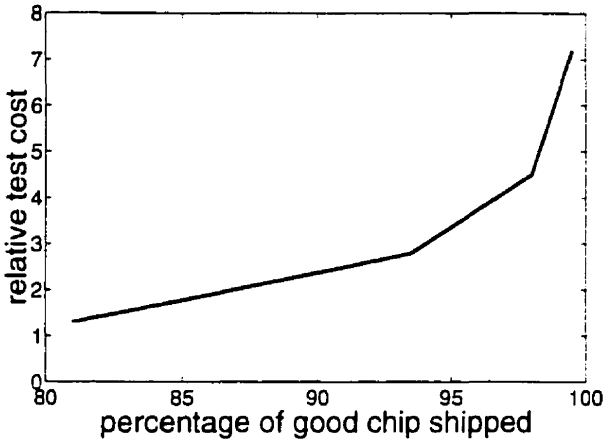

Figure 1.2: Relative test cost vs. percentage of good chips shipped [4].

itors the responses. Because of increased circuit performance, the cost of high-end automated test equipment is expected to reach ten million dollars half-way into the next decade [3]. Obviously, the test time of a single IC is an important parameter. However, it should not be considered in isolation from the other factors. Any test solution must address the whole test cost equation.

Another concept, fault coverage, is not represented in the above analysis but is pivotal to the economics of IC testing. It is defined as the number of faulty parts identified by the test procedure over the total number of defective devices. Obviously, it should be as close to 100% as possible. Undeniably, faulty parts not identified by the tester will result in repair or return cost. Nevertheless, silicon manufacturers usually settle for a fault coverage between 90% and 99.9% for a given test. This is because uncovering the remaining faults requires an inordinate amount of effort, both in test development and in test time as illustrated in Figure 1.2. The marginal cost of uncovering these hard-to-test faulty devices becomes larger than the penalty for letting them through.

1.2 Built-In Self-Test

With shrinking manufacturing costs and increasing performances, the significance of integrated circuit testing has ballooned. This issue has reached the critical level for a class of

device referred to as mixed-signal ICs. These devices are usually composed mostly of digital circuits which perform signal processing. However, to interface these computing engines to the physical world, analog circuits are required. Common examples of this type of circuit are analog-to-digital converters, digital-to-analog converters, phase-locked loops and sensors. While these circuits typically occupy less than 10% of the total silicon area of a mixed-signal IC, their test cost is approaching half the total cost of the product [5].

A number of solutions have been proposed by researchers, most of them adaptations of digital testing ideas. In fact, digital test methods have succeeded so far in coping with the exponential growth in circuit size. The main reason for this success is the integration of these test solutions in computer-aided design (CAD) software. In contrast, analog test rely mostly on ad-hoc solutions. Yet, it should be mentioned that analog test CAD tools have emerged recently but they fall short of industry's expectations.

A DFT concept which offers great promises is built-in self-test (BIST) [6]. Though the expression is used for widely different schemes, it generally implies the inclusion of some or

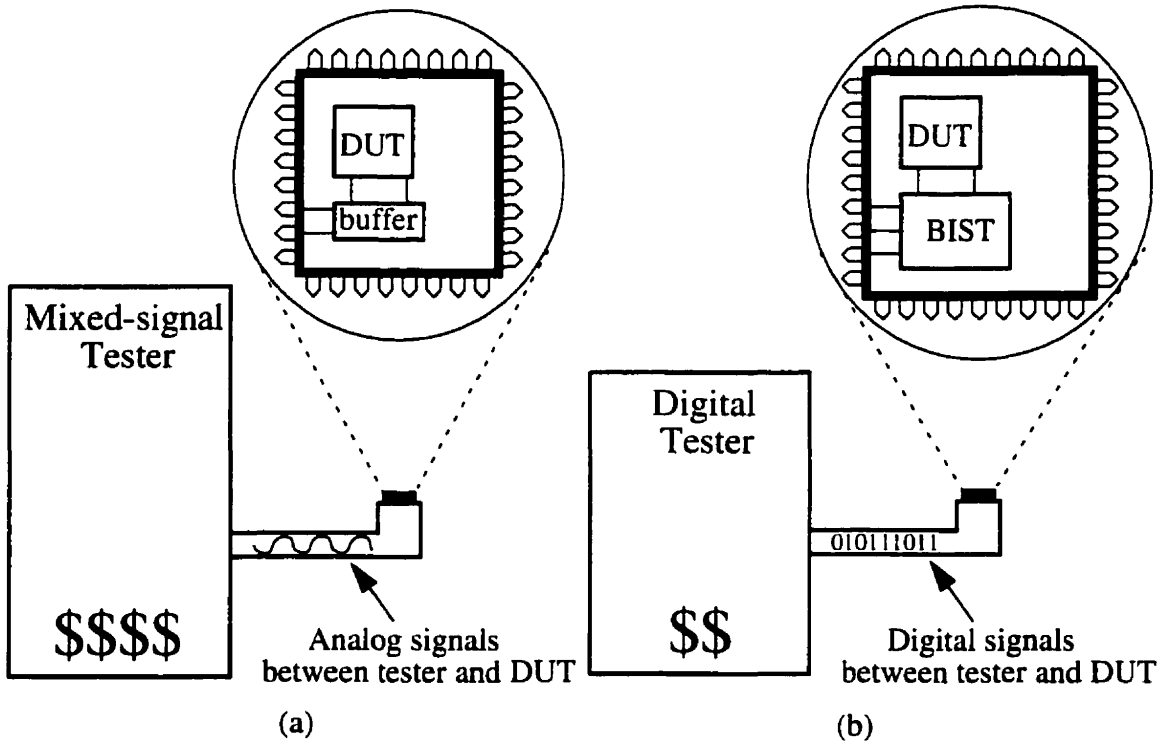

Figure 1.3: (a) Traditional mixed-signal testing. (b) Built-in self-test.

all of the test functions on the same die as the device-under-test (DUT). Test can thus be repeatedly performed throughout the manufacturing process without requiring complex testers. Furthermore, as the output analysis circuitry is on the die, BIST also alleviates the need for complex buffers to drive signals off-chip. This feature is most important for high speed analog signals. Yet, this solution is not free as an IC which is BIST capable will usually require more silicon area for implementation. However the tester needs no longer to be capable of generating analog signals. Significant savings in capital costs are thus expected. Figure 1.3 illustrates these differences between traditional mixed-signal testing and an ideal BIST scheme.

Placing analog signal sources and measuring instruments on-chip is not an easy task. Indeed, the integrated circuits economics dictate a minimal silicon area overhead to prevent the gains on the tester from being written off by a more expensive fabrication cost incurred for BIST. Furthermore, as these circuits will be used to characterize others, their functionality must be guaranteed. Again, to prevent overhead costs from rendering a BIST solution unattractive, they should involve little or no calibration. Fortunately, digital circuits address this latest concern. In fact, the functionality of digital circuits can be tested easily and they exhibit excellent immunity to noise sources. Moreover they are more flexible than analog circuits as they can be handily programmed or reconfigured. While analog sources or instruments may not be built out of digital circuits only, most of the signal processing can be handled in the digital domain and consequently, the dependency on analog component reduced. Therefore, digital circuits will be preferred over analog circuits wherever possible for circuit blocks implementing BIST.

1.3 Phase-Locked Loop Testing

The phase-locked loop (PLL) is a mixed-signal block used in a large number of applications such as frequency synthesis, phase demodulation, clock distribution and timing recovery [7]. As such, it is essential for systems like wireless phones, optical fiber links and micro-computers. PLLs are even expected to appear soon on other purely digital circuits such as field-programmable gate arrays (FPGAs) and digital signal processors (DSPs).

PLL-based circuits are difficult to characterize because the parameters assessing their performance may not be extracted easily from the signals. Precision analog signal sources and measuring instruments are necessary [8]. This problem is recurrent to the characterization or testing of almost any mixed-signal or analog circuit. However, as a PLL is the only analog block on many mixed-signal chip, its test requirements directly translate into inflated product cost. Looking for solutions, test engineers are tempted to break the task and verify sub-blocks independently. This strategy has been successfully applied to digital testing problems. On the other hand, analog circuits are much more complex and do not lend themselves very well to this methodology [9]. Furthermore, the nature of the phase-locked loop renders this solution unattractive as the tight feedback makes it difficult to relate the specifications from the PLL level to the block level. Test development is further hampered by the computational power required by even simple simulations of the PLL. This is due to the order of magnitude spread in the time constants of the device. Behavioral simulation could reduce the computing burden but, because of the feedback around these blocks, complex models are required [10]. There is thus a serious need to address the issue of PLL testing from a different angle.

The work herein is limited in scope to the class of PLL operating on digital signals. The phase information of these signals is contained in the transition times. While these transitions are asynchronous, the signals are quantized to two levels, i.e. binary. These PLLs usually make use of sequential phase detectors which require charge pumps and are thus called charge-pump PLLs [11].

1.4 Dissertation Goals

This thesis will investigate techniques to integrate some of the PLL test functions on-chip, thereby reducing the requirements on the tester. In following the philosophy developed in the mixed-signal test group at McGill, the dependency of the testing scheme on analog components shall be reduced to a minimum. As a matter of fact, digital circuits are more reliable and benefit from established design and test methodologies. In this thesis we explore the verification of PLL characteristics using mostly synchronous digital circuits. It

implies that the stimulus signals can only change at the clock edges and that the output signals may only be sampled at the same clock edges. This would seem like an unbearable constraint as jitter, both created and measured, is quantized to the test clock period. The results of any test would thus be severely limited in precision. Yet, it will be shown that this hurdle can be overcome using a technique called delta-sigma ($\Delta\Sigma$) modulation. The ultimate aim is to propose a complete BIST scheme for PLLs integrating both a signal generator and the output analysis function alongside the DUT. Devices compliant to this method would be testable with a simple and inexpensive digital tester.

1.5 Small Dissertation Outline

Chapter 2 will explain the problem of IC testing. This will help locate the proposed scheme in the spectrum of solutions introduced in recent years. The phase-locked loop will be the topic of Chapter 3. There, the notation in use in the core chapters will be explained. Various specifications which apply to PLLs will be described. Also, the literature will be reviewed for testing solutions targeting specifically PLLs. Chapter 4 will examine stimulus generation. After discussing solutions ported from analog instruments, digital methods similar to those employed in current manufacturing tests will be investigated. Finally novel solutions suitable for BIST will be proposed. Then, Chapter 5 will explain output analysis and test methodology. Again, a number of methods will be examined, leading to a very simple solution. Experiments are described in Chapter 6 where results are displayed. A brief summary and a look a future research will compose Chapter 7. Appendix A will explain the key technique of delta-sigma modulation and will take a look at the problem of generating digital signals on-chip. A novel circuit, the slope-limited delta-sigma modulator will also be described in this Appendix. Finally, the use of the phase modulation and PLL test techniques for the generation of bandpass analog signals will be discussed in Appendix B.

Chapter 2 : Mixed-Signal Testing

While testing of digital circuits has matured to the point where most of the engineering work is performed by CAD tools, analog and mixed-signal testing solutions are mostly ad hoc. To understand the reasons for this lag, it is necessary to take a step back and look at the global issues of integrated circuits testing. First the distinction between design validation, or characterization, and production test must be clearly established. Following this, functional testing, a manufacturing test strategy which insures all the specifications are satisfied, will be presented. However its weakness lies in the test duration which sometimes make this solution inapplicable. Thus two methods for reducing the test time, test set ordering and inductive fault analysis, will be introduced. Finally, an introduction to analog and mixed-signal BIST for measuring specifications will be provided. It should be noted that improving controllability and observability of internal nodes, a significant topic of mixed-signal testing, will not be covered here as it is not relevant to this dissertation. A section in the following chapter will specifically address PLL testing.

2.1 Design Validation Versus Production Test

Figure 2.1 illustrates a typical integrated circuit design flow. From specifications, a design is created in schematic form and verified using a circuit simulator such as Spice. This design is modified a number of times until the simulations show it complies to the specifications. However, computer simulations rely on models which can only represent reality to a certain extent. Furthermore, as circuit simulation can be computationally intensive, only a subset of the specifications are verified. For example, transient analysis is performed for a small number of frequencies which are thought to capture the important aspects of the device behavior. It is thus necessary to fabricate a small number of devices before full-scale production. Functionality must then be fully verified to insure that the models used for the

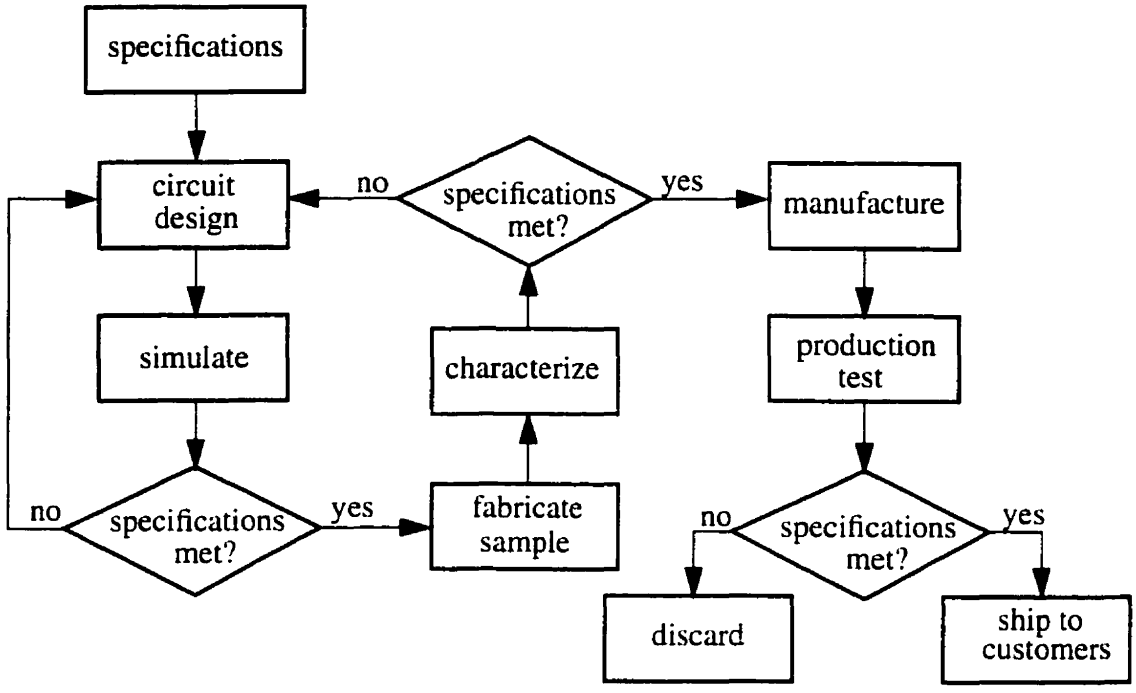

Figure 2.1: Typical integrated circuit design flow.

design were accurate and that the extrapolations which were not verified through simulations are correct. This operation is called design validation or characterization. It requires accurate digital and analog instruments such as signal generators, oscilloscopes or spectrum analyzers.

Once validated the IC may be volume manufactured. However, even after thorough design characterization, the fabricated devices may fail to meet specifications. A number of problems inherent to the semiconductor manufacturing process can occur. A dust particle for example could prevent the deposition of a metal trace and cause an open circuit. Alternatively, bad process control could result in a thin gate oxide causing transistors to break down at abnormally low voltages. Also random variations in the distance between two polysilicon layers could result in intolerable mismatch between two capacitors. These issues are unavoidable and can only be reduced but not eliminated by semiconductor process improvement. Regardless of their origin, the above mentioned problems are labelled defects. However, it should be noted that defects do not necessarily result in a faulty circuit. In fact, a goal of IC designers is to desensitize their design to process variations and defects. This

can be performed by leaving head room between the specifications and actual circuit performance, by employing circuit techniques such as negative feedback and by adding redundancy. Furthermore, multiple faults may cancel each others and have a reduced impact on performance. Hence searching for defects is sufficient for the purpose of testing but it is not the optimal strategy. It should be clear that the purpose of production test is to screen devices that no longer meet specifications because of defects, not to find defects per se.

Clearly, design validation and production test have different goals and different constraints. In design validation, all the specifications must be verified and, in case of failure, the location of the problem should be revealed. On the other hand, the cost of production test can be a significant part of the final device cost. The aim is therefore to minimize the total test cost while identifying most if not all defective parts. Verifying all specifications may not be necessary as a defect causing a circuit failure will likely affect many metrics and thus can be uncovered with any one of a number of measurements. This idea will be explained in more details in the following section.

2.2 Production Test Strategies

Various production test strategies may be called upon to reduce the test cost of an IC. As explained in Section 1.1, a large number of factors enter the test cost equation. Test time has traditionally received more attention as it impacts the number of expensive testers that must be bought by a manufacturer. However, the test cost equation should be studied globally.

2.2.1 Functional Testing

As the first goal of production testing is to insure product quality, the basic strategy, referred to as functional testing is to verify the compliance of the device to all specifications. The device under test (DUT) is treated as a black box as only its input-output characteristics are of interest. This procedure is similar to device characterization except that the cause or origin of an error is usually not meaningful. Functional verification is generally the industry's choice for analog and mixed-signal testing.

The main advantage of complete functional verification is that the fault coverage is by definition 100%. Also, test development is relatively straightforward. However, this type of test requires complex signal sources and measuring instruments to extract performance metrics such as frequency, power or jitter. Expensive mixed-signal testers are thus necessary. Also, the number of test vectors required to perform the test is usually quite large and the test time may be very long. Consider for example the verification of an integrated service digital network (ISDN) transceiver [12]. A full activation from cold start requires convergence of the echo canceller, convergence of the decision feedback equalizer and recovery of the data clock. The worst case activation time specified by the underlying standard is 15 seconds, clearly an unacceptable test time in a production environment. Furthermore it would require the ATE to generate and analyze complex communication waveforms. Therefore, alternative test strategies must be employed.

2.2.2 Test Set Ordering and Trimming

All the tests necessary to verify compliance to the specifications compose the test set of a DUT. Some of these tests will be failed more often as they are more sensitive to process variations. Since the test cost is directly proportional to the average time an IC spends on a tester, it makes sense to order the tests such that those with low yield are performed first. As soon as a device fails a test, it is dropped and another device begins the test sequence. Furthermore, some of the tests may be highly correlated. Intuitively, a fault (for example a short circuit) which degrades a measure will possibly affect negatively other measures. It is thus possible that a test will never uncover a fault if it is performed after other tests. They are thus dropped from the test set.

The procedure for modifying the test set is depicted on the left side of Figure 2.2. A layout of the DUT is the starting point. A large number of devices are fabricated and tested for all the specifications. Using the results obtained for the yield of the different tests, the test set may be ordered and trimmed. Algorithms for exact ordering [13] and heuristic ordering [14, 15] have been presented.

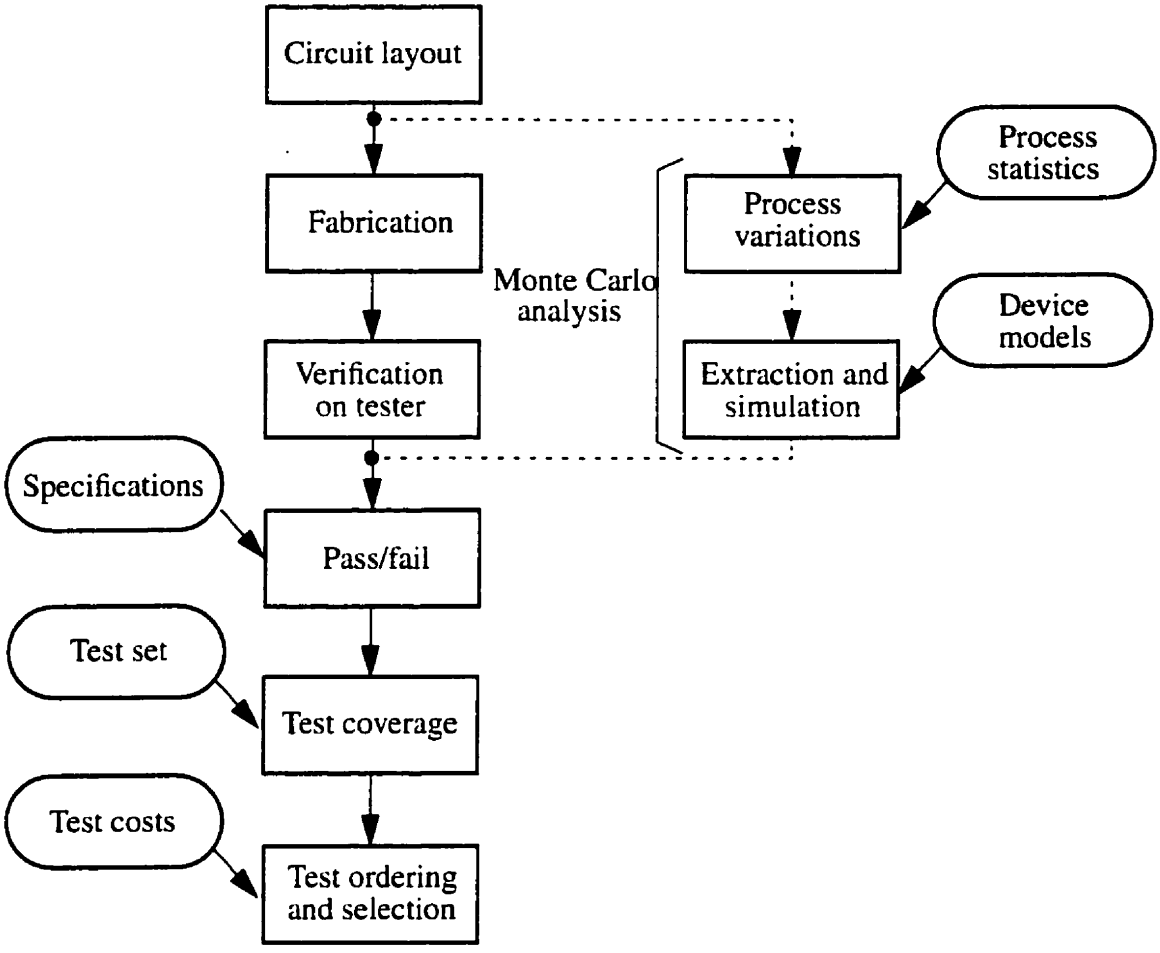

Figure 2.2: Test set trimming and ordering.

2.2.3 Inductive Fault Analysis

The drawback with test set ordering and trimming is that a large number of devices must be tested for all specifications before statistically sufficient data is obtained. The first fabricated devices thus require very long test times. As an IC tester is very expensive, a large sum can be saved by doing the test set ordering and trimming before fabrication starts. The manufacturing process is instead emulated on a computer with defects injected in the layout according to process statistics. The resulting virtual devices are simulated to obtain the parameters relevant to the specifications instead of being placed on a tester as shown on the right side of Figure 2.2. This analysis is performed on a large sample of random process conditions in a Monte Carlo fashion. This procedure has been labelled inductive fault anal-

ysis (IFA) and was first introduced for digital circuits [16] before being extended to analog circuits [17].

The major drawbacks of this technique are that a comprehensive fault model is difficult to obtain and that in maximizing the fault coverage, a large number of simulations are required. Significant research work is going on to reduce the computational requirements, for example by using hierarchical models [10]. However, the only experimental work published to date shows that IFA for analog circuits results in unacceptably high number of escapes (faulty devices that succeeded the tests) because of unmodeled faults [18]. It can be concluded that IFA is not suitable for circuits with high specification complexity where faults from process tolerances play a significant role. Nevertheless, IFA has been successfully implemented in CAD tools for digital circuits where gates show little sensitivity to process variations. Proponents of this technique for mixed-signal circuits expect the ever increasing power of computing machines to allow more complex models and therefore better results.

The introduction of IFA and the concept of fault coverage have led to the appearance of benchmark circuits, first for digital applications [19] and now for analog and mixed-signal devices [20]. These benchmarks however are not meaningful for functional tests where the fault coverage is inherently 100%. As the test solutions in this dissertation are intended to measure specifications, benchmark figures are not appropriate.

2.2.4 Structural Testing

An alternative to the black box approach of functional verification is to make use of the knowledge of the circuit to break down the verification problem into smaller, more manageable test problems. As most of the time the behavior of these sub-circuits may only be loosely related to specifications, IFA must be employed to compare the quality of the structural tests to the functional tests they intend to replace.

The advantage of this approach is that the number of test vectors is usually much smaller and therefore the test time and test costs are reduced. Also the stimulus and response signals may possibly be simpler thus requiring cheaper instruments. It does not come as a surprise

therefore that structural testing is the standard today in digital testing. Although it is still evolving, the fault model is relatively simple and digital simulations can be performed easily. In mixed-signal and analog verification, structural testing has been impeded in part by the excessive computational power required for simulations of faulty circuits but mostly by the difficulty in mapping physical defects to a set of analog faults [9].

2.3 BIST for Analog and Mixed-Signal Circuits

While BIST of memory arrays is widespread and BIST of datapath structures is finding its way into actual products, no commercial part yet integrates BIST for a mixed-signal circuit. This lag may be attributed to the large spectrum of analog circuits and specifications and also to the difficulty of performing reliable analog measurements on-chip. A number of schemes has nevertheless been proposed by researchers. However, most of them apply solely to a particular type of circuit such as switch-current or switch-capacitor filters and measure only a subset of the specifications. Because of their specialized applications, they will not be described here. However, the more general concept of instruments-on-a-chip deserves an explanation.

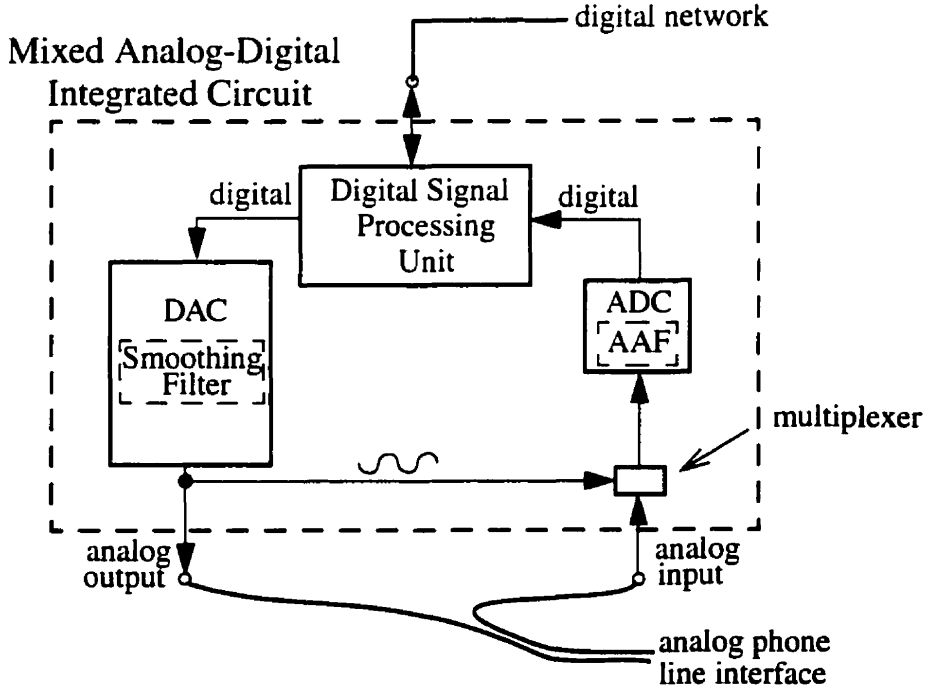

Figure 2.3: A typical mixed-signal IC for telephony applications.

2.3.1 Instruments-on-a-Chip

A straightforward implementation of BIST is to move the analog instruments, both the signal sources and the measuring units, to the integrated circuit. This usually involves some digital circuits, a digital-to-analog converter (DAC) and an analog-to-digital converter (ADC). In fact, some IC may already include these circuits. Such is the case for a telephone line chip, illustrated in Figure 2.3, also dubbed a voice CODEC (coder-decoder). It serves to interface the analog phone line and the digital phone company switching network. The DAC and the ADC may be used to perform analog measurements. However, before they do so, they must be verified to insure accuracy.

In a scheme introduced by Ohletz and labelled H-BIST, it is proposed to connect the output of the DAC to the input of the ADC so that a digital-input digital-output circuit is obtained [21]. A pseudo-random input is applied and a digital signature is compacted at the output. After test completion, this signature is compared with the value expected from a fault-less circuit. This scheme is a straightforward application of digital BIST. However it suffers from two shortcomings. First, a fault in one of the block could be offset by a fault in the other block. This problem is called fault-masking. The second pitfall exists because of the very nature of analog circuits. Indeed, while digital circuit behavior may be described with discrete values, the performance of analog circuit must be explained with continuous variables. Therefore, specifications usually list a nominal response and a tolerance band. Be-

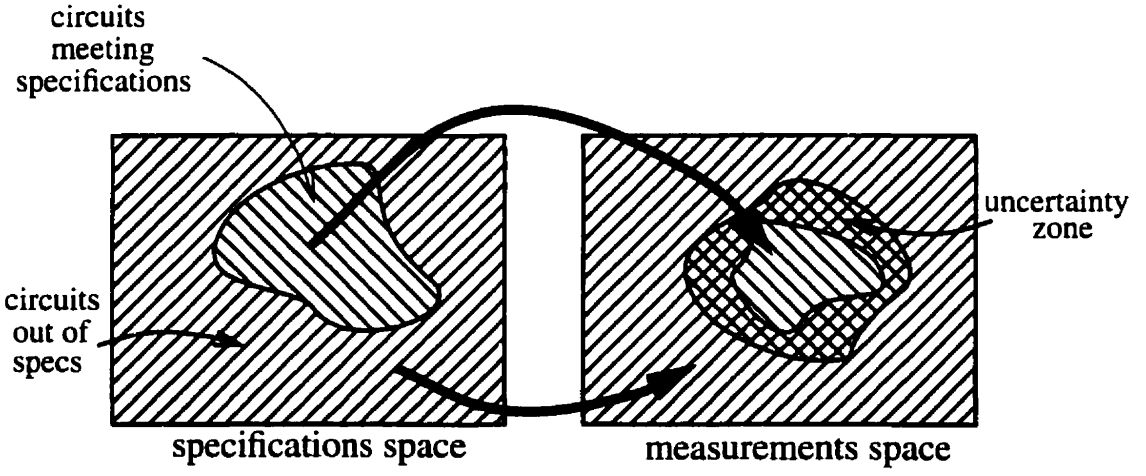

Figure 2.4: Mapping of analog circuits to non-specification measures.

cause of this variance on the accepted behavior, there may be many signatures indicating correct circuits. Furthermore, some signatures may originate from both valid and faulty circuits. This situation is illustrated in Figure 2.4. It can be seen that the boundary between good and faulty circuits is clear when dealing with specifications. On the other hand, for measurements unrelated to specifications, there may exist a zone where a decision can not be made regarding the functionality of a circuit.

2.3.2 MADBIST

The two shortcomings of the previous concept are addressed in a scheme called mixed analog-digital BIST (MADBIST) [22]. Fault masking is avoided by first verifying the ADC on its own before closing the loop. This, however, requires the generation of an analog signal to stimulate the ADC. Therefore, to implement the BIST capability in the mixed-signals IC of Figure 2.3, a lowpass delta-sigma ($\Delta\Sigma$) oscillator and a multiplexer have been added. This lowpass $\Delta\Sigma$ oscillator is an analog sinewave generator which is entirely digital except of a 1-bit DAC [23]. A description of this circuit may be found in Appendix A, Section A.5.2. A diagram of the test setup for the ADC is shown in Figure 2.5. A pulse density modulated (PDM) signal from the lowpass $\Delta\Sigma$ oscillator is applied to the input of the ADC. In the PDM bit stream is a high-quality low-frequency sinewave and high-frequency quantization noise. This noise is removed by the anti-aliasing filter (AAF) leaving a spectrally-

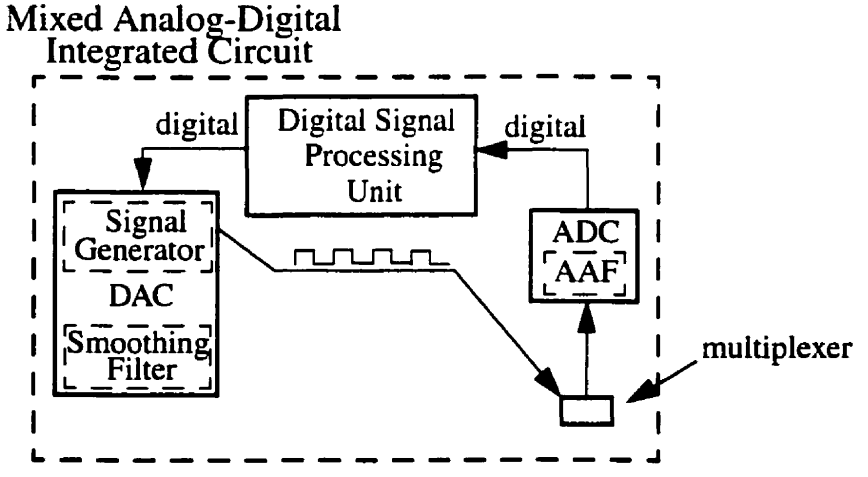

Figure 2.5: MADBIST test of an analog-to-digital converter.

pure sinusoid to excite the remaining circuitry. The digital output of the ADC is then analyzed by the digital signal processor (DSP) to separate the signal and the noise using a fast Fourier transform (FFT) or a digital filter. Analog measures such as signal-to-noise ratio (SNR), gain tracking and frequency response may then be obtained. Furthermore, the low-pass $\Delta\Sigma$ oscillator may be adapted to provide multi-tone signals [24]. This feature allows intermodulation distortion measurements [25] and rapid frequency response tests. A further advantage of MADBIST is that the results are digital representations of analog specifications. Therefore, the classification the devices or binning is straightforward, requiring only two digital comparators. The tolerance inherent to analog circuits is taken into account.

As most of the $\Delta\Sigma$ oscillator is digital, it can be tested almost fully using well established digital methods. This leaves only the 1-bit DAC and the multiplexer unverified in series with the DUT. These analog circuits are very small and thus much less likely than the DUT to be hit by defects. Nevertheless, should this occur, then the effect will either be catastrophic and easily noticed or subtle but causing signal degradation. The test results would show lower performances and a rejection of a good device is possible. However the acceptance of a bad part is unlikely.

After the ADC has passed all the tests, the DAC may now be verified by placing them in cascade. A digital signal, generated by the $\Delta\Sigma$ oscillator or the DSP, excites the DAC. The analog output is sent to the ADC which converts it back into digital form. A digital filter or FFT implemented by the DSP may then separate the noise from the signal. After the DAC is characterized, other analog circuits may be tested by placing them between the DAC and the ADC.

The MADBIST scheme illustrates the preeminent features to be found in a successful BIST. First, untested circuits should not be used to verify another circuit unless failure modes of the former are predictable and can be accounted for. Secondly, the continuous nature of analog specifications must be treated appropriately. That is, the tests must allow for a range of acceptance. Finally, digital circuits should be used as much as possible since they do not require calibration.

2.4 Summary

An overview of mixed-signal and analog production testing was provided. It was shown how, given a set of specifications, the functional test set of an IC is obtained. Test time may then be reduced by ordering and trimming the test set. As testers are expensive and because it is beneficial to optimize the test set before first silicon, inductive fault analysis is often used to simulate the fabrication and testing phases. However, this strategy has yet to gain acceptance for analog circuits. The concept of built-in self-test was explained and a solution for CODECs was described. Finally, desirable features of analog BIST schemes were reviewed.

Chapter 3 : Charge-Pump Phase-Locked Loops

The first description of phase lock appeared in the 1930's as a new method for radio reception [26]. However, the first widespread use of phase-locked loops (PLLs) was in the synchronization of horizontal and vertical scan in television receivers [27]. PLLs also proved useful early in receiving the weak signals subject to Doppler shift from spacecrafts [9].

The fundamental task of a PLL is to align the phase of an oscillator signal with the input signal. From this apparently simple behavior, important functions may be derived. Notably, PLLs are used for frequency multiplication and channel selection, as buffers and filters for clock signals and for recovering a periodic signal from a signal with missing transitions.

This chapter will first explain charge-pump phase-locked loops (PLLs) and study their mathematical model. Typical applications of these devices will then be presented. The metrics used in characterizing PLLs will be examined as well as the procedures for measuring them in a laboratory. Finally, the literature will be surveyed for papers related to PLL test.

3.1 Charge-Pump Phase-Locked Loop Structure

PLLs can be divided in two categories according to their input signal and consequently their phase detector. Historically, PLLs operated on sinusoids. They relied on combinational phase detectors such as analog multipliers and exclusive-or gates. Today, this type of PLL is employed mostly for high frequency applications such as frequency synthesis in wireless communication systems. The alternative is the use of a sequential phase detector. As the name implies, these phase detectors have some form of memory implemented with latches

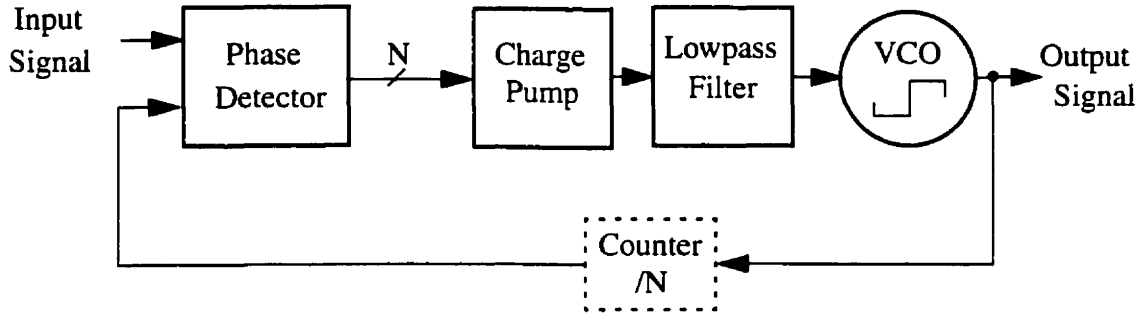

Figure 3.1: Charge-pump phase-locked loop functional block diagram.

or flip-flops. Since these devices are built from digital logic, only the signal transitions between two binary states are significant. They are thus particularly suited for digital signals where the phase information is coded in these transitions. Because the output of sequential phase detectors is discrete and multi-valued, a charge-pump is required for digital-to-analog conversion. Hence PLLs equipped with sequential phase detectors are commonly labelled charge-pump phase-locked loops.

A block diagram of a simple charge pump PLL [11] is shown in Figure 3.1. To the left, a sequential phase detector (PD) compares the transitions of the input and output signals. The output of the phase detector is a multi-valued digital signal and thus a charge pump is required for digital-to-analog conversion. A lowpass filter removes short-term variations and shapes the PLL characteristics. This filter is usually first-order and is implemented with passive components. The voltage-controlled oscillator (VCO) in turn generates a signal (a square wave for charge-pump PLL) whose frequency depends on the amplitude of its analog input. With the help of negative feedback, the VCO frequency is slowed down or increased until the phase of the input and output signal align. A counter may be inserted in the feedback loop to lock the VCO to a lower frequency reference signal.

Sequential Phase Detector

A number of sequential phase detectors have been introduced over the years. As PLL applications are numerous and diverse, different phase detector circuits are implemented in specific situations. A very popular one is the three-state phase-frequency detector (PFD)

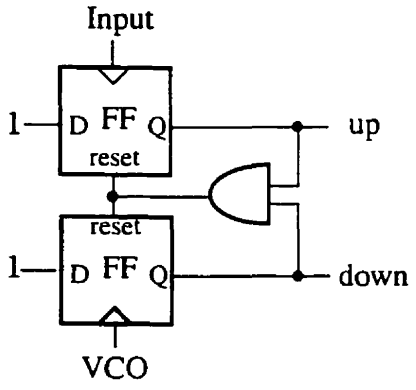

Figure 3.2: Three-state phase-frequency detector.

[28] illustrated in Figure 3.2. It is composed of two flip-flops triggered by one of the two input signals. Their reset line is controlled by an AND gate to insure that only one output may be high at any time. The PFD compares the edges of the input signal and the voltage-controlled oscillator (VCO) output and asserts the *up* signal if the input signal is leading or the *down* signal if it is lagging the VCO output. The pulse duration is equal to the time difference between the signal edges. This type of phase detector is thus linear with respect to phase difference. Furthermore, it is sensitive to a frequency difference as well, resulting in increased acquisition range and lock speed for the PLL. Typical waveforms out of this PFD are shown in Figure 3.3. It is assumed here that the flip-flops are positive-edge triggered. In part (a), the VCO output is late and the phase detector places a pulse train on the *up* line to increase the VCO frequency. In part (b), the VCO frequency is too high and the *down* output is triggered. This PFD however cannot tolerate missing transitions which can occur in some applications.

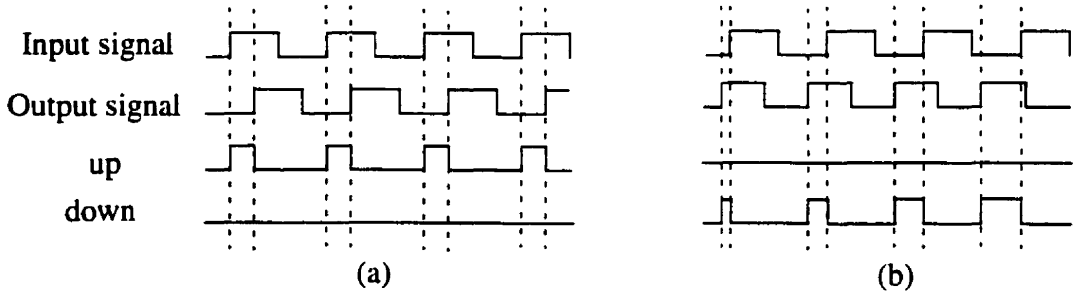

Figure 3.3: Phase detector waveforms: (a) Output lagging. (b) Output has higher frequency.

Charge Pump

Charge pumps may be of two kinds: voltage and current. Their schematics are illustrated in Figure 3.4 with their respective first-order passive filters. Current charge pumps are usually preferred because voltage charge pumps are not as precise since the amount of charge injected depends on the voltage at the output node. This makes it difficult to predict the charge-pump gain and consequently the PLL behavior.

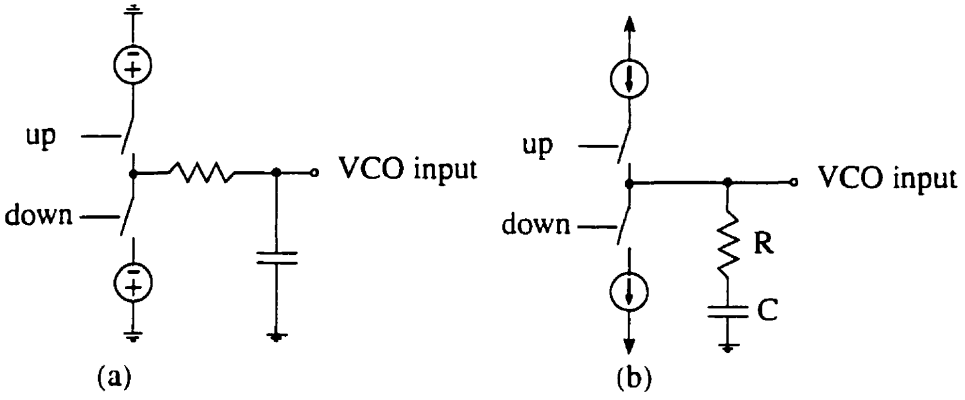

Figure 3.4: Charge pumps with first-order filters (a) Voltage. (b) Current.

Loop Filter

The purpose of the filter is two-fold. First it sets the closed loop behavior of the system. PLL characteristics such as jitter rejection and lock time will depend on the loop bandwidth. It must also filter high frequency components introduced by the charge pump.

The loop filter is traditionally implemented with passive components, as illustrated in Figure 3.4, except for rare implementations where DC gain is necessary. Its parameters are therefore dependent on the absolute value of resistors and capacitors which are very sensitive to process variations. In fact, for selected applications, filter components must be trimmed after the integrated circuit fabrication. The loop filter transfer function is typically first-order or second-order. However, the filter may also have to implement higher-order poles which have little influence on the loop characteristics but will remove the undesired high-frequency content of the pulses created by the charge pump.

Voltage-Controlled Oscillator

The voltage-controlled oscillator (VCO) generates a signal, either a square wave or a sinusoid, whose frequency is set by the input voltage. The characteristics of a VCO include the center frequency, the sensitivity, or frequency gain, and the linearity of its transfer function. VCO are also rated for their ability to reject power supply variations which is the mechanism accounting for most of the internal noise sources of the PLL.

The VCO must use some kind of unstable circuit to achieve oscillation. It is therefore no surprise that its characteristics will vary widely with process variations. For the purpose of this work, it is not required to understand its internal behavior. Interested readers are directed to the abundant literature on the subject. The two types of VCO used in charge-pump PLLs are the multivibrator [29,30] and the ring oscillator [31,32].

3.2 Phase-Locked Loop Mathematical Model

The PLL operation is generally non-linear. Nevertheless, a steady-state linear approximation is very useful in understanding its behavior and identifying the trade-offs involved. However, it should be remembered that any conclusion reached with the help of models must be validated with experimental results. This section will introduce the continuous-time and discrete-time model of the PLL. Then a simple example will be presented to compare the two models and highlight some performance measures.

The charge-pump PLL is a time-varying circuit because of the presence of switches. It seems therefore that the theory on linear time-invariant (LTI) networks may not be invoked. However, in most applications, the PLL changes state by only a small amount each cycle. As such it might be useful to average the behavior of the PLL over a few cycles so that LTI tools may be used. For this to give meaningful results, the bandwidth of the PLL must be significantly smaller than the input frequency [11].

3.2.1 Continuous-Time Model

The continuous-time model of a PLL, illustrated in Figure 3.5, is derived in a number of textbooks [7]. The terms $\theta_i(s)$ and $\theta_o(s)$ are the phase of the input signal and the output signal, respectively. K_P is the combined gain of the phase detector and charge-pump circuits given in A/rad or V/rad depending on the type of charge-pump. The gain of the VCO is denoted K_O and is stated in rad/V. If a counter is present then its effect is lumped into this constant by dividing the VCO intrinsic gain by the counter length (N in Figure 3.1). Using simple arithmetic, it can be found that the closed-loop equation governing the operation of the PLL is

$$H(s) = \frac{\theta_o(s)}{\theta_i(s)} = \frac{K \cdot G(s)}{s + K \cdot G(s)}, \tag{3.1}$$

where $K = K_P \cdot K_O$. $H(s)$ is also labelled jitter transfer function since the input and output variables refer to the phase of the respective signals. The system is lowpass with order one larger than the filter transfer function $G(s)$. The output of the phase detector is also called the phase error $\theta_e(s)$ and is defined as $\theta_i(s)$-$\theta_o(s)$. It may be useful to examine the relation between this value and the input phase signal. From the diagram shown in Figure 3.5, the following equation can be derived:

$$\frac{\theta_e(s)}{\theta_i(s)} = \frac{s}{s + K \cdot G(s)}. \tag{3.2}$$

This transfer function has a highpass nature owing to the fact that it reduces to zero as $s \to 0$.

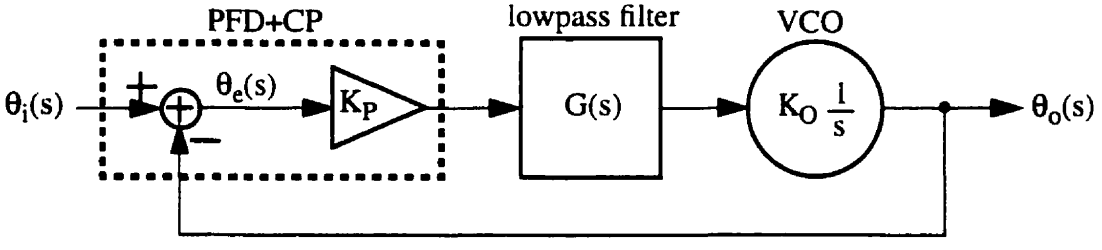

Figure 3.5: Continuous-time model of a phase-locked loop.

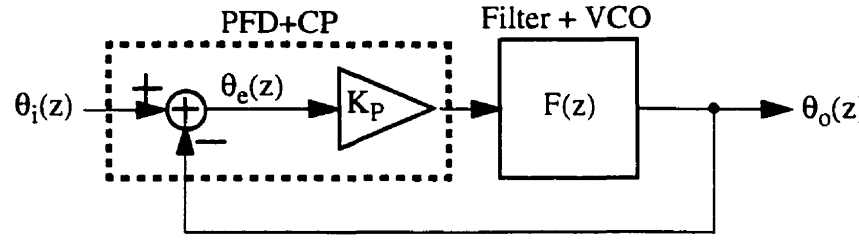

Figure 3.6: z-domain model of a phase-locked loop.

3.2.2 Discrete-Time Model

The continuous-time model yields satisfactory results for PLLs equipped with combinational phase detectors and therefore processing sinewaves. On the other hand, the phase of digital signals is contained in the signal transitions and is thus better represented as a discrete-time sequence. Therefore, the analysis of PLLs operating on digital signals should be performed using difference equations or z-transform tools. Indeed, it has been shown that the discrete-time model is more accurate for these devices, especially at high jitter frequency [33]. The z-domain representation of a charge-pump PLL is shown in Figure 3.6.

It should be noted that the gain of the phase detector (K_P) is scaled down by $\frac{1}{f_{vco}}$, the period of the VCO clock. The closed-loop equation is now

$$H(s) = \frac{\theta_o(z)}{\theta_i(z)} = \frac{K_P \cdot F(z)}{1 + K_P \cdot F(z)}, \tag{3.3}$$

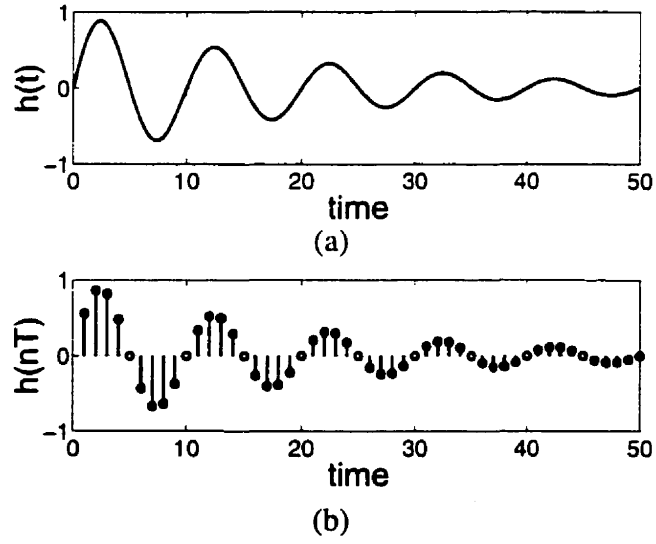

Figure 3.7: (a) Continuous-time impulse response. (b) Impulse invariant transform.

where $F(z)$ is the impulse invariant transform of $\frac{K_O G(s)}{s}$. This transform preserves the values of the impulse response of the continuous-time system at the sample times of the discrete-time representation. Figure 3.7 illustrates this relationship between a continuous-time impulse response and its impulse invariant transform.

3.2.3 Example: Second-Order Charge-Pump Phase-Locked Loop

From the equations presented above, it may be difficult to get a grasp on the PLL operation. To illustrate how this theory is applied, a simple PLL will be analyzed. A current charge pump and first order passive filter such as that seen in Figure 3.4 (b) are assumed. The filter transfer function is

$$G(s) = R + \frac{1}{sC}. \tag{3.4}$$

Using Eq. (3.1), the jitter transfer function is found to be

$$H(s) = K_P K_O R \frac{s + \frac{1}{RC}}{s^2 + K_P K_O R s + \frac{K_P K_O}{C}}. \tag{3.5}$$

While the filter is first-order, the system is second-order because the VCO act as an integrator on phase. This continuous-time transfer function will be compared to the discrete-time model. First, note that

$$F(s) = \frac{K_O G(s)}{s} = \frac{K_O}{C}\frac{RCs + 1}{s^2} = \frac{K_O}{C}\left[\frac{1}{s^2} + \frac{RC}{s}\right]. \tag{3.6}$$

The impulse response of $F(s)$ is found by taking the inverse Laplace transform, resulting in

$$h(t) = \frac{K_O}{C}[t + RC]u(t). \tag{3.7}$$

The impulse invariant property requires that

$$h[n] = \frac{K_O}{C}[nT + RC]u[n], \tag{3.8}$$

where T is the VCO period. The inverse z-transform of Eq. (3.8) is then computed and the result is

$$F(z) = \frac{K_O}{C}\left[\frac{Tz^{-1}}{(1-z^{-1})^2} + \frac{RC}{1-z^{-1}}\right] = RK_O z\left[\frac{z-(1-T/(RC))}{(z-1)^2}\right]. \tag{3.9}$$

For simplification, let α = 1-T/RC, resulting in

$$F(z) = RK_O \frac{z(z-\alpha)}{(z-1)^2}. \tag{3.10}$$

Substituting Eq. (3.10) into Eq. (3.3), the following transfer function describing the PLL in the discrete-time domain is obtained:

$$H(z) = \frac{K_P K_O R}{1+K_P K_O R} \cdot \frac{z(z-\alpha)}{z^2 - \frac{2-\alpha K_P K_O R}{1+K_P K_O R} z + \frac{1}{1+K_P K_O R}}. \tag{3.11}$$

f	155 MHz
K_P	8×10^{-7} A/rad
K_O	5×10^{8} rad/V*s
R	1 kΩ
C	220 nF

Table 3.1: Parameters for an OC-3 compliant phase-locked loop.

Typical parameter values for a device compliant with the OC-3 synchronous optical network (SONET) [34] specifications are shown in Table 3.1. Using these values in the expression for the continuous-time model (Eq. (3.5)), it is seen that the system is characterized by a pair of poles located at 730 Hz and 63 kHz. A left-hand-plane zero is also present at 723 Hz. The jitter transfer function using the continuous-time model (Eq. (3.5)) and the discrete-time model (Eq. (3.11)) are plotted in Figure 3.8. Both models yield similar results except at high frequencies. In Figure 3.8 (b), it can be seen that the jitter transfer function peaks before the roll-off. This can be attributed to the presence of a transfer function zero at a lower frequency than the poles. This behavior is unavoidable in classical PLL architectures. The maximum gain of the jitter transfer function is in fact called

jitter peaking. This is an undesirable feature as it leads to the accumulation of jitter in strings of repeaters for data communication networks [35]. For PLLs in these applications, a bound is placed on jitter peaking such as 0.1 dB for the OC-3 specification. The other parameter of interest in the magnitude response of the jitter transfer function is the PLL bandwidth defined as the frequency where the gain reduces by 3dB from the DC level.

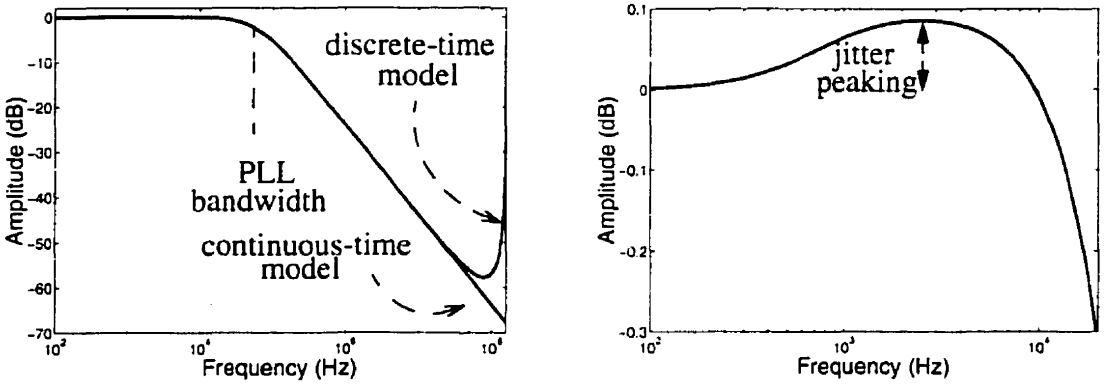

(a) (b)

Figure 3.8: Jitter transfer function (a) Up to Nyquist frequency. (b) Magnification of the jitter peaking region.

3.3 Applications of Charge-Pump Phase-Locked Loops

The design, and thus the test, of a particular PLL strongly depends on the application targeted. This section will thus examine the two foremost applications of charge-pump PLLs. They are the recovery of the clock signal in digital data transmission and the distribution of system clocks.

3.3.1 Timing Recovery

Most digital communications systems such as wireless phones, disk drives or fiber optic links transmit or store only data signals. Usually, it is not desirable or feasible to send or save the clock signal with the data. This clock signal is nevertheless embedded in the received signal through the data transitions. PLLs are essential in recovering this clock signal

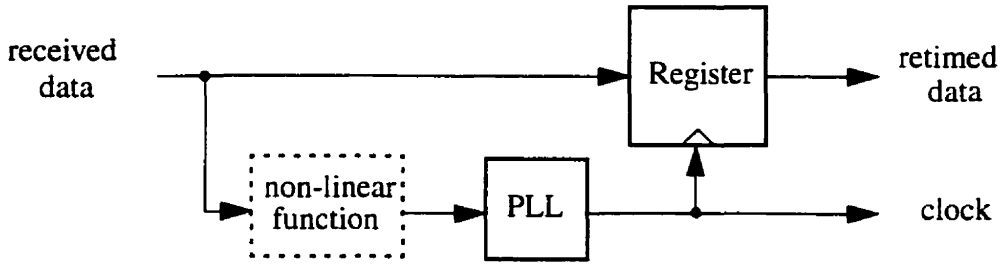

Figure 3.9: Digital receiver using a phase-locked loop.

at the receiver. The clock signal thus obtained is then used to regenerate the data signal with the help of a register as illustrated in Figure 3.9.

The received data signal may be strongly corrupted by the communication channel, adjacent symbols or noise. Furthermore, it may present long runs of the same symbol without transitions. Clock recovery thus necessitates state-of-the-art PLLs because of the high data rate and stringent jitter requirements. The effect of jitter for data regeneration can be visualized with the help of an eye diagram such as the one displayed in Figure 3.10. It is obtained with an oscilloscope triggered at the same frequency as the data rate. The four possible transitions of a binary alphabet are clearly visible on each side. These transitions are fuzzy lines as inter-symbol interference and noise from the communication channel are present. The PLL clock should be phase locked to the ideal sampling instant, where the noise margin is maximum. Any jitter is likely to reduce the tolerance on noise. In fact,

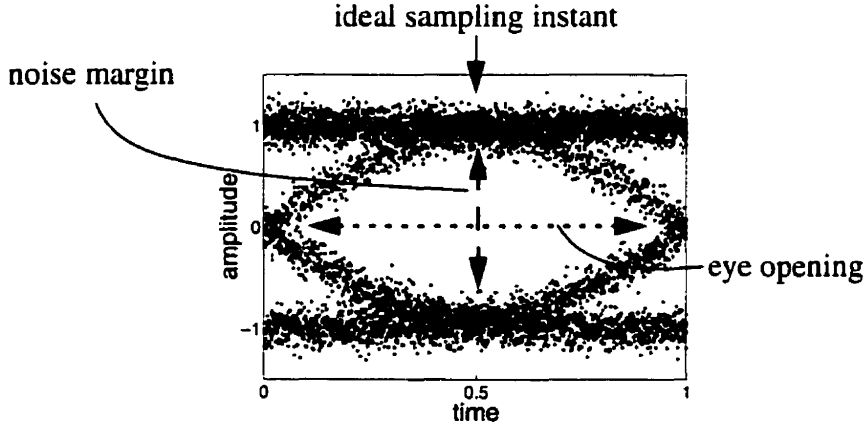

Figure 3.10: Typical eye diagram.

should the peak-to-peak jitter exceed the eye opening then errors will occur even in the absence of noise.

The data signal must exhibit significant power at the clock frequency or the PLL will not lock. Figure 3.11 shows the power spectral density of random signals for three different line signal codings [37]. The frequency scale is relative to the clock frequency of the transmitted signal. Two line signal codings have spectral null at the clock frequency. Therefore some non-linear processing would be performed before feeding this kind of signal to the PLL.

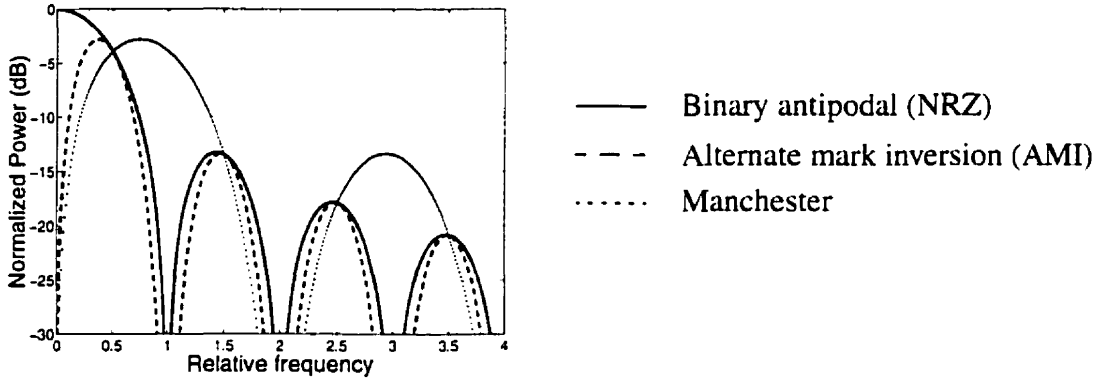

Figure 3.11: Power density of random signals using different line signal codings.

3.3.2 Clock Distribution

Timing information in large silicon systems has been traditionally distributed using a clock tree as depicted in Figure 3.12. Because of buffer delay, transmission line effects and load capacitance C_L, the clock signal at the leaves will be delayed with respect to the master clock. Furthermore, as other chips have different clock trees and different skews, it becomes increasingly difficult to synchronize data exchanges.

A PLL may be used to reduce this skew. The circuit of Figure 3.13 phase locks the distributed clock to the reference clock [38]. This arrangement also allows the distribution of a frequency multiple of the reference clock with the addition of a counter in the feedback loop. Having a low-frequency reference clock on a printed circuit board reduces electro-

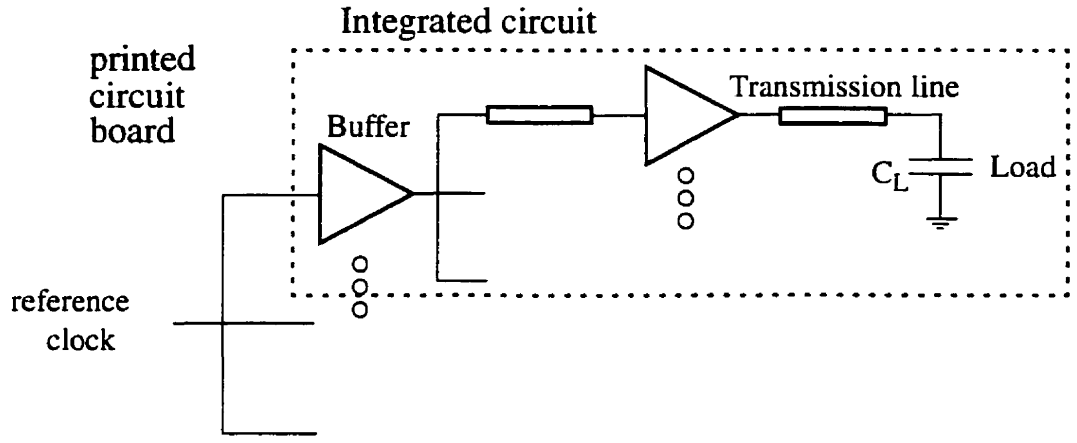

Figure 3.12: Clock tree on integrated circuits.

magnetic interference and power consumption. PLL-based clock distribution is essential in large micro-processors as internal clock frequencies reach a few hundred megahertz [39]. Other digital ICs such as digital signal processors (DSP) and field-programmable gate arrays (FPGA) are expected to follow suit.

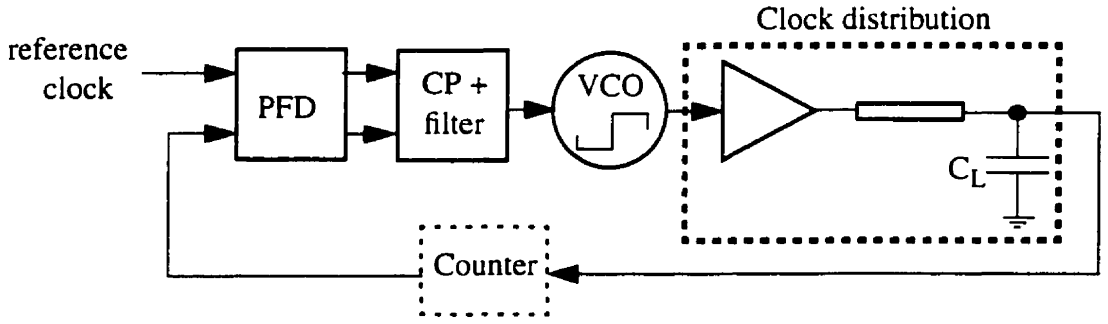

Figure 3.13: Clock skew reduction using a phase-locked loop.

3.4 Phase-Locked Loop Metrics

Specifications for PLLs may be stated in numerous ways according to the applications. In fact, the abundance of metrics may confuse newcomers to the PLL area. It is therefore useful to classify the specifications in three groups: first those measuring the input-output characteristics of jitter in steady-state operation, then the metrics assessing the internal noise sources of the PLL and finally the specifications qualifying the dynamic or transient oper-

ation. In this section, various metrics of PLL performance will be described as well as the test setup used to measure them.

3.4.1 Transfer Function in Steady-State Operation

The specifications in this group describe the jitter input-output behavior of the PLL after lock has been acquired. They will be most useful in applications where the PLL operates on a signal at the output of a noisy channel or for clock recovery. On the other hand, these specifications would be less important when the PLL is driven by a crystal for example.

Jitter Transfer Function

The jitter transfer function is the relation between the jitter in the input signal and jitter at the VCO output as described in Eq. (3.1). The parameters of interest on the curve are the bandwidth and the jitter peaking. Figure 3.14 illustrates a test setup for the measurement of the jitter transfer function [8]. The signal source generates an RF signal which is modulated with the source signal from a spectrum analyzer using an Armstrong phase modulator [40]. The resulting signal is used as the clock signal for a bit generator. The clock signal recovered by the PLL is then downconverted to obtain the phase argument.

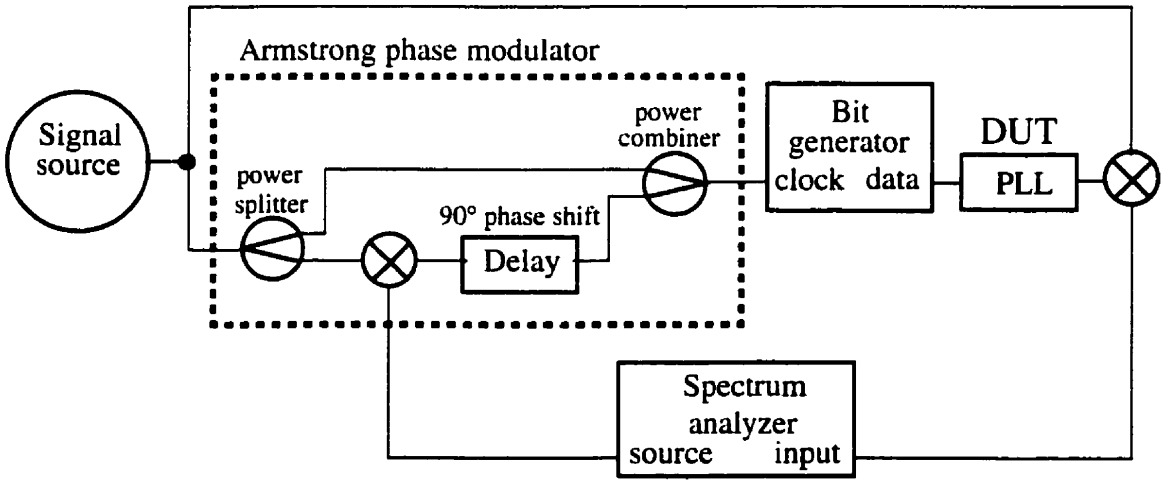

Figure 3.14: Jitter transfer function test fixture.

Jitter Tolerance

Jitter tolerance is a measure of amount of phase drift in the input signal that can be accommodated in clock recovery circuits before significant bit errors result. A frequency modulated sinewave is used as the clock for stimulus generation. The mechanism which leads to errors can be understood by studying the linear model of Figure 3.5. The wrong data will be recovered when the phase error, $\theta_e(t)$, corresponding to the time difference between the VCO output and the input signal exceeds the eye opening of the data (illustrated in Figure 3.10). Yet, for the VCO to track the input signal, a sinusoidal phase error $\theta_e(t)$ must be created. The jitter tolerance test thus measures the maximum input jitter which can be generated without the phase error going above a threshold set by the eye opening. Intuitively, at low frequency, the open loop gain is large and the maximum error signal will be able to help the VCO track large frequency deviations. On the other side, the small open loop gain at high frequency makes the same phase error result in a much smaller frequency deviation at the output of the VCO. Jitter tolerance is traditionally quoted as the peak-to-peak jitter ($jitter_{pp}$) in unit intervals (ui)[1] instead of the maximum frequency deviation (Δf) of the FM signal used for the stimulus. To find the relation between these two measures, first note that the instantaneous frequency of a modulated signal is

$$f(t) = f_c + \Delta f \cdot \cos(2\pi f_m t), \tag{3.12}$$

where f_c is the carrier frequency and f_m is the modulating signal frequency. The phase is the integral of the frequency and is thus defined as

$$\phi(t) = 2\pi\int_0^t f(\tau)d\tau = 2\pi f_c t + \frac{\Delta f}{f_m} \cdot \sin(2\pi f_m t) \text{ (rad)}. \tag{3.13}$$

Finally the peak-to-peak jitter will be twice the phase deviation:

$$jitter_{pp} = 2\frac{\Delta f}{f_m}(rad) = \frac{\Delta f}{\pi f_m}(ui). \tag{3.14}$$

The transfer function defined by the maximum amplitude of the input jitter tolerated is in fact the inverse of the phase error transfer function as defined in Eq. (3.2). As $\theta_e(s)$ is fixed

1. 1 ui = 2π rad

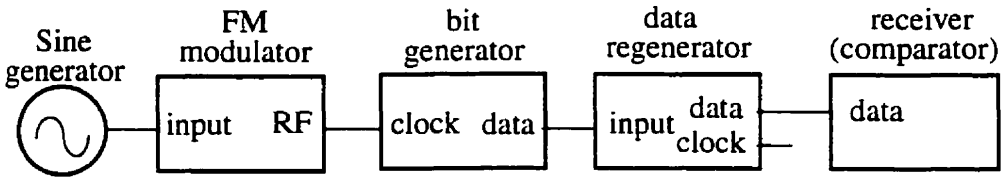

Figure 3.15: Jitter tolerance test setup.

by the size of the eye opening, the corresponding input jitter $\theta_i(s)$ is determined and thus the transfer function $\theta_i(s)/\theta_e(s)$ is computed. The corner frequency is therefore the same as the bandwidth of the jitter transfer function. It differs from this last metric in that the VCO frequency range will limit the jitter tolerance at low frequencies. A standard test set-up used to measure jitter tolerance [8] is illustrated in Figure 3.15. The two left most blocks generate a frequency modulated sinewave which is used as the clock input of the bit generator. The data is then sent to the device-under-test (DUT) which recovers the clock and makes decisions on the data (refer to Figure 3.9). The bit receiver then compares this with the data sent. The amplitude of the sinewave is increased until the BER reaches the threshold. The procedure is repeated for different sinewave frequencies.

3.4.2 Internal Noise Sources

The following set of metrics measures the phase noise generated by the PLL blocks which appears at the output. The main source is generally the VCO. However, the pattern jitter introduced by the phase detector and by the high-frequency content of the pulses from the charge pumps also contribute to phase noise.

Jitter Histogram

To assess the jitter present in the output clock or in the retimed data, a histogram is constructed from the timings of the zero crossings over a large number of periods. In the case of clock recovery circuits, the input of the DUT can be a signal containing a maximum number of transitions or a pseudo-random bits signal (PRBS). The latter type will usually generate more jitter because long runs of zeros or ones reduce the rate at which the VCO input is adjusted. From the histogram, three parameters can be extracted. The most important one is the standard deviation or root-mean-square (RMS) value of the jitter distribu-

tion. The mean of the jitter distribution with respect to the zero crossings of a reference signal may also be stated. Finally, the peak-to-peak jitter is sometimes quoted. These quantities can be stated in absolute time values, or relative to the clock frequency and expressed in degrees.

Phase Noise

The phase noise is a measure of the signal power near the carrier frequency on the power density spectrum. This power density denoted in units of dBc/Hz where the subscript c denotes that the measure is relative to the carrier power is quoted at a given frequency offset from the carrier frequency. Figure 3.16 illustrate how phase noise is measured from a power density spectrum plot. The power has been normalized to a 1 Hz resolution bandwidth. In this example, the phase noise happens to be 60 dBc/Hz at 100 Hz. This metric is mostly used for frequency synthesizers.

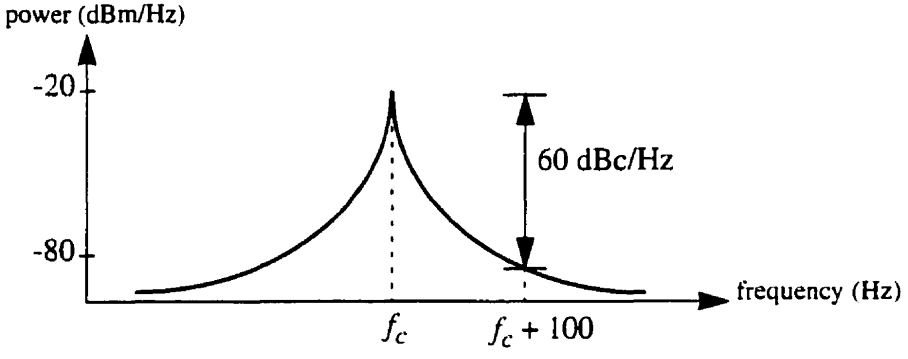

Figure 3.16: Measurement of phase noise on a power density spectrum plot.

3.4.3 Dynamic Operation

The final set of metrics to be discussed describe the characteristics of a PLL when out of lock. Transient measurements are used to obtain these metrics. Generally, the results of these tests are predictable using the steady-state measurements and thus they are not necessary to characterize the device. However in some applications, there exist mechanisms to aid in acquisition which are later disabled and thus not measured if the device is always in phase-lock during the test.

Phase Step or Frequency Step Transient

A phase step is applied at the input of the PLL and the error signal at the output of the charge-pump is monitored. Of particular interest in many applications are the acquisition and settling times. For example, in frequency synthesis for wireless communication the transient will result in power leakage into other channels. In portable computing devices, some chips are powered down when idle and upon powering up, must be able to acquire timing lock quickly.

Capture Range and Lock Range

To realize these tests, a sine wave (jitter free signal) is applied at the input of the PLL. The frequency is varied and the range in which the PLL is in lock is recorded. The frequency range obtained when going from out-of-lock to lock is defined as the capture range while the frequency range in which lock can be held is denoted the lock range. The lock range depends only on the frequency range of the VCO while the capture range is also a function of the phase detector. For sequential phase detectors, the capture range is the same as the lock range. These metrics are usually not quoted for charge-pump PLLs. However, lock range will set an upper amplitude bound on jitter tolerance as it limits the maximum frequency deviation of the FM signal used in that test.

3.5 Production Test of Phase-Locked Loops: Literature Review

Relatively few papers have appeared in recent years on production test of PLL. Other analog circuits such as ADCs, DACs or filters have generated larger bibliographies. The papers on PLL may be classified in four broad categories. Those in the first set suggest improvements to automated test equipments (ATE) in order to maintain measurement accuracy in the face of ever higher PLL operating frequencies. In the second group are communications regarding test strategies for system ICs incorporating PLLs. Then, the application of induc-

tive fault analysis (IFA) and structural testing of PLL are grouped together. Finally, the last set includes papers on the topic of on-chip measurement of PLL specifications.

3.5.1 Improving Automatic Test Equipment for PLL Test

With reduced silicon device dimensions, the test timing accuracy cannot keep up with the fast decrease of the device under test (DUT) clock period [3]. ATE manufacturers must resort to alternative test strategies to be able to measure the ever faster devices. In [41], the authors propose a digital time interpolator to achieve a timing accuracy better than the tester clock frequency. This solution will be examined in details in Chapter 5.

Another suggestion is to use a vernier strategy to improve the output signal sampling resolution [42]. The output is sampled at a clock rate slightly larger than the phase-locked loop frequency. In the example given in the paper referenced above, a 32-bit sequence is repeated at a 125 Msymbol/s rate for a 256 ns periodic signal. The output is captured using a clock with a 256.0625 ns period. Thus each sample will have a 62.5 ps offset within the period when compared to the previous one, as illustrated in Figure 3.17. Thus 4096 samples will cover the whole sequence. Now since jitter noise is not deterministic, a large number of equivalent cycles are captured to obtain a probability distribution of jitter. This method alleviated the limitation of the sampling rate but requires very good timing accuracy nevertheless.

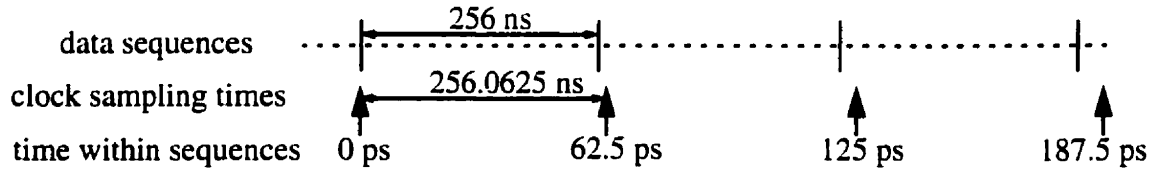

Figure 3.17: Measurement using a vernier strategy.

3.5.2 Integrated Circuits with PLL: Test Strategies

One paper discusses test of a PLL as a building block of an IC [12]. This IC is an integrated services digital network (ISDN) U-transceiver. The U-interface is located at the contact point between the access line and customer's equipment such as a modem. A PLL is required to derive a synchronization clock from the incoming data. As a very low loop band-

width is required, the authors have ruled out functional testing because of the long transient time. Fortunately, most of the PLL components are external to the chip and may be tested separately using conventional methods. Thus, the filter and the VCO which are usually the most difficult blocks to test are not a concern here. Only biasing circuitry of the VCO needs be verified. As this situation is not representative of the level of integration found today, this strategy is of little interest.

3.5.3 Inductive Fault Analysis for PLL

PLLs offer great challenges for IFA. Some components such as the VCO are very sensitive to process parameters. Therefore small modeling errors will render the result meaningless. Also the simulation of PLL is known to be computationally very expensive, owing to the mathematical stiffness of the problem. Two papers address the use of hierarchical models to speed up simulations involving PLLs [43, 44]. However, in the absence of a thorough study of the effectiveness of IFA for mixed-signal circuits, IC manufacturers are reluctant to adopt this technique.

3.5.4 Design For Test Techniques For Phase-Locked Loops

Papers on improving the measurement of PLL specifications with circuit structures on the die have appeared only recently. One such design for test (DFT) scheme targets ring oscillators which are commonly employed as VCOs [45]. A diagram of a circuit adapted for this technique is shown in Figure 3.18. During normal operation of the PLL, the VCO input controls the bias level of the inverters and the feedback loop is closed. To measure the VCO

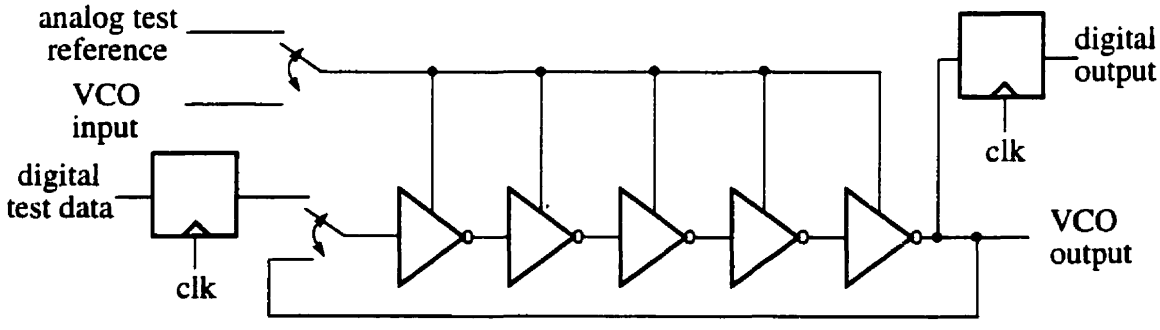

Figure 3.18: Design for test strategy for ring oscillator (VCO).

frequency range, the ring oscillator is turned into an open loop circuit by disconnecting the feedback. The delay through the inverter chain is then measured using two flip-flops at each end by varying the clock frequency. The delay is thus compared against different values of clock period. Measurements are repeated for different bias levels as set by the analog test reference input.

On the other hand, a company named LogicVision recently announced a built-in self-test scheme for PLLs. Specifications measured would include the lock range, the lock time and the jitter transfer function. Details have not been released yet as a patent has been applied for. However the method is claimed to be digital and non-intrusive [46].

3.6 Summary

The charge-pump phase-locked loop is an important building block for integrated circuits. With the ever increasing complexity of ICs and the divergence of on-chip and board clock speed, PLLs will likely be part of an ever increasing proportion of chips. There exists a large number of metrics to evaluate PLL performance according to the host application. The level of research activity on PLL test has traditionally been low. However, with the announcement of a commercial product, the interest level should take off.

Chapter 4 : Stimulus Generation

The traditional procedure for testing an analog device is to apply a signal at the input and measure the output response. For example, to obtain the frequency response of a filter, a sinewave is used to stimulate the circuit and the amplitude and phase of the output signal are evaluated. However, testing an arbitrary analog device may be very expensive. The signal sources and the measuring instruments can be very complex. Notably, it is the case for PLLs as explained in Section 3.4. Thus, methods for characterizing this device on a test bench may not be directly ported on-chip. Unfortunately, the test circuitry would be very sensitive to process variations, requiring a significant test effort itself, and the area overhead would be too large. Yet, for circuits based on a negative feedback loop such as a PLL, other options are available for test stimulus and measurement. In fact, because of the loop dynamics, a test signal may be injected at any point and stimulate the whole circuit. Furthermore, the loop behavior may, in theory, be measured at any node. While the problem of signal measurement in the PLL will be addressed in the following chapter, the different strategies for PLL stimulation will be studied here.

Figure 4.1 illustrates the continuous-time model of a PLL. The four possible nodes for test signal injections are labeled with the variables $\theta_i(s)$, $\theta_e(s)$, $\theta_v(s)$ and $\theta_o(s)$. Note that two nodes appear to have been left out. They are in fact scaled versions of other nodes and, as

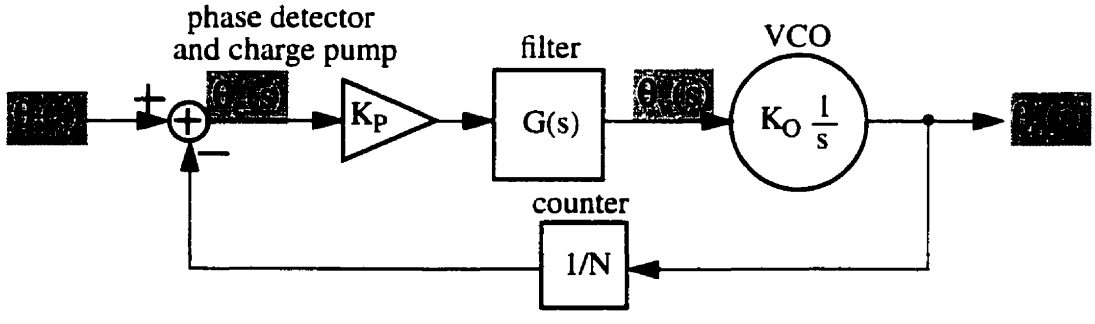

Figure 4.1: Phase-locked loop continuous-time model.

such, are not considered distinct. Now, the list of candidate nodes for device stimulation should be curtailed by an important criterion, the possibility of introducing a summing operation for the test signal injection at the appropriate point in the circuit. Each of the aforementioned node will be examined for this requirement. First, as the input variable, denoted $\theta_i(s)$, refers to jitter on a digital signal, a summing operation at this point of the circuit may be performed by a variable delay circuit. This technique will be described in greater length in upcoming sections. The output variable, $\theta_o(s)$, is of the same nature and therefore a similar solution may be applied. On the other hand, it is very difficult to sum a signal at the input of the VCO, labelled $\theta_v(s)$ on the diagram. The quantity represented by this variable is a voltage for which no simple circuit exists for addition. Finally, the last candidate node is the input of the loop filter, equivalent to $K_P\theta_e(s)$. It turns out that its possible use for test signal injection is conditional on the type of charge pump and the details of the filter implementation. For a current charge pump or the combination of a voltage charge pump and an active filter, the summation of signals may be implemented easily. However, no trivial solution is available for a voltage charge pump coupled with a passive filter. Fortunately, this configuration is seldom employed in modern designs as it results in non-linear PLL characteristics.

This chapter explores the three different methods for generating a stimulus suitable for a PLL. Each proposition will be evaluated for its amenability to built-in self-test (BIST). Explicitly, signal sources must have four essential characteristics to qualify for on-chip integration. First, circuits which require calibration using off-chip analog instruments should be avoided. Otherwise, the test of the PLL with external instruments is merely replaced by a similar effort to calibrate the signal source. Also, the silicon area required for the source implementation should be relatively small when compared to the device under test. Finally, the signal source should also provide sufficient accuracy and offer some degree of programmability.

4.1 PLL Input Test Signal

The characterization of a charge-pump PLL through its input signal requires the generation of a square wave (digital signal) whose phase, sometimes called jitter, can be controlled. Conceptually, this signal may be obtained from a sinusoidal signal fed to a comparator. This signal, *s(t)*, is thus expressed as

$$s(t) = \mathrm{sgn}(\sin[2\pi f_c t + \theta(t)]), \tag{4.1}$$

where *sgn* is the signum (sign) function, f_c is the carrier frequency and $\theta(t)$ is the instantaneous phase. The amplitude of the signal is irrelevant as *s(t)* is a binary digital signal. Here the two levels are assumed to be +1 and -1. The instantaneous frequency of the argument of the signum function, f_i, is defined as

$$f_i = f_c + \frac{1}{2\pi}\frac{d}{dt}\theta(t). \tag{4.2}$$

Returning to Eq. (4.1), the instantaneous phase should be a sinusoidal function because the jitter transfer function which is the aim of this work (refer to Section 3.4.1) is a frequency domain measure. It is therefore of the form

$$\theta(t) = A\sin(2\pi f_j t), \tag{4.3}$$

where f_j is the jitter frequency and A is its amplitude. Using this definition in Eq. (4.1), the resulting signal may thus be represented as

$$s(t) = \mathrm{sgn}(\sin[2\pi f_c t + A\sin(2\pi f_j t)]), \tag{4.4}$$

Replacing $\theta(t)$ in Eq. (4.2) with the corresponding expression in Eq. (4.3), the frequency of a carrier which is phase modulated with a sinusoidal signal is

$$f_i = f_c + Af_j\cos(2\pi f_j t). \tag{4.5}$$

It can be seen that the desired stimulus is equivalent to a carrier which is frequency modulated (FM) with a sinewave. Fortunately, this type of signal is common and a number of circuits have been introduced over the years for its generation. Some of these circuits will be reviewed next. Their limitations will be exposed and the need for a different approach

will be argued. Finally, delta-sigma (ΔΣ) digital phase modulation, a fully digital scheme, will be introduced.

4.1.1 Analog Frequency Modulated Signal Generation

Frequency modulated signals may be generated with analog circuits such as those shown in Figure 4.2. In part (a), a low-frequency sinewave is fed to the input of a voltage-controlled oscillator (VCO). The carrier frequency is equal to the VCO output frequency for a zero input value. Oscillators and VCOs are very sensitive to process variations and therefore, calibration is mandatory before these circuits can be used. The VCO is especially difficult to calibrate so a different arrangement, illustrated in Figure 4.2 (b), is usually preferred for broadcast FM generation. In this circuit, the carrier is amplitude modulated (AM) by the low-frequency sinusoid. As the narrowband FM spectrum is very similar to the AM spectrum, the output signal is a sufficient approximation. This scheme still requires the tuning of two analog oscillators. Furthermore, it is difficult to make them programmable to obtain FM signals of varying characteristics.

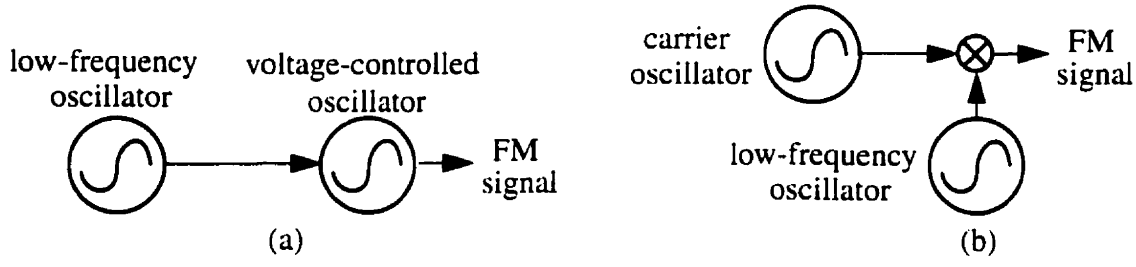

Figure 4.2: FM signal generator: (a) With VCO. (b) With amplitude modulation.

4.1.2 Digital Synthesis of a Frequency Modulated Signal

As discussed in Chapter 2, sensitive analog circuits are not suitable for BIST applications. Digital circuits are a better choice as they do not require calibration, are very accurate and may be programmed for various signal characteristics. Therefore, the preferred solution for analog signal generation in BIST is to combine signal processing in the digital domain and digital-to-analog conversion. In fact, the latter operation may also benefit from a shift of the signal processing burden from analog to digital. A technique which realizes this is delta-

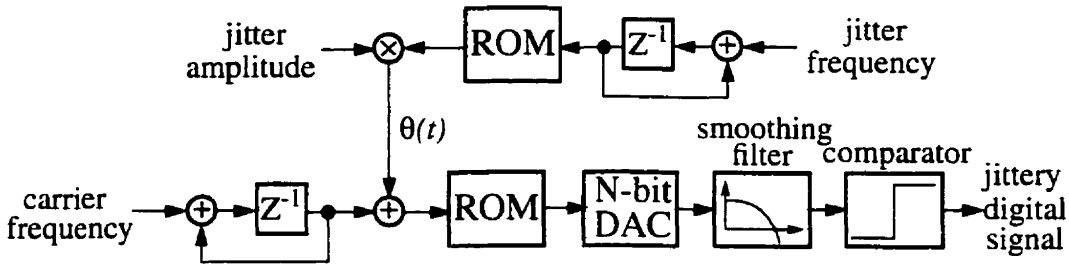

Figure 4.3: ROM-based digital frequency synthesis.

sigma ($\Delta\Sigma$) modulation. It can be viewed as a coding scheme for compressing analog or N-bit signals into single-bit streams. A detailed explanation is provided in Appendix A.

The most straightforward and versatile digital signal generation scheme is a ROM-based frequency synthesizer followed by a digital-to-analog converter (DAC). A diagram of an FM synthesizer using this technique is illustrated in Figure 4.3. At the upper right-hand corner, the jitter frequency argument is integrated to obtain a phase signal. This signal then references a ROM where sine samples are stored. The output is a sinewave scaled by the desired jitter amplitude to obtain the instantaneous phase $\theta(t)$. The result is then added to the phase of the carrier and a digital FM signal is generated with the help of a second look-up operation. The modulated signal is converted to the analog domain with the help of an N-bit DAC. A comparator finally generates the jittery square wave. There are two problems with this proposition for a BIST application. First, the ROM implementation requires a large silicon area and so does the multiplier. However the biggest problem lies with the DAC and the smoothing filter. As these circuits are analog, they cannot be verified with digital methods before the PLL test. In fact, the DAC may be more difficult to test than the PLL.

A possible solution would be to move the comparator just after the ROM, before the digital-to-analog conversion. Then the input of the DAC would become a single-bit signal. This arrangement is illustrated in Figure 4.4 (a). One-bit DACs are very simple to design and are inherently linear. To examine the effects of this modification, signals from this circuit are compared with signals from the circuit of Figure 4.3 and the results are displayed in

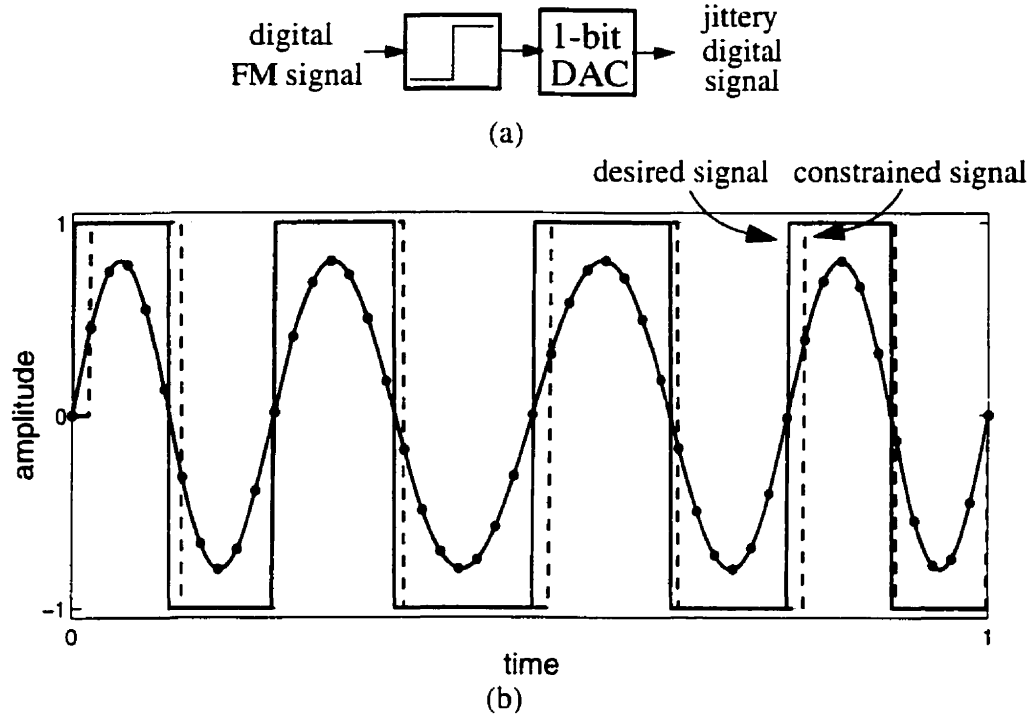

Figure 4.4: Performing the comparison digitally: (a) Block diagram. (b) Signals.

Figure 4.4 (b). The dots show the discrete values of the digital sinusoidal signal. These points are connected by the continuous line representing the signal that would result from a perfect digital-to-analog conversion. The solid square wave is derived from it using an idealized sign operation. This is the desired signal. The dashed square wave shows the result when the comparison is performed in the digital domain. Obviously, significant errors result because the signal transitions are constrained to the clock rising edges. To reduce the magnitude of theses errors and thus improve accuracy, a solution is to increase the ratio of the clock frequency to the PLL frequency. However, in most applications, the PLL frequency is too large to afford a ratio that would lead to sufficient signal quality.

4.1.3 Delta-Sigma Digital Phase Modulation

Looking at Figure 4.4 (b), it can be seen that performing the comparison in the digital domain results in errors in the zero crossings. They occur because these zero crossings are constrained to the clock edges. This situation is analogous to the general problem of digital-to-analog conversion. Indeed, this latter operation aims at representing a continuous signal

with discrete values. Similarly, for jitter generation, it is desired to use a square wave with transitions constrained to the clock rising edges to represent an FM signal with zero crossings which fall in between clock events. There is thus an exact mapping between the two problems where one domain is voltage or current level and the other is time (phase). It is therefore natural to examine how the quality of digital-to-analog conversion with coarse quantizers can be improved and see if any solution can be adapted to jitter generation. It turns out that the technique named delta-sigma ($\Delta\Sigma$) modulation is suited for this purpose [47]. It is described in details in Appendix A. In a nut shell, the error which occurs because of quantization at a given time is taken into account for the following quantization operations. If the bandwidth of the device using the signal is much lower than the rate at which the quantizations are performed, then the adjustments to the errors make the quality of the signal almost as good as if no quantization occurred. $\Delta\Sigma$ modulation is also referred to as noise-shaping because, in the frequency domain, the power from quantization error is located mostly at high frequencies while the signal is low frequency.

$\Delta\Sigma$ modulation may be applied to jitter generation in the following way [48]. For each PLL cycle, a value for the instantaneous phase is quantized to a clock edge. The error created is recorded and will influence the phase of later PLL cycles. The use of $\Delta\Sigma$ modulation on phase is pictured in Figure 4.5. The vertical dashed lines are rising edges of the clock signal with the thicker ones defining the zero reference for phase in each cycle. Six different cy-

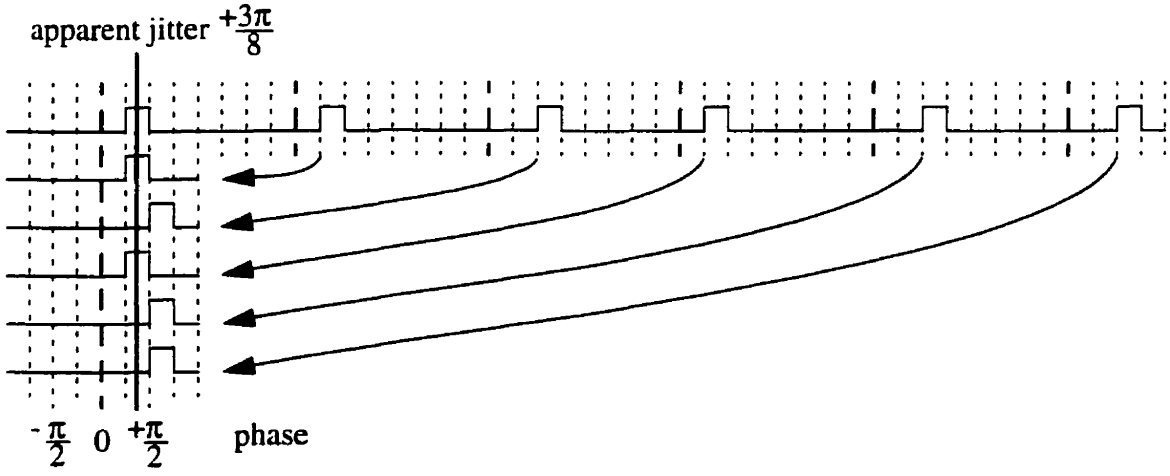

Figure 4.5: Delta-sigma modulation of jitter.

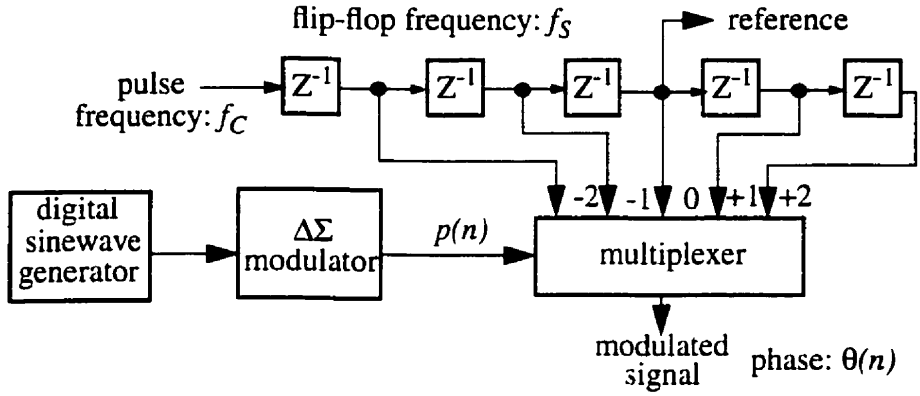

Figure 4.6: Digital phase modulation circuit.

cles of a signal where transitions are restricted to the rising edges of the clock are shown at the top of the figure. Stacking them vertically while aligning the reference phase, the principle of $\Delta\Sigma$ phase modulation may be visualized. Here, by toggling the phase between $\pi/4$ and $\pi/2$ radians, on average a jitter of $3\pi/8$ radians is created. This averaging is performed because the bandwidth of the PLL which will use this signal is much smaller than the rate of phase switching. A similar principle was demonstrated for the generation of clock signals with programmable frequency [49]. This scheme is not limited to constants but may also generate complex signals such as sinewaves.

Figure 4.6 illustrates a circuit that can realize this scheme. The signal at the upper left is a pulse whose period is an integer multiple of the test clock period. Using a string of registers, denoted z^{-1}, signals with different phases are created. The output of one of these registers is arbitrarily denoted *reference* and is assigned an index of 0. These signals are routed to a multiplexer where a multi-bit digital signal, denoted[1] $p(n)$ selects the desired phase. This signal is labelled the phase index signal and is updated at every PLL cycle. Because $p(n)$ is encoded using a $\Delta\Sigma$ modulator, the phase of the output signal, denoted $\theta(n)$, will be noise-shaped.

1. $p(n)$ is the time-domain representation. $P(z)$ is also used to refer to the same signal but in the frequency domain. Other variables will also abide by this notation.

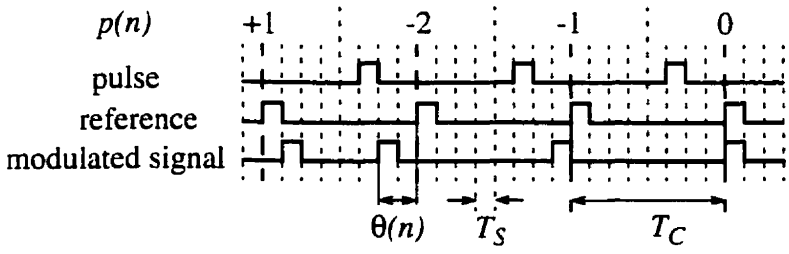

Figure 4.7: Typical digital phase modulated circuit waveforms.

The operation of this circuit can be better understood by examining its signals. Figure 4.7 shows typical waveforms in this circuit for a test clock (f_S) 8 times the reference signal frequency (f_C). In fact, the ratio of the two frequencies is denoted as the oversampling ratio D. T_S and T_C are the periods of the test clock and the reference signal, respectively. The dotted lines represent the test clock edges with the thick ones indicating zero phase references. It can be seen that the reference signal is delayed by a constant three clock periods with respect to the pulse signal. On the other hand, the positions of the pulses of the modulated signal are dictated by the values of $p(n)$ in each cycle. Specifically, the phase index signal $p(n)$ represents the number of clock cycles the output signal is delayed with respect to the reference signal. It is limited to the range $-D/2$ to $+D/2$ as will be explained later.

For each PLL cycle, the time delay relative to a pulse at the reference position is therefore $p(n) \cdot T_S$. The instantaneous phase of the output signal is then expressed as

$$\theta(n) = p(n)T_S\frac{2\pi}{T_C} = 2\pi p(n)\frac{f_C}{f_S} = p(n)\frac{2\pi}{D} \text{ (rad).} \qquad (4.6)$$

It is obvious from the example in Figure 4.7 that the possible positions of the pulses in the modulated signal are very limited within a period. This observation raises questions about the significance of any test result obtained with such a coarse signal. However, PLLs are frequency selective with respect to jitter. In fact, the jitter transfer function is lowpass and the device filters high-frequency jitter. When incorporating a $\Delta\Sigma$ modulator in the signal generator, quantization noise is shaped to high frequencies. In fact, the encoded signal from a lowpass $\Delta\Sigma$ modulator, such as $p(n)$ in Figure 4.7, contains a high-quality low-frequency sinewave and high-frequency quantization noise as illustrated with the top curve of

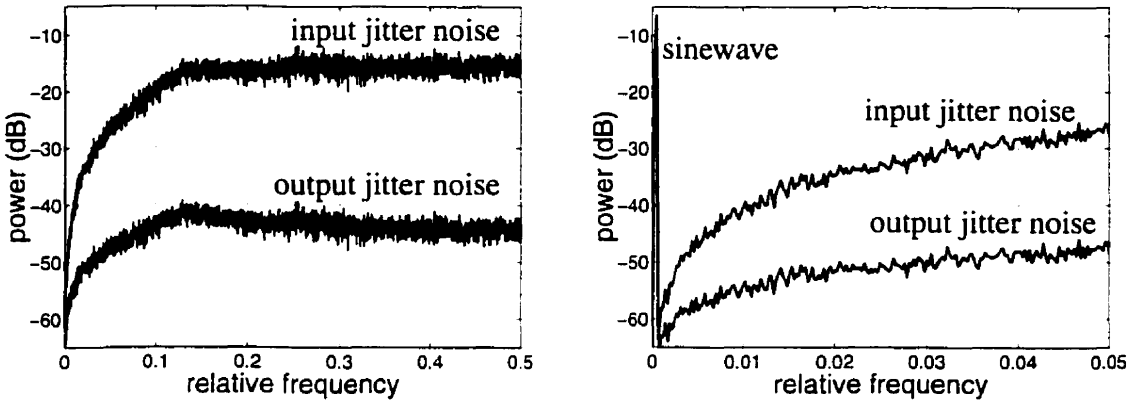

(a) (b)

Figure 4.8: Phase power density spectrums: (a) Nyquist interval. (b) Signal band.

Figure 4.8 (a). Because the PLL is lowpass and is designed to reject high-frequency components in the loop, the quantization noise will be mostly filtered as shown with the lower curve. Figure 4.8 (b) shows the signal band in more details. A similar filtering principle was demonstrated for a voice CODEC, another analog lowpass circuit [22]. Because of the low noise level, the output phase of the PLL may be considered to be a pure sinewave. This is of course an approximation that will be verified later. Section A.5 of the Appendix examines in details how, practically, a noise-shaped sinusoidal signal can be generated on chip.

4.1.4 Improved Delta-Sigma Digital Phase Modulation

In the circuit of Figure 4.6, $p(n)$ may vary between $-D/2$ and $+D/2$ as larger values could result in collisions with pulses from adjoining cycles. The digital phase modulation is thus limited to jitter values between $-\pi$ and $+\pi$. Fortunately, a different circuit may be used to generate arbitrarily large jitters. However a fundamental limitation, causality of edges,

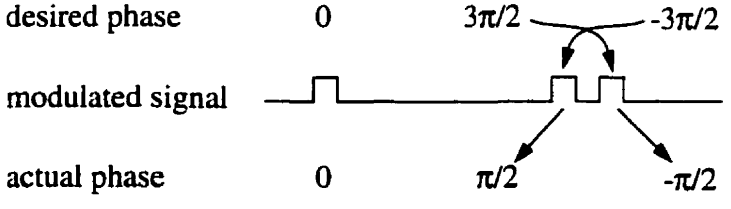

Figure 4.9: The effect of non-causal pulses.

must still be respected. That is, the pulses must arrive in the correct order and not at the same time. A violation of this rule is illustrated in Figure 4.9. The desired phase at each cycle is displayed on the top line with the resulting modulated signal below. However, pulse arrival must be causal and the actual phase as experienced by the PLL is different. In Figure 4.7, the pulse signal is active for the duration of the clock period, T_S. At the output, pulses should not be produced in consecutive clock periods or they will be meshed together in a large pulse, occluding one rising edge. On the other hand, if pulses of half the clock period (return-to-zero pulses) are used then the two pulses may appear in consecutive clock cycles. This will be the case from now on.

Using simple modulo arithmetic, the aforementioned condition may be stated in a more formal language. The causality of pulses forces $p(n+1)$ to be larger than $p(n)$-D+1. This condition must be taken into account when generating $p(n)$. Figure 4.10 shows a block diagram of the improved phase modulation method. The two left most blocks are also found in the previous scheme except that the $\Delta\Sigma$ modulator now enforces the constraint. The last two blocks replace the multiplexer implementation which would be too bulky here as it would extend over many PLL periods.

The diagram of Figure 4.11 shows an implementation of the last two blocks. The operation is fairly simple. From the phase index signal $p(n)$, the signal $d(n)$ corresponding to the delays between the pulses is computed. After a pulse is generated, a new value of $d(n)$ is loaded into the delay register. When the content of the delay register is not zero, the delay register multiplexer is set in reverse counting mode and the modulated signal is idle. The delay register is decremented until it reaches zero. At this point, three events are triggered. First a new value for $p(n)$ is fetched. Then the multiplexer at the input of the delay register is set to load a new value of $d(n)$. Finally, a pulse is generated and incorporated in the mod-

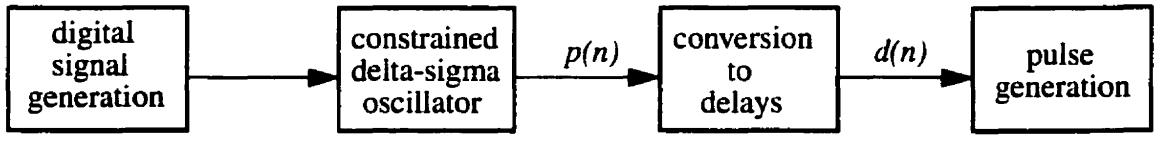

Figure 4.10: Improved phase modulation method.

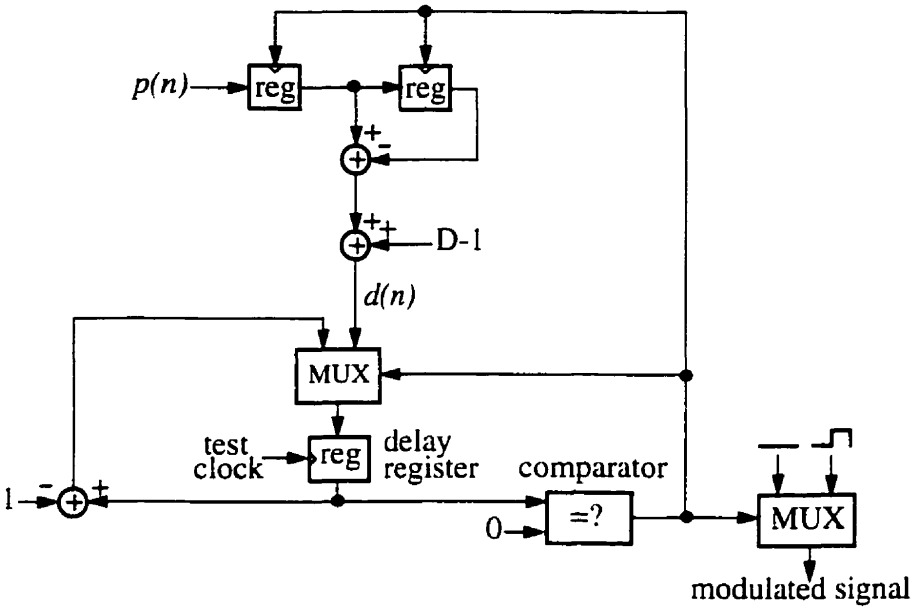

Figure 4.11: Delay extraction and pulse generation circuits.

ulated signal. Note that there is no explicit zero phase reference anymore as the position of each pulse is relative to the previous one.

Figure 4.12 shows signals that could be obtained from the circuit of Figure 4.11 for a value of D equal to 4. It can be seen that $d(n)$ represents the number of clock cycles the modulated signal is idle between two pulses. There is however an important implementation issue with this scheme. The circuit generating $p(n)$ must be capable of supplying data within a test clock period, T_S, to satisfy the circuit of Figure 4.11 while the circuit of Figure 4.6 consumes only one value of $p(n)$ for each PLL period, T_C. At first glance, it may thus appear that the circuit generating $p(n)$ in Figure 4.11 must now be D times faster. However, it should be pointed out that on average, this circuit will require an input value at a rate equal to the PLL frequency. In practice, the circuit generating $p(n)$ will be operated at the PLL

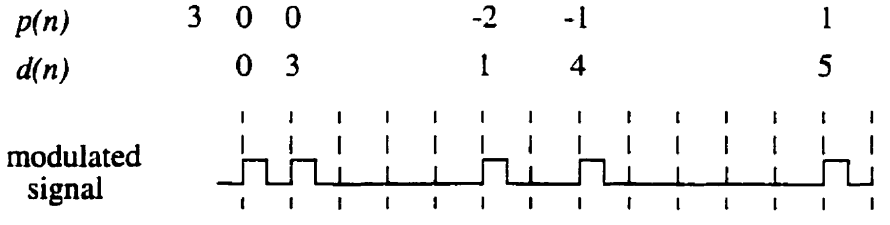

Figure 4.12: Example signals from improved phase modulation circuit.

frequency, f_C, and the variable input data rate of the circuit of Figure 4.11 will be accommodated with a first-in first-out (FIFO) buffer.

The phase index signal, *p(n)*, is an arbitrary waveform, generally a sinewave, encoded with a slope-limited delta-sigma modulator. This device, described in Appendix A, Section A.4, is a $\Delta\Sigma$ modulator which implements the constraint discussed above. Except for the use of a special delta-sigma modulator, the generation of *p(n)* may be performed using any of the three methods explained in Section A.5 of the Appendix. For the methods where *p(n)* is computed off-chip, the delays, *d(n)*, defined as

$$d(n) = p(n) - p(n+1) + D - 1 \tag{4.7}$$

may be stored on-chip instead of *p(n)*. The area of the circuit will be reduced as *d(n)* requires less bits for storage and because of the removal of the arithmetic operations shown in the upper part of Figure 4.11. Furthermore, the limitation on the amplitude of *d(n)* introduced by the finite number of bits allocated for its storage can be expressed as another constraint, in addition to causality, on the operation of the $\Delta\Sigma$ modulator.

The number of quantization levels can now be very large. In fact, it is only limited by the stability of the $\Delta\Sigma$ modulator operating with the above constraints. This topic is covered in Appendix A, Section A.4. It should be noted that the quantization noise power is approximately constant when going from the circuit of Figure 4.6 to this new implementation if the quantization step remains the same. Therefore as the signal amplitude increases, so does the signal-to-noise ratio (SNR) and better test results are obtained.

4.2 Modulation of the PLL Feedback

The second access point for stimulus injection in the PLL is the output. It may seem odd to actually use an output terminal to stimulate a device but it should be reminded that this circuit is a closed loop. Therefore, feedback ensures that every block is stimulated by the test signal.

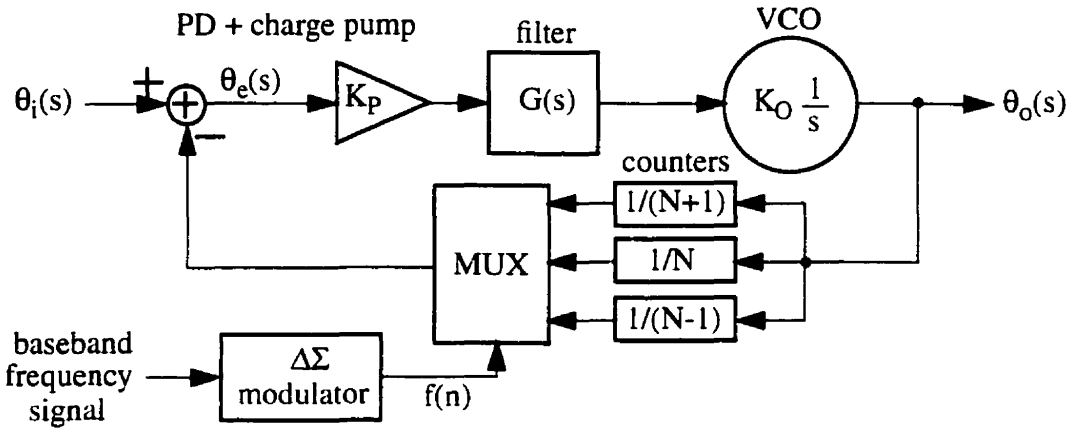

Figure 4.13: Delta-sigma fractional-N synthesizer.

Since the output signal is a jittery digital signal, the addition of a stimulus is performed much like that for the input signal as described in Section 4.1. A variable-length counter implements a programmable delay allowing jitter to be added or subtracted. A block diagram of the resulting PLL is illustrated in Figure 4.13. A baseband signal encoded with a $\Delta\Sigma$ modulator selects the counter length and therefore the instantaneous delay for each PLL cycle. This circuit was introduced a few years ago for frequency synthesis under the label of delta-sigma fractional-N synthesizer [50]. In that application, the purpose of the variable length divider is to allow the synthesis of sinewaves whose frequencies are not limited to integer multiples of the input frequency. Therefore, in synthesizers, a counter with two lengths is sufficient. Consequently, the $\Delta\Sigma$ modulator which encodes the baseband signal will have a single-bit output for toggling between the two lengths. On the other hand, for clock distribution applications the output frequency will be an integer multiple of the input frequency. In this case, two supplemental counter lengths, one larger and one smaller are necessary to allow jitter with both positive and negative values to be added. This requirement that both positive and negative delays be produced imposes the constraint that a counter must already be present in the PLL. Otherwise only positive values of jitter can be added and thus sinewaves may not be injected without an undesired DC component.

Modulation of the PLL feedback is equivalent to digital modulation of the PLL input as illustrated in Figure 4.14. For simplicity, only the phase detector and signal injection point

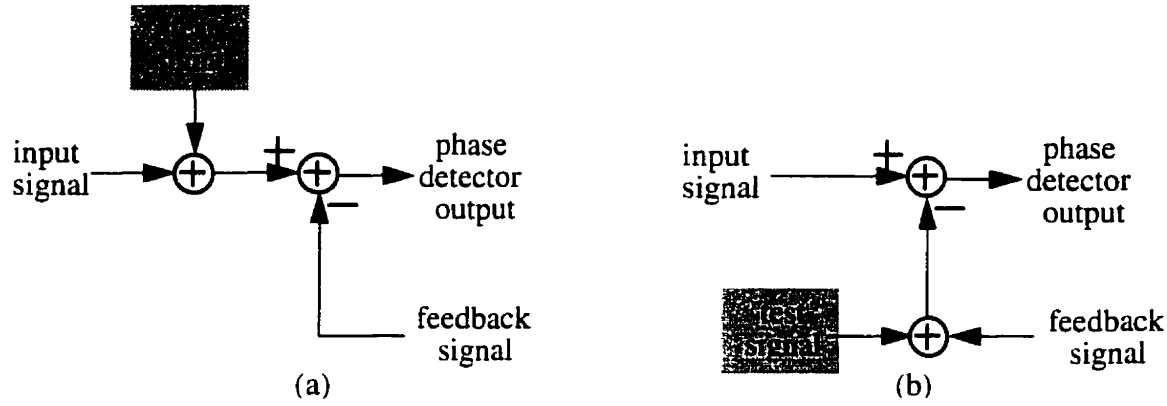

Figure 4.14: Comparison of digital modulation methods: (a) input. (b) feedback.

are shown. Obviously, test signals injected in the input signal or in the feedback path will have the same effect on the circuit except for a negligible sign change. Therefore, results obtained with any one of these methods applies to the other. They are interchangeable from a theoretical point of view. Practically, however, digital modulation of the input signal necessitates an external reference (clock) frequency of a much higher frequency. On the other hand, it does not require the presence of a divider in the feedback loop.

Furthermore, the modulation of the PLL feedback is a non-linear operation and distortion may result. This comes about because the quantization step in the time-domain, as set by the counter in the feedback path, is the output signal period, denoted T_o. This value is not constant but rather depends on the output frequency, f_o. On the other hand, the time quantization step at the input is the clock period, T_S, which is a constant. An expression can be obtained linking T_o to the output jitter. First, the output signal, s_o, is defined as

$$s_o(t) = \mathrm{sgn}(\sin[2\pi N f_C t + \theta_o(t)]), \tag{4.8}$$

where f_c is the frequency of the reference input and N is the average value of the counter length. The output period is defined as

$$T_o(t) = \frac{1}{f_o(t)} = \frac{2\pi}{\frac{ds_o(t)}{dt}}. \tag{4.9}$$

Using Eq. (4.8) in Eq. (4.9), the following expression is obtained,

$$T_o(t) = \frac{1}{Nf_C + \frac{1}{2\pi}\frac{d\theta_o(t)}{dt}}. \tag{4.10}$$

It should be noted that the frequency deviation is much smaller than the center frequency of the VCO, a relation denoted as

$$\frac{1}{2\pi}\frac{d\theta_o(t)}{dt} \ll Nf_C. \tag{4.11}$$

Therefore Eq. (4.10) simplifies to

$$T_o(t) \approx \frac{T_C}{N}\left(1 - \frac{1}{2\pi N f_C}\frac{d\theta_o(t)}{dt}\right). \tag{4.12}$$

The second term in Eq. (4.12) is the relative frequency deviation. Its maximum value is defined as the modulation index in FM signals. While insight is difficult to grasp with the previous equations, their behavior can be understood using simulations. Figure 4.15 shows the power density spectrum, obtained with the help of simulations, of the output jitter, $\theta_o(s)$, at low frequencies. The bottom curve was obtained with digital phase modulation of the input while the top curve represents modulation of the feedback. Although not shown on this graph, the two curves are identical at larger frequencies. It can be seen that the second method results in a higher noise floor at low frequencies and that the second harmonic is visible.

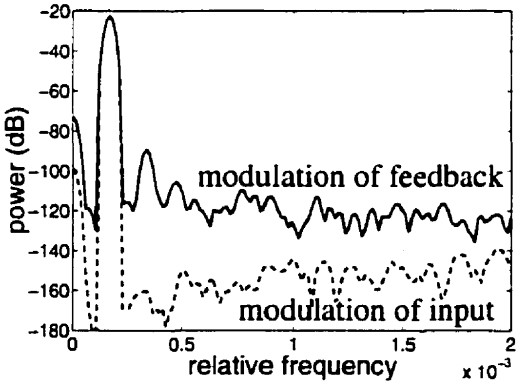

Figure 4.15: Phase power density spectrums of two signals using different phase modulation techniques.

Nevertheless, this does not degrade significantly the SNR over the Nyquist band because the noise level at higher frequencies is larger and similar for both. However, here the modulation index of the output signal is a small 0.4%. The effect of the non-linear behavior of the digital modulation of the PLL feedback will become significant for larger maximum frequency deviations.

4.3 Test Signal Injection with a Charge Pump

The last option for the stimulation of a PLL is the injection of a test signal at the input of the loop filter as shown in Figure 4.16. This signal source, represented here by the variable $\theta_x(z)$, is injected through a second charge pump with gain K_X. It should be understood that this signal source is not a jittery digital signal but an analog signal represented by a voltage or a current. However, this signal, when referred back across the phase detector is equivalent to input jitter. Meanwhile, the PLL input signal, also called the reference signal, is a constant frequency square wave and whereby the input jitter $\theta_i(z)$ will be zero.

A test with this signal injection method is not theoretically a functional test as access to an internal node is required as well as some knowledge about the charge pump implementation. However the measures obtained can be directly related to PLL metrics. This setup will be used to evaluate the characteristics of the PLL through the measure of the transfer func-

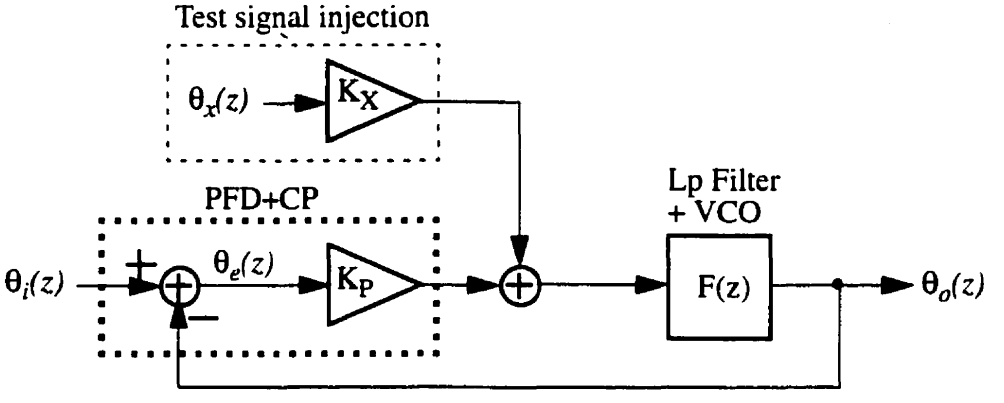

Figure 4.16: Discrete-time model of phase-locked loop modified for signal injection.

tion from $\theta_x(z)$ to $\theta_o(z)$. Examining the model of Figure 4.16, it can be seen that this transfer function will be equal to

$$\frac{\theta_o(z)}{\theta_x(z)} = \frac{K_X}{K_P} \cdot \frac{K_P \cdot F(z)}{1 + K_P \cdot F(z)}. \tag{4.13}$$

It is thus equivalent within a multiplicative constant to the jitter transfer function defined previously in Eq. (3.1).

The test signal injection is simpler than it may seem at first thought. A second current charge pump is placed in parallel with the phase detector charge pump and both outputs are tied in a current summing node as shown in Figure 4.17. In this schematic, the impedance Z_F implements the loop filter and V_C is the controlling voltage of the VCO. The accuracy of this analog-to-digital conversion is a function of the matching of the two current sources I_P and I_X, typically ranging between 0.1% - 1% in monolithic form. As it will be shown later, the charge-pump mismatch will lead to a jitter gain offset which can be subtracted from the results. Another method for nulling a mismatch would be to repeat the test with the input of the charge pumps interchanged. That is, what used to be the test charge pump

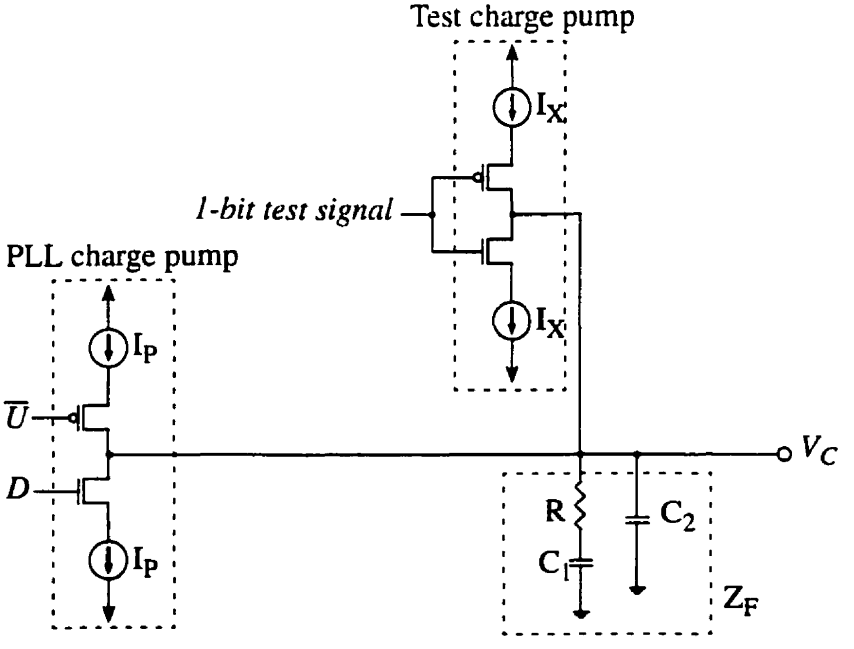

Figure 4.17: Circuit modifications for signal injection in the phase-locked loop.

is now connected to the phase detector output while the PLL charge-pump is used for test signal injection. This has the effect of averaging the mismatch error.

The test signal needs to be encoded on a single bit and this may be accomplished by ΔΣ modulation (see Appendix A). For a spectral test such as the jitter transfer function, the desired test signal is a sinewave. The one-bit test signal will thus contain a very high-quality sinewave and high-frequency quantization noise. Most of this noise is filtered out by the inherent lowpass characteristics of the PLL as explained in the section 4.1.3.

It is important to note that the clock period for signal injection must be an integer multiple of the reference signal period to prevent aliasing of the quantization noise back in the PLL passband. This condition implies that the signal injection frequency cannot be higher than the reference signal frequency. Because of this limitation on the signal injection frequency, simplification of the scheme through the use of a single charge pump and time-sharing is not possible. Therefore, in applications which cannot afford the extra charge pump or require high accuracy, another method should be selected.

4.4 Maximum Jitter Amplitude in the Presence of Quantization

A PLL has a limited range of input frequencies. This is usually set by the tuning capability of the VCO and is referred to as the lock range. Any PLL test must therefore ensure that the frequency of the stimulus signal remains within the lock range of the DUT. This restriction is more of a concern for noise-shaped signals such as those used in the jitter generation

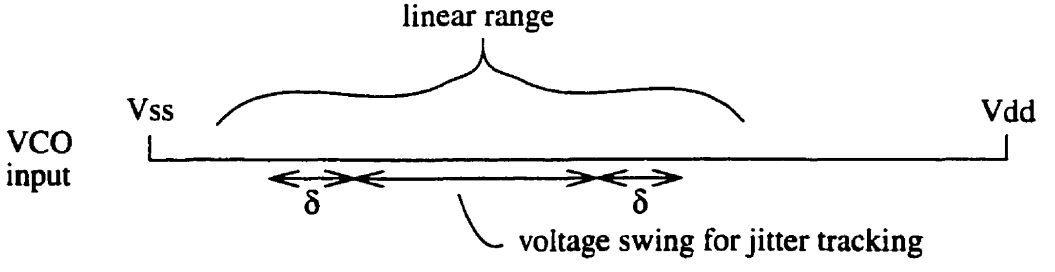

Figure 4.18: Effect of quantization noise on the signal swing of sensitive nodes.

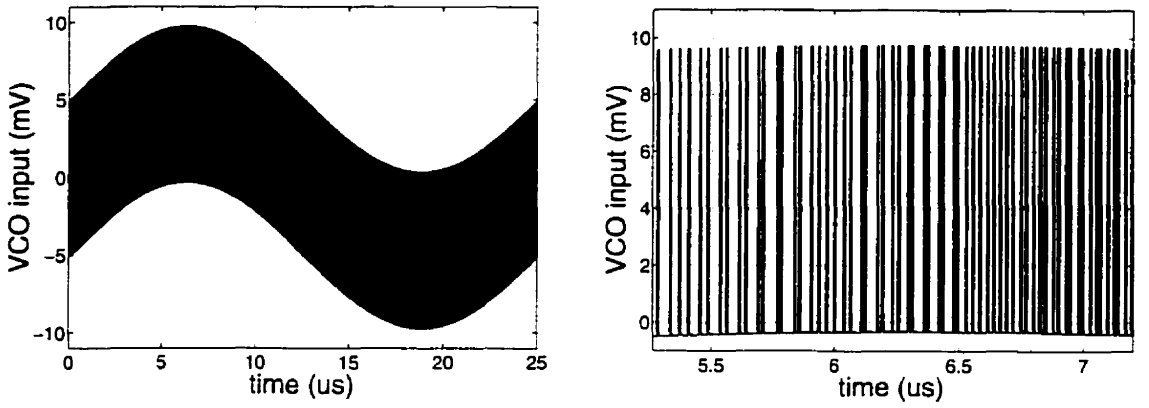

(a) (b)

Figure 4.19: Typical signal at the input of the VCO: (a) One cycle. (b) Zoom.

schemes of this chapter. At many PLL nodes, the residual quantization noise in the test signal will increase the signal swing. Here, the residual quantization noise refers to the input signal noise that is not attenuated by the loop dynamics. To avoid meaningless test results, this effect must be accounted for so that all signals in the PLL remain within their linear range.

Figure 4.18 illustrates the situation for a critical PLL node, the input of the VCO. The voltage at this node will vary to allow the VCO to track the jittery signal. However, quantization noise will increase the voltage swing by an amount δ at each end. The test designer must make sure that the resulting voltage deviations remain within the range of input to the VCO for linear operation. A typical signal that could be found at the input node of a VCO for a PLL driven by a $\Delta\Sigma$ modulated signal is shown in Figure 4.19. It was obtained from simulations of a PLL with the passive filter described in Figure 3.4 (b) on page 22 and the parameters listed in Table 3.1 on page 27. Part (a) illustrates a full cycle while part (b) shows in more details the time interval where the maximum is located. The signal consists of a high-frequency square wave superimposed on a sinewave. While the sinewave amplitude is only 9.5 mV peak-to-peak, the quantization noise is pushing the voltage range to almost 20 mV peak-to-peak.

To obtain an expression for the added voltage swing, δ, it should be noted that the maximum disturbance will occur when the sinewave is at its maximum or minimum and a quan-

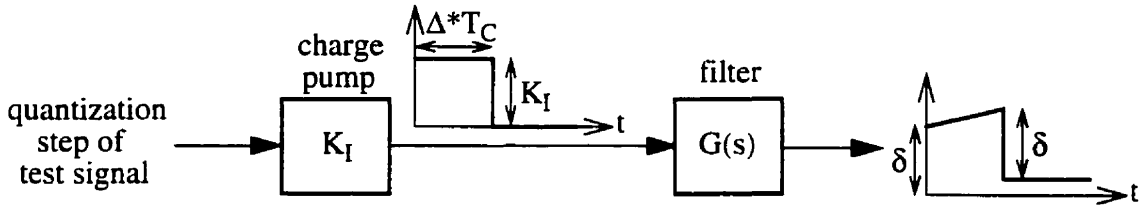

Figure 4.20: Open loop model for voltage swing due to quantization.

tization step of relative value Δ is added. This event is contained within a single period of the VCO signal. Therefore, as the VCO output is constant, the circuit can be analyzed in open loop configuration. Figure 4.20 illustrates the model. A quantization step in the test signal is turned into a pulse of duration Δ times the period of the phase detector, T_C. The pulse amplitude is denoted K_I, representing the gain of the test signal injection circuitry. It will generally be equal to K_P, the gain of the phase detector and charge-pump circuits of the PLL except when an extra charge pump injects the test signal (see Section 4.3). Finally, this pulse is shaped by the PLL loop filter.

The input to the filter, $x(t)$, is the sum of two steps of equal magnitude, one of which is inverted and delayed. This signal is thus expressed as

$$x(t) = K_I[u(t) - u(t - \Delta \cdot T_C)]. \tag{4.14}$$

The output of the filter, $y(t)$, will be its convolution with the filter impulse response, $g(t)$, and is denoted

$$y(t) = x(t) \bullet g(t). \tag{4.15}$$

This signal will have two discontinuities at $t = 0$ and at $t = \Delta \cdot T_C$. The magnitude of δ is the amplitude of the discontinuities in the output signal. It will thus be proportional to the immediate response of the filter to a step function. This is denoted in the time domain as

$$\delta = \lim_{0 \leftarrow t} x(t) \bullet g(t) = \lim_{0 \leftarrow t} K_I u(t) \bullet g(t) \tag{4.16}$$

The equation may be reformulated in the frequency domain resulting in

$$\delta = K_I G(s)|_{s \to \infty}. \tag{4.17}$$

As an example, this theory can be applied to the simple passive RC filter of Figure 3.4 (b) with test signal injection using the PLL charge pump. This could be, for example, the method of digital phase modulation of the input. The transfer function of the filter is

$$G(s) = R + \frac{1}{sC}. \tag{4.18}$$

Using this function into Eq. (4.17), the following expression is obtained for the disturbance

$$\delta = K_P\left(R + \frac{1}{sC}\right)\Big|_{s \to \infty} = K_P R. \tag{4.19}$$

For the values quoted in Table 3.1 on page 27, the maximum bound on δ is found to be

$$\delta = 8\text{x}10^{-7}\ \text{A/rad x } 2\pi \text{ x } 1\ \text{k}\Omega = 5\ \text{mV}. \tag{4.20}$$

This value is consistent with the simulation results shown in Figure 4.19.

4.5 Summary

Various techniques for the generation of square waves with controlled sinusoidal jitter were presented. First an analog solution consisting of an FM generator and a comparator was presented. Direct digital synthesis methods were then explored. It was argued that these techniques, used in analog instruments and in automatic test equipments, are not suitable for on-chip implementation. Therefore, three methods for stimulating PLLs were introduced. The first one is dubbed digital input phase modulation using delta-sigma modulation. It is non-intrusive but requires a test clock frequency at least four times the PLL frequency. The second one is the modulation of the signal in the feedback path of the PLL. It requires a lower external reference frequency but the PLL must have a counter in the feedback loop. Finally, the last method proposed is based on injecting a test signal directly at the input of the loop filter in the PLL. While it requires an extra charge pump, it allows testing the PLL using a test clock of the same frequency as the reference signal. Finally, the limitation on jitter amplitude because of quantization noise was derived.

Chapter 5 : Response Analysis

The second component of any test, after the generation of a stimulus, is the measurement of the response from the device-under-test (DUT). Generally, evaluating jitter involves measuring the instantaneous period of a signal with a counter. However, as the operating speed of devices get closer to the test clock controlling the counter, the accuracy achieved is not sufficient. This chapter will examine various schemes for the analysis of jitter in a digital signal.

5.1 Direct Jitter Quantization

The measurement of time intervals, and thus jitter, is traditionally performed with a counter. This device records the number of rising edges of the clock in a time span. Obviously, the accuracy is then limited by the period of the counter. When sampling a signal with a test clock, a transition occurring at any time between the two clock events will be sensed on the second one. The mapping between the instantaneous phase, $\theta(t)$, and the measured phase, $\tilde{\theta}(t)$, is therefore not one-to-one as depicted in Figure 5.1. The quantization step, Δ, is defined as

$$\Delta = \frac{2\pi f_C}{f_S} = \frac{2\pi}{D} \text{ (rad)}, \tag{5.1}$$

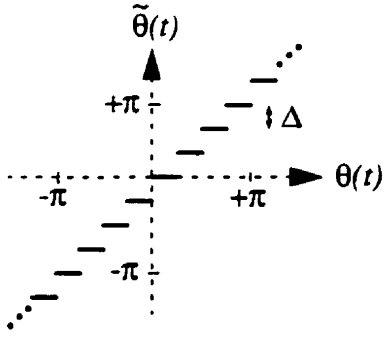

Figure 5.1: Quantization of jitter in sampling measurements.

where f_S is the test clock frequency and f_C is the PLL reference frequency. D is equal to f_S/f_C and is denoted the oversampling ratio. It assumes a value of 8 in Figure 5.1 as 8 quantization steps appear in a 2π interval. The quantization step is thus set to $\pi/4$.

The problem of measuring the jitter in a signal is thus similar to analog-to-digital conversion. The theory developed for this latter application may be used to evaluate the quality of measures obtained. Different factors such as the signal amplitude, the quantization step and the number of points recorded will have an influence on this. Under reasonable assumptions, the quantization noise power, P_N, is equal to

$$P_N = \frac{\Delta^2}{12} = \frac{1}{12D^2}. \tag{5.2}$$

For a sinusoid with amplitude A, the signal power, P_S, will be

$$P_S = \frac{A^2}{2}. \tag{5.3}$$

Hence, the signal-to-noise ratio (SNR) is

$$SNR = \frac{P_S}{P_N} = \frac{6A^2}{\Delta^2} = 6A^2D^2. \tag{5.4}$$

A large value for A or D is necessary to yield a decent measurement accuracy. Figure 5.2 shows the measured SNR with respect to the signal amplitude for two different values of

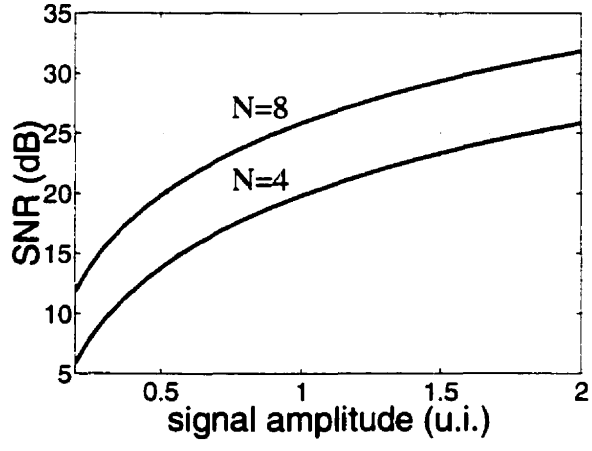

Figure 5.2: Signal-to-noise ratio of measured jitter signal.

clock ratio. The acceptable minimum measurement SNR is defined by the accuracy desired. Typically, a 40 dB SNR is aimed for. To achieve this, a solution is to increase D which is equivalent to increasing the test clock frequency. However, for a fast PLL, a large value of D may not be an option. Yet, because the signal is narrowband, some techniques may be used to somewhat improve the quality by isolating the signal from the noise using a filter or a fast Fourier transform (FFT) but these solutions are computationally expensive.

If an FFT is used, then a multi-tone test can be performed. In a multi-tone test, sinewaves of different frequencies are summed and simultaneously applied to the input of the device under test. The FFT is performed on the output to extract the PLL gain at each sinewave frequency. Thus, with a single input signal, the transfer function may be measured at many frequencies. However, there is a practical issue with concurrent FFT processing. The output of the counter and thus the input to the FFT is a variable-rate signal whose frequency is equal to the instantaneous frequency of the PLL. However, the average rate of this signal is the PLL input signal carrier frequency and on-chip processing may be performed with the help of a first-in first-out (FIFO) buffer. Alternatively, this signal may either be stored on the chip with a RAM and later processed or dumped to the tester. Another solution is to send the digital data directly to the tester but it then requires a high-speed link.

In conclusion, direct measurement of jitter using a counter is not possible for fast PLLs because the test clock frequency can not be made significantly larger than the PLL frequency. However, it is feasible for low-frequency PLLs where it would allow fast testing using multi-tone measurements.

5.2 Analog Phase Measurement

The accuracy achievable with measurements of jitter using a counter is not sufficient. Methods to increase the accuracy beyond the clock period should thus be investigated. In the PLL itself can be found a device that performs time measurements. By definition, the phase-frequency detector (Figure 3.2) compares two signals and asserts one of the two output lines for a duration equal to the time difference between the rising edges of the input

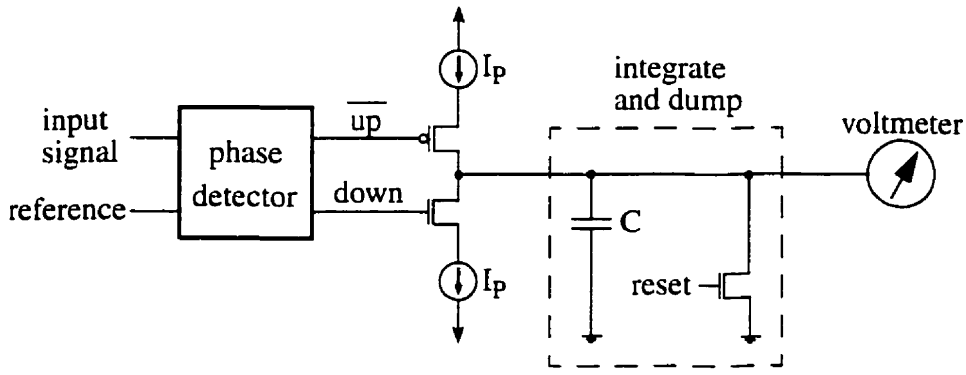

Figure 5.3: Analog phase measurement.

signals. To convert this time difference to a voltage, a charge pump and a filter are used. Figure 5.3 illustrates an analog time interpolator based on this principle. The capacitor integrates the current at the output of the charge pump and the time difference may then be read as a voltage. Once a measure is performed, the capacitor is discharged to allow for the next measurement. The limitations of this measuring scheme are obvious. First, it suffers from integral non-linearity because the absolute values of the current sources and the capacitor are highly dependent on process variations. Furthermore the reset phase in which the capacitor must discharge severely limits the rate of operation. However, analog phase measurement has been used to measure static jitter with a filter for averaging in place of the integrate and dump circuit [51].

5.3 Digital Time Interpolation

The limitation of the clock period may be more easily overcome using digital measurements. The circuit which allows this feat is the calibrated delay line [52]. A typical delay

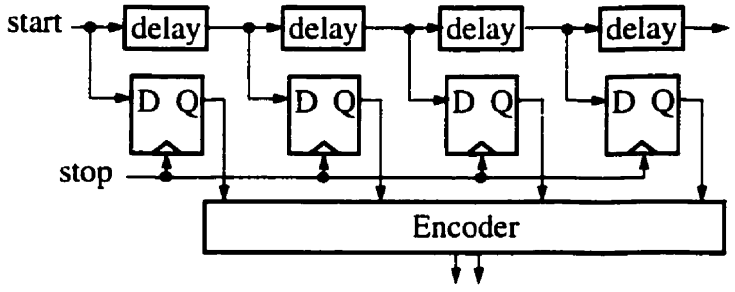

Figure 5.4: Time digitization using a delay line.

line for the digitization of short time interval is shown in Figure 5.4. The stop signal is used to capture the position of the start signal propagating through the delay line. Because the delays are process and operating conditions dependent, the delay line must be continuously calibrated, an operation usually involving a delay-locked loop. Yet, matching errors will cause significant differential non-linearity. Furthermore, digitization of small time intervals using delay lines demands significant silicon area and requires very careful layout. Calibrated delay lines have reportedly been used in automatic test equipment (ATE) [53]. Nevertheless, the area overhead is too large for this measurement method to be used in a BIST scheme.

5.4 Comparing Jitter With a Threshold

It was shown in the previous sections that the exact gauging of the jitter response $\theta_o(z)$ is thus rather difficult. This is especially true within the limitations imposed by a BIST application. Fortunately, using the test clock, a measure that can be made with good accuracy is whether $\theta_o(z)$ is above or below a phase threshold. However, this analysis strategy requires a modification of the usual measurement procedure. Generally, an input signal is applied to a DUT and the output is measured exactly. Here, for a given jitter frequency, an output jitter level is selected and the amplitude of the input jitter is varied. Ultimately the test signal amplitude that results in the selected output jitter is obtained and used to compute the jitter transfer function. The edges of the test clock, assumed to be jitter free, will be used to establish this threshold. Figure 5.5 illustrates how this can be accomplished. The dashed lines represent the rising edges of the test clock with the bold ones indicating the rising transitions of the reference clock. The output signal is sampled at positive test clock edges until two adjacent samples yield a 0 followed by a 1, indicating a rising edge of the signal. If this

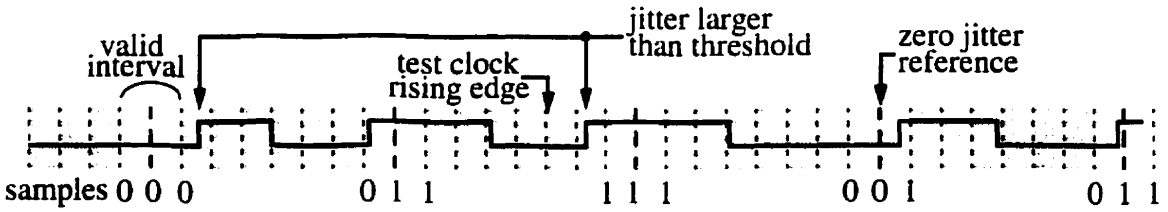

Figure 5.5: Comparing a square wave against a jitter threshold.

rising edge occurs in intervals immediately before or after the reference test clock edge, then jitter is below threshold; otherwise, an error is generated. Returning to Figure 5.5, a vector is composed of the three samples obtained using the positive test clock edges which are closest to the reference clock rising edge. Obviously, when this vector is equal to 000 or 111, it can be concluded that the jitter is larger than the threshold. Over a time interval, the number of errors is counted and a bit error rate (BER) measure can be obtained. This averaging is done to prevent glitches and noise sources from significantly affecting the final result. In Figure 5.5, a ratio of the test clock frequency over the reference signal frequency of 8 allows a minimum value for the phase threshold of $\pi/4$. Larger values could also be used for the threshold by allowing more test clock cycles in the valid interval.

This jitter analysis method requires a test clock frequency at least three times the reference frequency. Three periods of the test clock for a PLL period is the limiting case as two of the clock periods must compose the valid interval, leaving one to catch jitter above threshold. A different scheme however can be used that compares jitter against a π rad threshold using a 50% duty cycle test clock of the same frequency as the reference clock [54]. Should the test clock fail to exhibit a 50% duty cycle then the results will be affected. The severity of the error will depend on the extent of the static and random components of the duty cycle

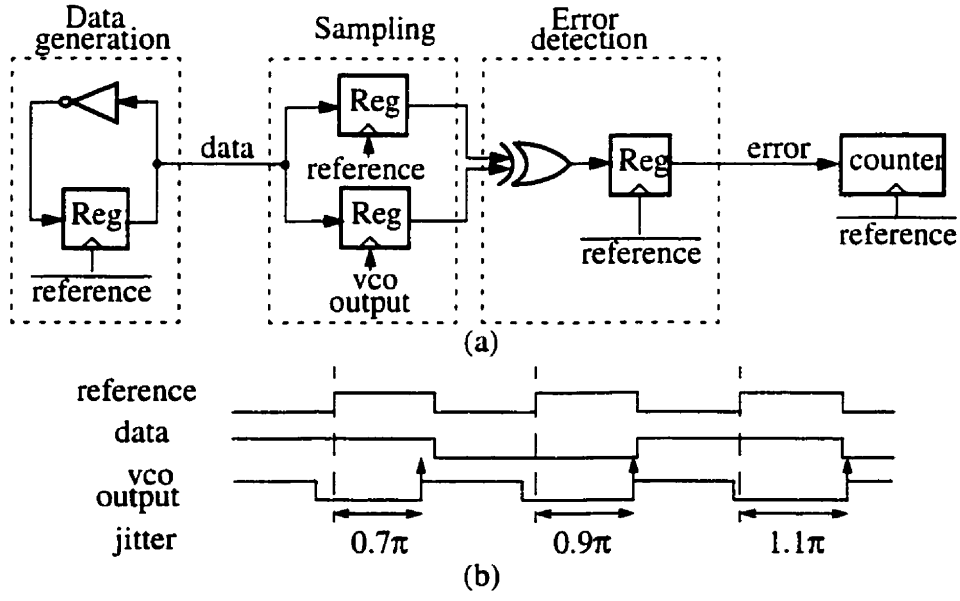

Figure 5.6: (a) Circuit to evaluate a π rad jitter threshold. (b) Typical waveforms.

variations. This topic will be explored in Section 5.5. Figure 5.6 (a) illustrates how it can be implemented with a few gates and registers. To the left, a dummy data signal is first generated with the help of a toggle register clocked with the negative edge of the reference signal. The threshold is then verified by sampling this data signal with both the reference clock, used at the input of the PLL, and the jittery VCO output signal. When the output jitter exceeds π rad, sampling errors will happen because the output of the PLL will slip a data cycle. This is shown in Figure 5.6 (b) where an error occurs when the jitter goes from 0.9π to 1.1π. This circuit is somewhat similar to the circuit used for the jitter tolerance test (see Section 3.4.1).

A single frequency jitter test is performed with the following procedure. An input jitter, $\theta_i(z)$ or $\theta_x(z)$, is applied to the PLL and its amplitude is modified until the maximum value resulting in output jitter below threshold is found. The fastest strategy for obtaining this input jitter amplitude is to use a binary search algorithm. The initial amplitude, denoted A_0, is zero and the initial increment, denoted Δ_0, is 0.5. A test is performed with amplitude $A_0+\Delta_0$. If the VCO output jitter is lower than the threshold then the next amplitude A_1 will be $A_0+\Delta_0$. Otherwise, the amplitude remains the same ($A_1=A_0$). The increment is then divided by two and the procedure is repeated until the desired accuracy is achieved. The measurement uncertainty associated with the signal amplitude following k such steps will be $\Delta_k/2$, or $2^{-(k+1)}$. However, arbitrary accuracy can not be achieved as noise sources are always present. The identification of these sources and the evaluation of their effect is the topic of the following section

5.5 Accuracy of Jitter Threshold

The accuracy of the results obtained with a jitter threshold measurement technique will depend on many factors. The sampling of the PLL output signal by the test clock is the critical action. Figure 5.7 depicts this operation. The output of the PLL is captured at discrete time instants by a flip-flop with the help of a test clock. Six different sources of errors are located on the diagram. These factors may be divided in two categories. The first one includes effects leading to a constant (static) jitter. The second set groups sources of random jitter. As

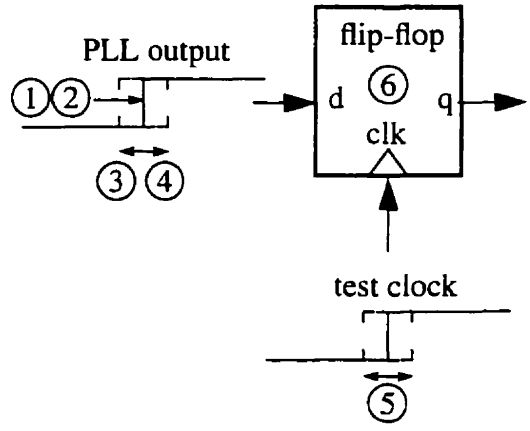

Figure 5.7: Factors contributing to measurement accuracy.

they result in different types of errors in the measurements, they will be analyzed separately.

5.5.1 Sources of Static Jitter Errors

PLL static jitter may be caused by mismatch and offsets in the PLL. Static jitter will add to the jitter from the input signal and thus modify the effective output jitter threshold. Its effect can only be reduced by using a jitter threshold much larger than the jitter noise. Alternatively, as it results in a gain offset independent of frequency, it could be accounted for and subtracted from the final results. This will be discussed in the next chapter. Another important issue is the analog behavior of digital components at high speed. For example, the delay of the flip-flop used to sample the output signal may become significant at very high frequencies. This also leads to static jitter. Finally, for the signal injection method, the matching of the two charge pumps is obviously a cause of errors. While it does not create a static jitter, its effect on test results is similar.

5.5.2 Sources of Random Jitter

The primary source of random jitter at the output is the jitter quantization noise from the input signal which passes through the PLL. While for simplification it has been assumed that the output jitter is a sinusoid, it is in fact corrupted by this noise. A measure of the quality of the output jitter is the signal-to-noise ratio (SNR). A number of parameters will in-

fluence the effective output jitter SNR: the jitter creation clock frequency, the $\Delta\Sigma$ modulator noise transfer function, the number of quantization bits and the PLL order and bandwidth. The effect of the PLL bandwidth and of the number of quantization levels can be seen in Figure 5.8. This graph shows the simulated SNR of the output jitter with respect to the ratio of the cutoff frequency of the PLL over the operating frequency. Curves are displayed for three quantization granularities. The bottom one is a 1-bit signal for the jitter injection technique. The top two curves are multi-bit signals suitable for digital phase modulation. A second-order $\Delta\Sigma$ modulator is employed and the output jitter amplitude is held constant at π. Two obvious conclusions may be drawn. First, the signal quality increases with the number of quantization levels. Secondly, as the loop bandwidth increases, more noise is passed through and the output jitter sinewave quality decreases. For reference, digital data communication systems such as SONET mandate a relative loop bandwidth of about 0.1%. It is interesting to note that for each frequency point of a jitter transfer function, the SNR of the measured output jitter will remain constant. As a proof, first note that, under a white noise assumption for the quantizer error, the noise present in the output jitter does not vary with the signal frequency. Furthermore, as the output jitter threshold is fixed, the output signal amplitude is constant and so is the signal power at that node. The SNR is thus constant over the whole frequency range.

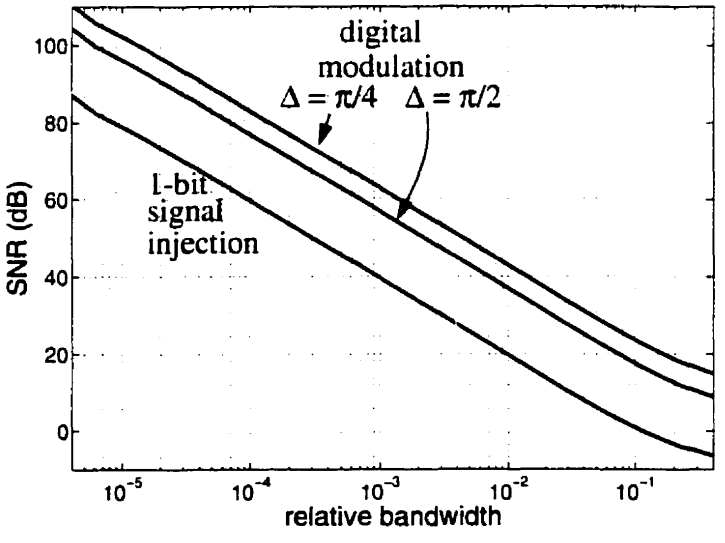

Figure 5.8: Signal-to-noise ratio of the output jitter versus the PLL relative bandwidth.

Another parameter that will affect the accuracy of the results is the amount of jitter present in the test clock. In fact, the timing accuracy is a fundamental limitation of any digital test. However, the clock signal should be generated by a tester and, provided a sound floorplan, this source of error should be negligible.

A more significant problem is the random jitter generated internally by the PLL. Much like static jitter, random jitter from different sources will add to the desired signal and modify the effective threshold. However, because of the randomness, the exact effect on the results may not be predicted. The only available solutions, albeit partial, are to reduce random jitter contributions wherever possible, to increase the jitter threshold and to average out the effects of noise by repeating the test many times.

5.5.3 Effect of Random Jitter on Results Variability

Any measurement will entail some uncertainty. Yet, the accuracy of the measures is a very important factor of test design. The objective of this section is to find the accuracy that can be achieved with the use of a threshold for measuring sine wave amplitude.

First, the signal at the output of the PLL, $s_o(t)$, can be described in the time domain as

$$s_o(t) = \mathrm{sgn}\{\sin[2\pi f_c t + A\sin(2\pi f_j t + \phi_i) + n(t)]\}, \tag{5.5}$$

where f_c is the carrier frequency. A is the jitter signal amplitude while f_j is its frequency and ϕ_i is its initial phase. The jitter signal described by these three parameters is the input jitter stimulus modified by the PLL transfer function. $n(t)$ is a phase noise (jitter) term. This last component is assumed to be unrelated to the signal and stationary. Because the output is a digital signal, only the positions of the rising edges of this signal are of interest. A discrete-time notation may thus be used where k is a time index of these events. They will occur when the expression inside the square brackets evaluates to a multiple of 2π. Now, let the phase of the jitter sinusoidal function be denoted $\phi(t)$. It is equal to

$$\phi(t) = 2\pi f_j t + \phi_i. \tag{5.6}$$

The analysis can be simplified by assuming that $\phi(t)$ is uniformly distributed at the rising edges of the VCO output. This can be justified by noting that the carrier frequency is much

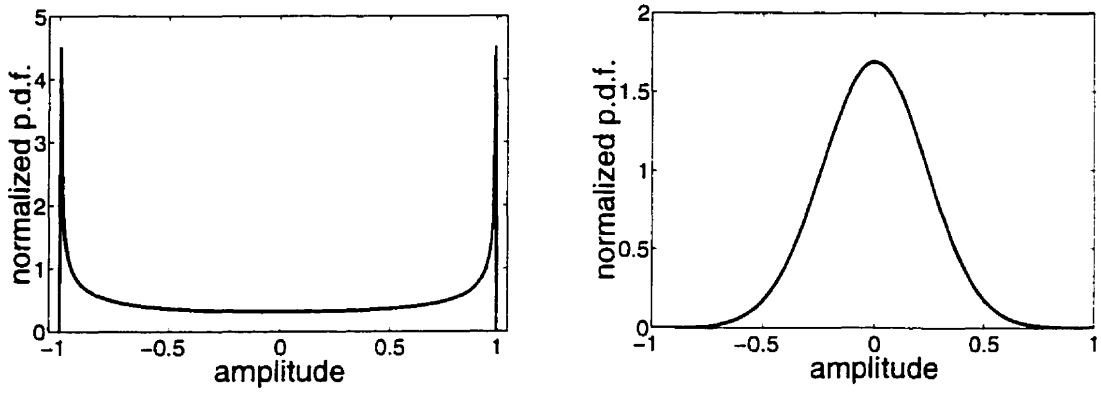

(a) (b)

Figure 5.9: Probability density functions: (a) Sinewave. (b) Gaussian noise.

larger than the jitter frequency and that the test duration encompasses a large number of sinusoidal phase periods. The analysis method consists in counting the number of errors for a given input jitter frequency and amplitude. Now, an error will be recorded when the absolute value of the signal jitter is larger than the jitter threshold, J_T, a condition denoted as

$$|A\sin(\phi[k]) + n[k]| > J_T. \tag{5.7}$$

Note that the time index is irrelevant here because the first component of the left hand side is cyclo-stationary while the second is stationary. Typical probability density functions of the two components are shown in Figure 5.9. The probability density function (pdf) of the left side of the inequality will thus be the convolution of the pdf of a sinewave and Gaussian distribution. An example is illustrated in Figure 5.10

The determination of the jitter transfer function at a given frequency to a l-bit precision involves performing l measures to find the input signal jitter amplitude resulting in the selected output signal jitter threshold. For each measure, an error rate, E, defined as

$$E = \sum_{N=0}^{M} \frac{\{|A\sin(\phi[k]) + n[k]| > J_T\}}{M}, \tag{5.8}$$

is computed where M is the number of PLL cycles in the test. This error rate is compared with a predetermined error level, E_T, and the current bit is set to 1 if it is below that threshold and to 0 otherwise. A bit will be set incorrectly when the measured value of the error

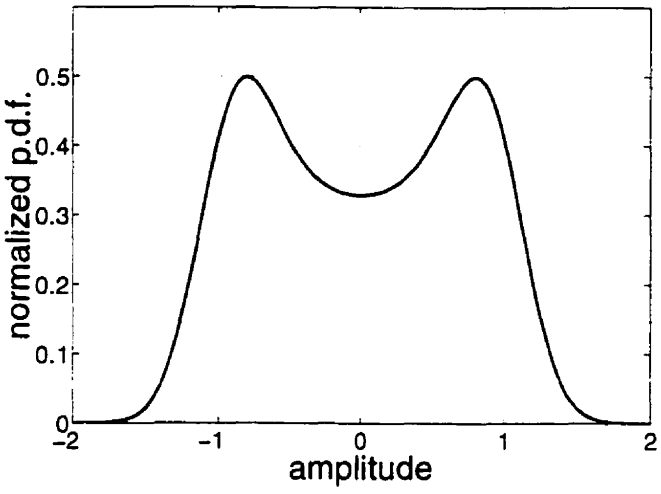

Figure 5.10: Probability density functions of sum of sinewave and Gaussian noise.

rate is on the wrong side of the error threshold. To quantify the test accuracy, the probability of a bit being set to zero, $P(E > E_T)$ with respect to the amplitude must be computed. Since the solution of Eq. (5.8) is very elaborate, Monte Carlo analysis is used. A simulation was thus performed where $P(E > E_T)$ was calculated for various ratios of signal amplitude versus threshold level. The results in the form of a cumulative density function are shown in Figure 5.11 for five different values of SNR. For each curve, two parameters are of interest. The first one is the median of the pdf which is determined by the amplitude at which the

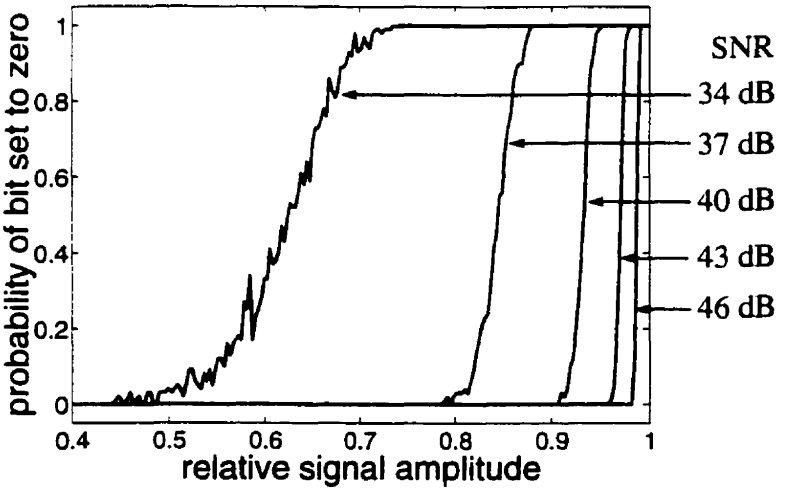

Figure 5.11: Probability of threshold test succeeding versus signal amplitude.

probability is 0.5 on the graph. Since this value is different from unity, where the amplitude is equal to the threshold, it implies that a systematic error will be introduced. For example, since the amplitude leading to 0.5 probability for the 34 dB SNR curve is 0.63, a 4 dB average measurement error can be expected. However, this systematic error will remain constant over the whole frequency range and may thus be accounted for when interpreting the results. Of more concern is the spread of the curves as it leads to a variability in the test results. It is obvious from the graph that the uncertainty will increase with decreasing SNR.

5.6 Summary

Various methods for analyzing jitter in a square wave were presented. It was argued that analog methods can not achieve sufficient accuracy while digital phase interpolation requires too much silicon real estate. The simplest implementation turns out to be a jitter threshold. However, it requires a different methodology. The factors contributing to the accuracy of the measurement were analyzed.

Chapter 6 : Experiments

Experiments are essential in verifying the models relied upon in a synthesis work. This is especially true of PLLs as many simplifications are necessary to obtain the generally accepted model. Furthermore, experimental work is useful in evaluating the extent of second-order effects which could otherwise only be assessed with extensive simulations. Obviously, the ultimate validation of a proposed built-in self-test (BIST) method would be to append it to a commercial product. However, this opportunity does not exist in an academic environment. The alternative is to design the device to be tested before adding the test circuitry. This is clearly not possible within the time frame of a graduate degree. As the design cycle of an IC (conception, fabrication and test) approaches 10 months, errors become very costly. Furthermore, it requires students to master all the subtleties of IC design on top of their core competency as the device under test (DUT) must approach commercial ICs performances. Ultimately, a student could be performing design work from the system level down to the transistor. Clearly, this is not possible. Therefore, in practice, the implementation of system ideas in silicon is restricted to university groups which are vertically integrated. The remaining acceptable solution is to prototype the device using off-the-shelf components on a bread-board. While a discrete-component implementation presents different characteristics from ICs in terms of matching, tolerance, parasitics and speed, the measured data is nevertheless closer to reality than simulation.

6.1 Experimental Set-up

A self-test capable PLL prototype was built on a bread-board to verify the jitter transfer function measurement scheme. An overview of the complete test setup is provided in Figure 6.1. Only off-the-shelf analog components were designed in and populated the bread-board. On the other hand, the digital circuits are coded in VHDL, a hardware descrip-

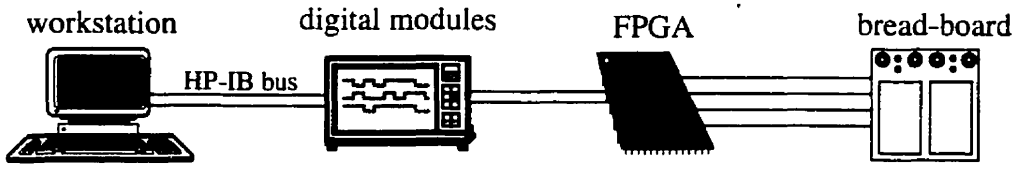

Figure 6.1: Prototyping system hardware.

tion language, and then compiled for an FPGA target. Digital modules provide parameters to the digital control circuitry and read results from the FPGA according to a program running on a workstation.

The schematic of the analog portion placed on the bread-board is shown in Figure 6.2. The DUT is centered around a VCO from a 74HC4046 monolithic phase-locked loop. However, because this IC uses a voltage charge pump and since its phase detector could not be separated from this block, the phase detector is implemented on the FPGA. The charge pump is built out of discrete NPN and PNP transistors, a resistor and analog switches from a 74HC4066. The charge pump circuit is in fact the critical circuit for determining the maximum frequency of the test. It was operated at 100 kHz as parasitics of the board and the time constants of some components, notably the analog switches, did not allow for a higher frequency. However, the measurement scheme should be extendable to much higher frequencies as the prototype circuits are similar in nature to the PLL components.

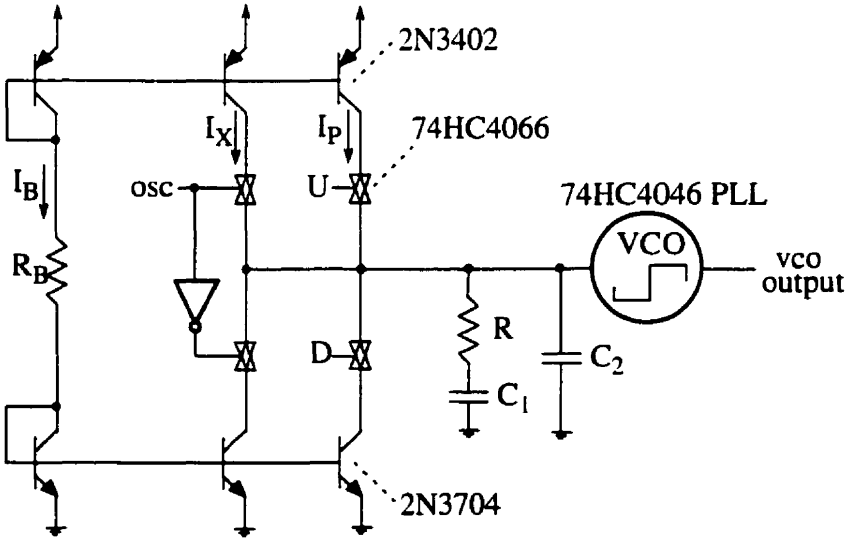

Figure 6.2: Analog portion of experimental set-up on bread board.

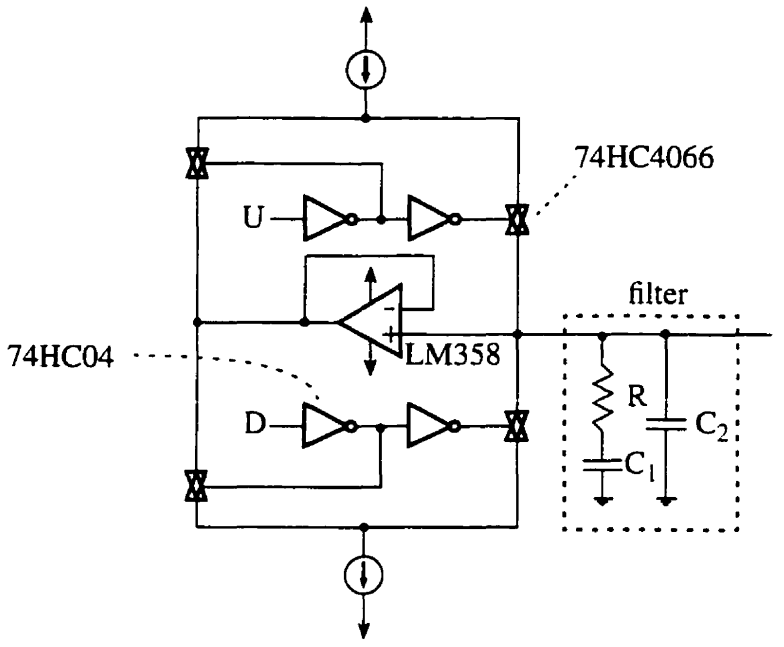

Figure 6.3: Circuit to maintain the current sources in their linear region.

A circuit is also required to maintain the transistors composing the current sources inside their linear region when the switches are open. It is shown in Figure 6.3 for one of the two charge pumps. On the right side are the components implementing the active charge pump and the filter, while on the left is a current sink based on an op-amp to avoid charge build up when switches are open and thus saturation of the transistors composing the current sources.,

Two different PLL architectures are used to qualify the BIST method. The first one has the VCO locking to the input reference frequency and therefore no counter. This architecture is used to verify the digital modulation of the input method and the test signal injection with a charge pump method. Since the other PLL stimulus technique, digital modulation of the output, requires a counter in the feedback path, a different architecture is necessary to verify its behavior.

6.2 Experiments Without a Counter in Feedback Path

A diagram of the digital portion of the test setup for the first architecture is shown in Figure 6.4. Two types of digital signals for jitter creation (*p(n)* and *osc*) are generated by a lowpass ΔΣ oscillator programmed on the FPGA. It uses 24-bit buses to achieve a tunability of 55 parts per million of its clock frequency. A 3-bit quantizer in addition to the standard 1-bit quantizer makes the signal generator capable of multi-bit output (refer to Figure A.10) for the purpose of digital phase modulation. It should be noted that apart from the quantizer, ΔΣ modulator circuitry is not duplicated as it will be operated in a time-shared mode at double speed for multi-bit operation.

The input to the PLL can be set to accommodate both jitter generation methods. For the loop jitter injection scheme, a 100 kHz square wave is presented to the input of the phase detector. Alternatively, this input can also be the same signal phase modulated with the help of a 800 kHz test clock. Eight jitter steps are therefore possible, resulting in a π/4 quantization.

Various jitter threshold circuits are also implemented on the FPGA. The π jitter threshold circuit of Figure 5.6 (a) will be employed in conjunction with the loop jitter injection. On the other hand, thresholds of π/2 and π/4 are implemented for the digital phase modulation method, taking advantage of the higher test clock frequency.

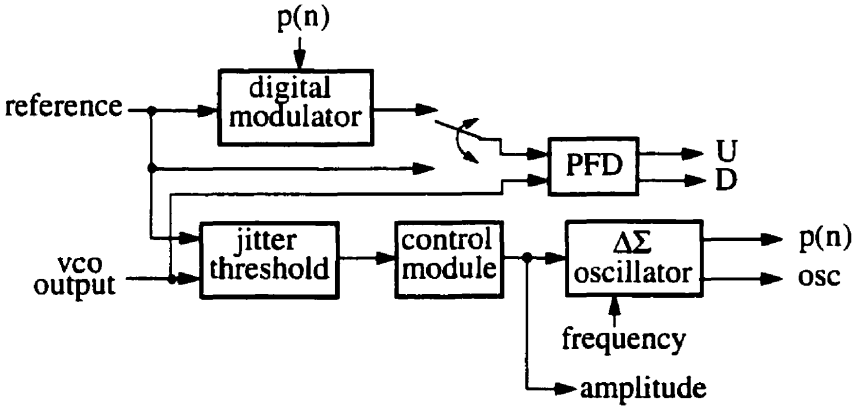

Figure 6.4: Digital portion of first architecture.

For each test, a warm-up stage of 2^{14} data cycles is executed to remove transients before a 2^{16} data cycles test stage is performed. The error threshold is set to 64, corresponding to a bit-error rate (BER) of 10^{-3}. A control module built around a finite state machine selects the amplitude of the input jitter for the ensuing test according to the output of the jitter threshold circuit, using the binary search algorithm. At each frequency point, the amplitude is resolved to an accuracy of 15 bits within 13 seconds. The entire digital circuitry for all the experiments requires 81 percent of the resources of an XC4010 FPGA. This experimental setup is connected to a workstation through I/O modules to allow a driving software to set the lowpass $\Delta\Sigma$ oscillator frequency as well as read the amplitude after a frequency test.

The jitter transfer function measurement is carried out for both the jitter injection and the digital phase modulation techniques on three PLL configurations having different bandwidths and damping values. Table 6.1 summarizes the main parameters of these experiments. The same current value is used for both charge pumps ($I_X = I_P$).

	Experiment 1	Experiment 2	Experiment 3
R	1.0 kΩ	2.2 kΩ	3.3 kΩ
C_1	2.0 µF	5.1 µF	10 µF
C_2	100 pF	480 pF	100 pF
I_X	47 µA	47 µA	47 µA
K_P	7.5 µA/rad	7.5 µA/rad	7.5 µA/rad
K_O	175 krad/V	175 krad/V	175 krad/V

Table 6.1: PLL without counter experiments parameters.

The results of the experiments on the first configuration are shown in Figure 6.5. A measured jitter transfer function for each jitter generation method is displayed. Note that the transfer functions are presented in the continuous-time domain as designers are more familiar with this representation. The dotted line represent the theoretical jitter transfer function as predicted from the values of the components displayed in Table 6.1. The phase modulation scheme used a $\pi/2$ threshold for this experiment. The curve shows a 0.4 dB offset which can be attributed mostly to the static jitter of the PLL. A 0.4 dB value for this offset

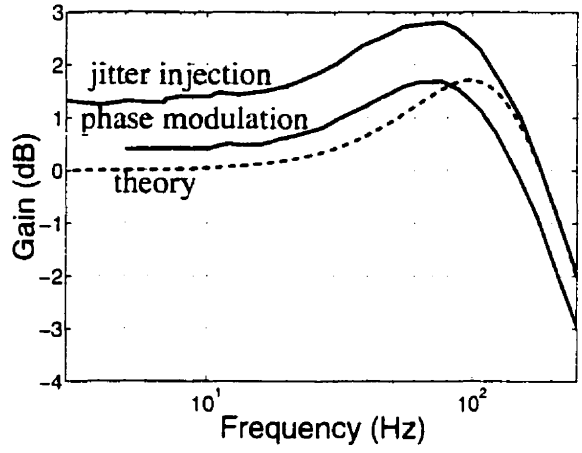

Figure 6.5: Jitter transfer function: no counter, experiment 1.

represents a 450 ps (4.5%) static jitter. For the other jitter creation scheme, the signal injection clock was chosen to be 50 kHz, that is half the PLL rate, in order to demonstrate the flexibility in selecting this parameter. The offset is larger, possibly because of mismatch between the two charge pumps realized out of discrete transistors. In fact, subtracting the static jitter value calculated previously, the charge pump mismatch is evaluated using the results of Figure 6.5 to be 8.8%.

Table 6.2 summarizes the features of the curves after removal of the offsets for all the experiments. The results of the first experiment can be found in the second and third column. Both methods yield similar results for the PLL bandwidth and jitter peaking. The theoretical predictions are slightly off, most probably because of the parasitics of the setup which were not accounted for in the calculations.

	R = 1.0 kΩ C_1 = 2 µF		R = 2.2 kΩ C_1 = 5.1 µF		R = 3.3 kΩ C_1 = 10 µF	
	3 dB (Hz)	j.p. (dB)	3 dB (Hz)	j.p. (dB)	3 dB (Hz)	j.p. (dB)
signal injection	238	1.49	435	0.42	---	---
phase modulation	234	1.29	395	0.24	541	0.04
theory	283	1.71	466	0.22	694	0.05

Table 6.2: PLL without counter results summary.

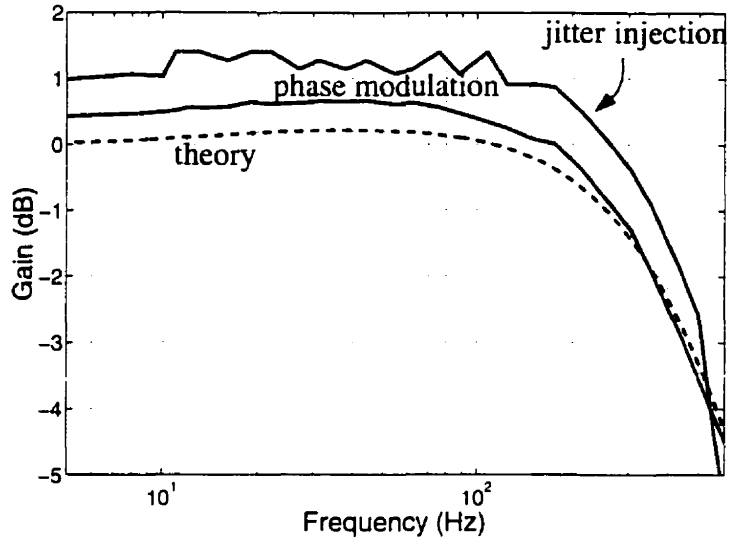

Figure 6.6: Jitter transfer function: no counter, experiment 2.

The jitter transfer functions measured in the second experiment are shown in Figure 6.6. This PLL exhibits a larger bandwidth and is more damped. It can be seen that the curve obtained here with the jitter injection technique is of lesser quality. This came about because the larger bandwidth yields a lower output jitter SNR. From the graph of Figure 5.8, it can be seen that this SNR is barely over 20 dB. On the other hand, the digital phase modulation still shows a smooth curve because of the 3-bit quantization which results in lower noise levels. The meaningful parameters are summarized in the two middle columns of Table 6.2.

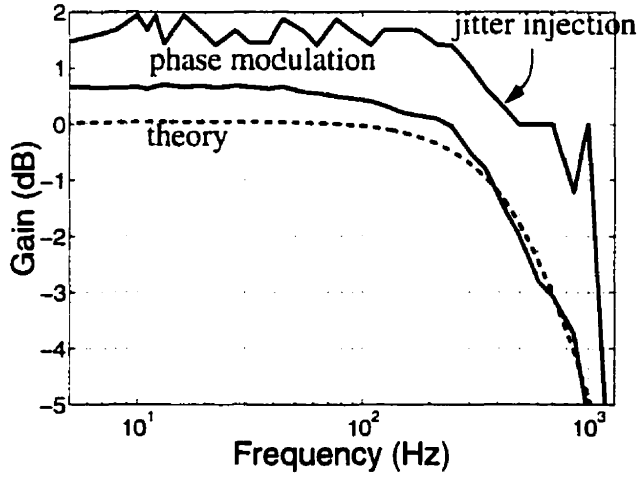

Figure 6.7: Jitter transfer function: no counter, experiment 3.

The filter impedance was further increased for the third experiment. The measured jitter transfer functions are shown in Figure 6.7. Obviously, the curve obtained with the jitter injection technique is of poor quality. This problem is caused because a current is being injected in a large impedance. As a result, the circuit node where the test signal is injected cannot handle the large voltage swings. In other words, the charge pump and the VCO probably venture out of their linear range because of the large charge injected on each cycle. This issue was discussed in Section 4.4. On the other hand, the digital phase modulation scheme is less disturbing to the circuit because of the 3-bit quantization. Indeed, multi-bit quantization results in lower noise levels. Even though the high frequency input jitter noise does not show up in the output signal, it must still be absorbed by the PLL. This noise reduces the dynamic range of sensitive circuit nodes. To further attenuate this problem for phase modulation, a lower threshold value of $\pi/4$ was selected as it allows for smaller input jitter amplitudes. The resulting curve is much smoother. Again, the meaningful parameters are displayed in the two right most columns of Table 6.2. Entries for the jitter loop injection method are missing because the quality of the curve does not allow for the extraction of its appropriate parameters.

6.3 Experiments With a Counter in Feedback Path

A diagram showing the digital portion of the test setup for the second architecture is drawn in Figure 6.8. A 800 kHz test clock is required for the jitter threshold circuit and it must

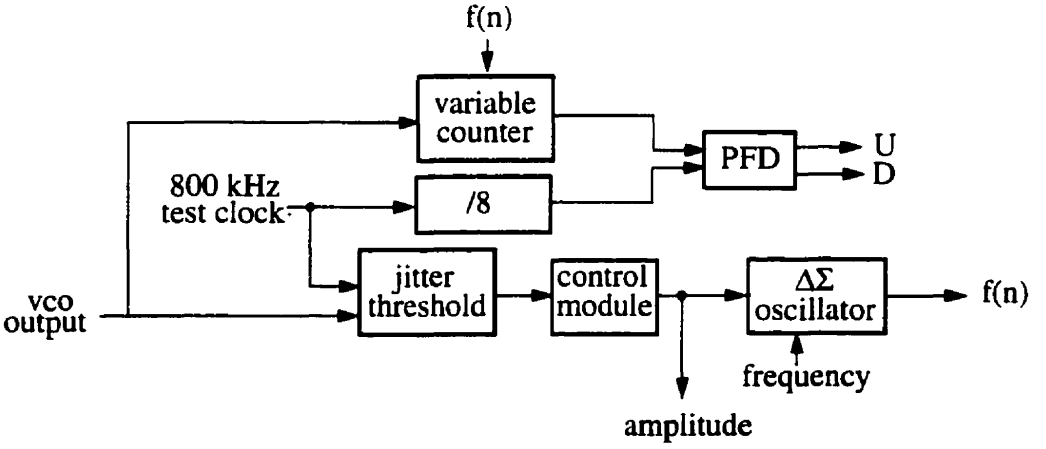

Figure 6.8: Digital portion of second architecture.

therefore be divided by 8 to generate a 100 kHz reference signal. The signal controlling the variable length counter is generated by a lowpass ΔΣ oscillator programmed on the FPGA. In contrast to the circuit in the previous section, it uses an 8-to-1 multiplexer to obtain a 3-bit signal instead of time-multiplexing the ΔΣ modulator. The average counter length is 8 and therefore the VCO carrier frequency will be 800 kHz.The jitter transfer function was measured for two PLL configurations having different bandwidths and damping values. The value of the components used in the experiments are summarized in Table 6.3.

	Experiment 1	Experiment 2
R	470 Ω	1.2 kΩ
C_1	2.0 μF	10 μF
C_2	100 pF	100 pF
I_X	47 μA	47 μA
K_P	7.5 μA/rad	7.5 μA/rad
K_O	175 krad/V	175 krad/V

Table 6.3: PLL with counter experiments parameters.

The results of the first experiment are shown in Figure 6.9. The measured jitter transfer function is plotted with a solid line. For comparison, a transfer function obtained using the measured values of the discrete components and PLL theory is plotted with a dashed line.

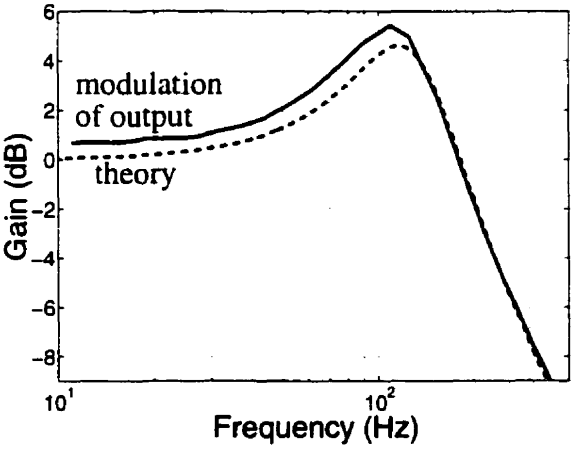

Figure 6.9: Jitter transfer function: counter in feedback, experiment 1.

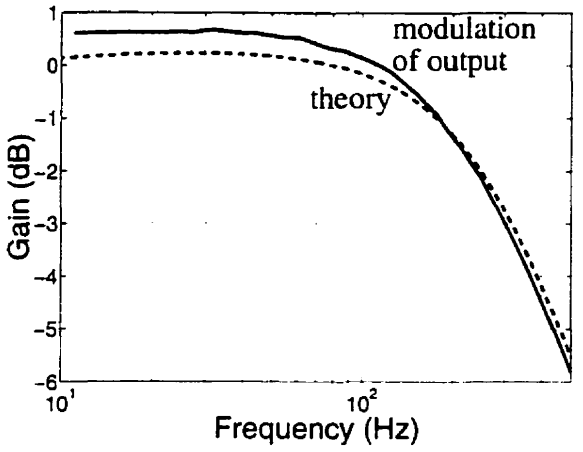

Figure 6.10: Jitter transfer function: counter in feedback, experiment 2.

It can be seen that the experimental results are quite similar to what is predicted by theory. However, a 0.6 dB offset is noticeable. It can be attributed to the static jitter of the PLL arising from mismatch in the *up* and *down* charge pumps. The measurement of a jitter transfer function with a larger damping value (increased R and C) is shown in Figure 6.10.

The features of the measured jitter transfer function are compared with the theoretical predictions in Table 6.4. Much like the results of the previous section, the match is quite good considering the fact that parasitic effects were not factored in the theoretical predictions.

	$R = 470\ \Omega$ $C_l = 2.0\ \mu F$		$R = 1.2\ k\Omega$ $C_l = 10\ \mu F$	
	3 dB (Hz)	j.p. (dB)	3 dB (Hz)	j.p. (dB)
experimental results	208	4.73	264	0.07
theory	219	4.61	317	0.23

Table 6.4: PLL with counter results summary.

6.4 Summary

The functionality of the jitter transfer function measurement scheme was verified with experiments. The test prototype was built out of discrete analog components and an FPGA. Measurements were first made with both jitter injection using a charge pump and digital phase modulation of the input for three different PLL configurations. The experiments were then repeated with a PLL including a counter and the digital modulation of the output method for two configurations. Comparison of the test results with theoretical predictions obtained from PLL components values show that the method is sound. However care must be exercised in selecting the jitter stimulus method and the test parameters, the test clock frequency for example, to avoid results of low quality.

Chapter 7 : Conclusion

7.1 Short Summary

A PLL jitter transfer function measurement technique has been presented. It is entirely digital except for the possible addition of a charge pump. The technique is suitable for built-in self-test since it does not require trimming. Three methods were introduced for the creation of jitter, allowing trade-offs between test clock frequency on one side and loading, complexity and accuracy on the other side. Output jitter is analyzed using a jitter threshold. Experimental results were presented which suggest this scheme could be successfully implemented on silicon.

7.2 Overhead for BIST Implementation

For any integrated measurement scheme, the area overhead is obviously a major concern. While the digital portion of the experimental setup of Chapter 6 uses a large portion of the FPGA, a much more compact implementation is possible. Indeed, a $\Delta\Sigma$ oscillator was selected as the signal generator because of its versatility since a complete jitter transfer function was sought. However, in many applications, a smaller number of signal frequencies and amplitudes are necessary and a fixed-length periodic byte stream could generate the signals for the cost a few kilo-bits of RAM. For example, to verify that the bandwidth of the PLL is smaller than some value, only one test is required. Exact figures for overhead, or measurement time depend heavily on the PLL application. Moreover, one should not have a dogmatic stance about overhead as the addition of the on-chip measurement circuits adds value to a system. The economic gains from a BIST may far overweight the cost of extra silicon.

7.3 Significance of Original Results

The jitter transfer function measurement method in this work allows the verification of PLL compliance to specifications using a low-speed digital interface instead of high-speed analog instruments. This is especially important for a number of digital integrated circuits which are devoid of analog blocks except for a PLL. While these ICs would normally require a high-speed mixed-signal tester, our BIST proposal makes them testable with a much less expensive low-speed digital tester. Furthermore, our proposed method is the only work published that addresses the measurement of PLL characteristics as a whole.

Three different methods have been presented for stimulus generation allowing trade-offs between speed, accuracy and complexity. A wide variety of PLL configurations may thus be accommodated. Two of the signal generation methods are novel while the third one has never been employed in test applications. In particular, digital phase modulation of the input is the first reported use of delta-sigma modulation in the phase (time) domain for digital-to-analog conversion. In fact, all these signal generation methods rely on delta-sigma modulation. The effect of the quantization noise inherent to this technique on the dynamic range of internal PLL nodes was studied and bounds useful for test design were obtained. The output signal analysis method is very simple. It thus has a small impact on the test overhead. The accuracy of this jitter measurement method was examined. All the error sources were identified and their effects on the results were described.

7.4 Future Work

Three directions are possible for the future of this project. The first one is to consolidate the gains by implementing the jitter transfer function measurement scheme alongside a PLL from a commercial part. Integrating the stimulus generation and jitter analysis circuits with a sub-micron process and performing the test at a few hundred megahertz would dissipate any doubt about the possibility of BIST for high speed PLLs. However, this requires a significant effort as well as collaboration with an industrial partner to obtain a state-of-the-art device. The second research perspective is the possibility of measuring the jitter output us-

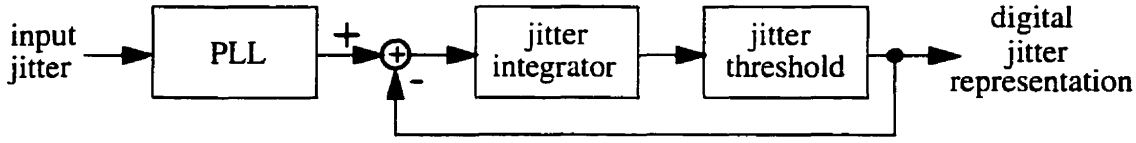

Figure 7.1: Applying delta-sigma modulation concept to jitter measurement.

ing some application of delta-sigma modulation. This concept is illustrated in Figure 7.1. The jitter threshold circuit only performs a coarse quantization. However, the decision is fed back and subtracted from the input. The effect is that the error occurred because of quantization is taken into account for the following decisions. It is expected that, with a suitable jitter filter, an integrator for example, quantization noise can be shaped outside the signal band. To recover the desired signal, decimation is performed on the digital output. The result is a more accurate jitter measure and the possibility to perform multi-tone tests. However important issues must be solved to actually implement the concept of Figure 7.1. Foremost is the analog nature of the jitter summer and the jitter integrator which may require calibration. Also note that such an apparatus would not be able to measure PLL jitter noise as this signal is highpass with a large bandwidth.

The final direction is the exploration of the use of the tools disclosed in this work to measure other specifications of the PLL such as lock range, lock time or jitter noise.

Appendix A: Delta-Sigma Modulation

Delta-sigma modulation[1] is a technique that allows the encoding of a bandlimited signal, either digital or analog, onto a very small number of bits. An error is created by this quantization operation but it is filtered such that its power appears mostly at high frequencies, away from the frequency band occupied by the signal. This feat is realized by feeding back the quantization error to a filter and summing it with the input before the next quantization is performed. With the help of delta-sigma modulation, arbitrary signals, such as a sinusoid, can be represented by a square wave (1-bit signal) with the difference being composed mostly of high-frequency noise. This noise may later be filtered to recover the original signal plus minimal in-band noise. The one-bit format is especially useful in analog signal generation as one-bit digital-to-analog converters (DAC) are easy to design and are inherently linear.

A.1 Delta-Sigma Modulation Basics

$\Delta\Sigma$ modulators, or oversampling converters, are extensively used in data conversion circuits. They have replaced traditional schemes, labelled Nyquist converters[2], in most applications where the clock frequency may be made significantly larger than the signal bandwidth. Even though $\Delta\Sigma$ modulators used in DACs and in analog-to-digital converters

1. Also referred to as sigma-delta ($\Sigma\Delta$) modulation (see [47]).
2. Nyquist converters use a clock frequency twice the maximum frequency component of the input signal in accordance with Nyquist theorem.

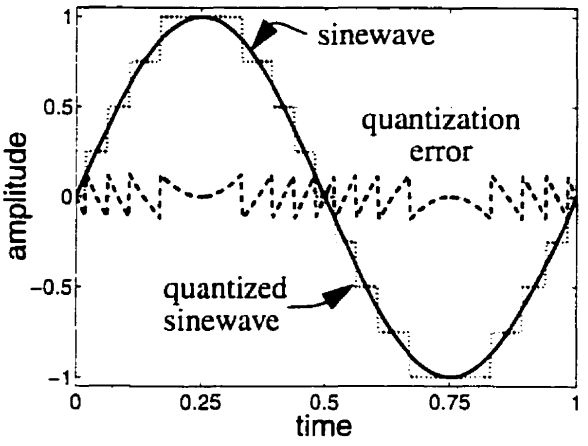

Figure A.1: Quantized sinusoid and quantization noise.

(ADC) share many features, we will be mainly concerned with the former as our goal is analog signal generation.

The purpose of a $\Delta\Sigma$ modulator in the context of digital-to-analog conversion is to encode the input such that there is a reduced number of quantization levels at the output. The output signal may ultimately be encoded into a single bit stream. However, reducing the number of quantization levels involves adding an error signal corresponding to the difference between the quantized signal and the original signal. This error signal is labelled quantization error or quantization noise. Figure A.1 illustrates the quantization error for a sinusoid of amplitude equal to one with nine quantization levels. It can be seen that this quantization error is quite large.

In fact, the power of the quantization error power for a 1-bit encoding is so large that it would not be practical with Nyquist conversion. Instead oversampling and noise shaping techniques must be used. The difference between Nyquist conversion and oversampling can be highlighted by observing their characteristics in the frequency domain. Consider a single tone input to each converter. The power spectral density appearing at the output are those seen in Figure A.2 (a) and (b). In Nyquist conversion, as the sampling frequency (f_S) is equal to twice the bandwidth of the signal (f_b), the quantization noise power, illustrated by the shaded area, is distributed between DC and f_b. Consequently, the noise floor, represented by a broken line, is high. Note that even though the quantization noise is shown here

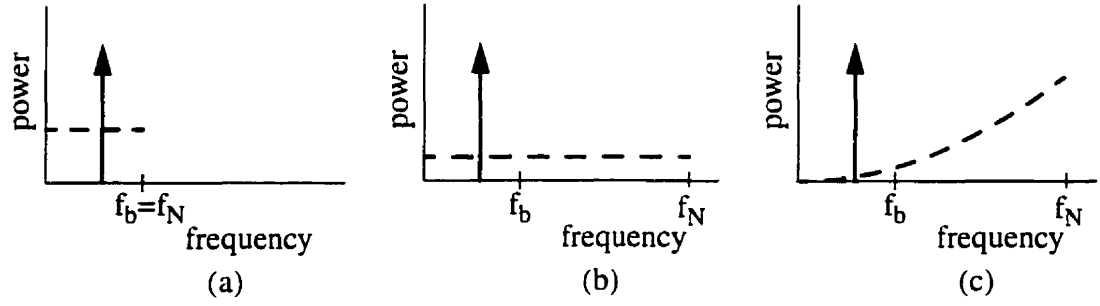

Figure A.2: (a) Nyquist conversion. (b) Oversampling conversion (c) Noise shaping.

as white, it is usually not the case for most signals. This simplification is used to help visualize the difference between the types of conversion. When sampling at a larger frequency (oversampling), the quantization noise power is the same, but is spread over a larger frequency band. Therefore the in-band noise, defined as the noise falling between 0 and f_b, is reduced and the noise floor is lower. The ratio of the Nyquist frequency (f_N) to the bandwidth of the signal frequency band is termed the oversampling ratio (*OSR*):

$$OSR = \frac{f_N}{f_b} = \frac{f_S/2}{f_b}. \tag{A.1}$$

$\Delta\Sigma$ modulation further filters out most of the quantization noise in a signal band by using noise shaping techniques resulting in the power spectrum shown in Figure A.2(c). It is evident that the in-band noise is significantly reduced.

To achieve noise shaping, a feedback loop is built around the quantizer. While the transfer function of the signal to the converter output is unaffected, at least in the signal band, the transfer function from the quantizer to the output represents a filter with significant attenuation in this signal band. This topic will be further developed in the following section. A generic $\Delta\Sigma$ modulator topology is shown in Figure A.3. The blocks *M(z)* and *N(z)* are discrete-time blocks that implement the required noise shaping behavior. $\Delta\Sigma$ modulators may be divided in different categories according to the location of the signal band. $\Delta\Sigma$ modulators shaping the quantization noise to high frequencies, as illustrated in Figure A.2 (c), are labelled lowpass (LP) and were introduced first [55, 56]. They are widely used in low frequency data conversion applications such as telephony and digital audio. Other flavors of

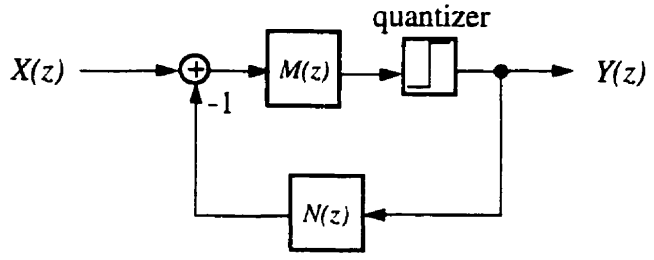

Figure A.3: Generic delta-sigma modulator.

ΔΣ modulation were later introduced. Bandpass ΔΣ modulators for example shape the quantization noise to both the low and high ends of the Nyquist interval [57]. The signal band can therefore be located at any ratio, up to one-half, of the clock frequency.

A.2 Mathematical Model of Delta-Sigma Modulators

ΔΣ modulators are highly non-linear circuits because of the presence of a quantizer. However, the theory on linear systems is much more developed than that for non-linear systems. Linear analysis is thus applied to ΔΣ modulators to provide insight into their behavior. For linear analysis, the quantizer shown in Figure A.4(a) is thus replaced by a two-input summer circuit with one input coming from the original source and the other attached to a signal source Q(z) representing the quantization noise, as shown in Figure A.4(b).

However, the characteristics of the quantization error strongly depends on the quantizer input signal. Nevertheless, it has been observed that in most instances, the quantization error, for a signal below the saturation level of the quantizer, approximates additive white noise. It should be emphasized that this assumption is not always valid since the quantization error is sometimes strongly correlated with the input signal. However, it is standard practice to use this assumption in analyzing circuits involving quantizers, most notably delta-sigma

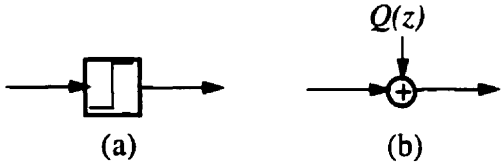

Figure A.4: (a) One-bit quantizer (b) Linear model of quantizer.

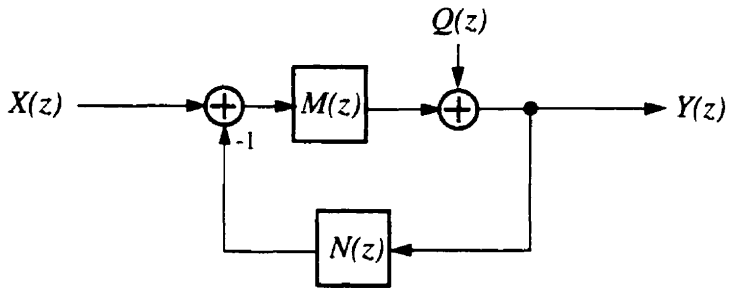

Figure A.5: Linear model of generic delta-sigma modulator.

modulators. Generally, results obtained from simulations agree to a large extend with predictions done using this linear model.

Modeling the quantizer of the network of Figure A.3 as a source of additive white noise, the linear model of the generic ΔΣ modulator, shown in Figure A.5, is obtained. The transfer function from the quantization noise source, $Q(z)$, to the output, $Y(z)$, is called the noise transfer function, while the response of the output to the input, $X(z)$, is designated the signal transfer function. ΔΣ modulators are therefore characterized by these two functions, according to:

$$Y(z) = STF(z) \cdot X(z) + NTF(z) \cdot Q(z). \tag{A.1}$$

For the generic ΔΣ modulator of Figure A.3, the signal transfer function (*STF*) is given by

$$STF(z) = \frac{M(z)}{1 + M(z)N(z)}, \tag{A.1}$$

while the noise transfer function (*NTF*) is given

$$NTF(z) = \frac{1}{1 + M(z)N(z)}. \tag{A.2}$$

The order of the NTF function indicates the order of the ΔΣ modulator. NTFs of higher order usually translate to superior conversion performances. However not any transfer function may be used for noise shaping. Indeed, ΔΣ modulators may be unstable, even though their linear model predicts otherwise. Instability is observed when the signal amplitude at an internal node exceeds an arbitrary large bound. Thus far, there has been no absolute method for determining stability of a ΔΣ modulator other than extensive simulation.

A.3 Second-Order Lowpass Delta-Sigma Modulator

A very common architecture for a lowpass ΔΣ oscillator is the double-integration design [58] illustrated in Figure A.6. Using the additive white noise model for the quantizer and performing linear analysis, it can be shown that the signal transfer function is z^{-1} while the noise transfer function is

$$NTF(z) = (1 - z^{-1})^2. \tag{7.1}$$

It is a second-order lowpass filter with two zeros located at z=1. This design has been found to be very robust to instability.

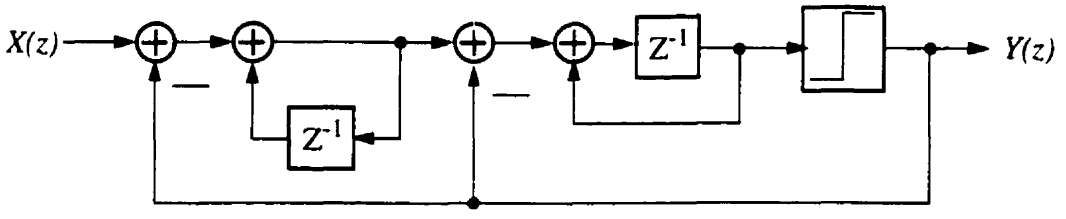

Figure A.6: Second-order lowpass delta-sigma modulator.

A.4 Slope-Limited Delta-Sigma Modulator

In some applications, it is desirable to encode the output of a delta-sigma modulator differentially. An example is the circuit for the generation of high amplitude jitter signals of section 4.1.4. For correct encoding on a finite number of bits, the ΔΣ modulator must enforce restrictions on its output signal, *p(n)*. These constraints are formulated as two inequalities expressed as

$$-D + 1 \le p(n) - p(n-1) \le -D + 2^b. \tag{A.1}$$

In the context of the above mentioned application, the left hand side is the requirement that the pulses be distinct and appear in the correct order (causality). D is in fact the ratio of the test clock frequency over the PLL operating frequency. The right hand side inequality is introduced because of the finite number of bits, b, allocated for the differential encoding and hence the limited range of $p(n)-p(n-1)$. These constraints can be incorporated in the

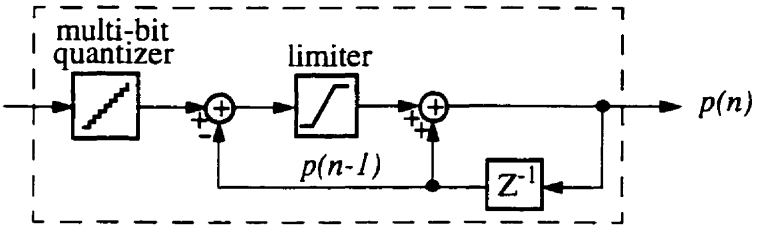

Figure A.7: Modified quantizer for slope-limited delta-sigma modulator.

feedback loop of the ΔΣ modulator by modifying the quantizer as shown in Figure A.7. The limiter in fact restricts the slope of the output signal. The resulting device is thus labelled a slope-limited ΔΣ modulator. The beneficial effect of incorporating this modified quantizer in the feedback loop of a ΔΣ modulator is the noise-shaping of the error signal introduced by the limiter.

As the circuit of Figure 4.11 does not restrict the amplitude of the phase jitter, the limiting factor is the stability of this type of ΔΣ modulator. Figure A.8 shows the maximum input which can be applied to a slope-limited delta-sigma modulator for two different byte widths and thus two different ranges for the limiter. The frequency axis is relative to a 256 time-step period. Higher input amplitude than the graph indicates would result in distortion or instability. As expected, the curve decays with increasing frequency to counter-balance the larger slope. This behavior is similar to what is observed for slew-rate limited op-amps. An

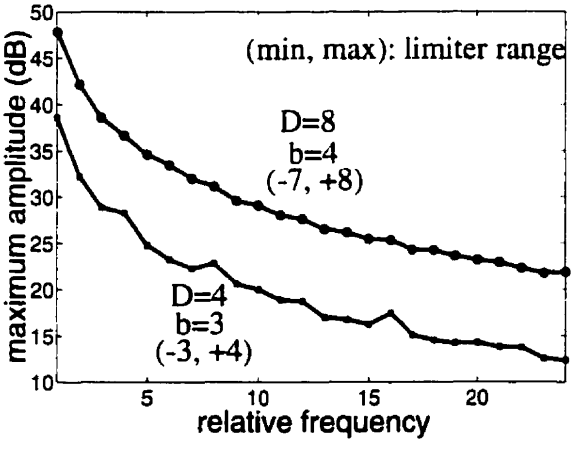

Figure A.8: Maximum amplitude of signal at input of slope-limited delta-sigma modulator.

interesting behavior, though not yet investigated, is the bumps in the curve at factors of the period (4, 8 and 16). Note that the possibility of generating signal for a frequency ratio of only 2 was explored. However, it was found that stable signals may not be obtained for this parameter value.

A.5 Noise-Shaped Sinusoid Generation

For the measurement of transfer functions, sinusoids are used to stimulate the device-under-test. In this section, three methods for generating sinewaves encoded with ΔΣ modulation are presented. The selection of one of these for an implementation is based on the available resources in the system, silicon area, PLL speed and required accuracy.

A.5.1 Direct Digital Synthesis

The most straightforward and versatile signal generation scheme is a ROM-based digital frequency synthesizer [59] followed by a lowpass ΔΣ modulator as illustrated in Figure A.9. To the left, an integrator converts a frequency argument to a phase signal. This signal is used to address a ROM where the sine samples are stored. The output of the ROM is a large word which is compacted on a smaller number of bits by the ΔΣ modulator. The final result may be a one-bit signal or a multi-bit signal [60] as required by the application. However, the direct frequency synthesis solution requires a large silicon area and is thus usually not a good choice for BIST. Yet, if a large block of RAM is present in the system, then it could be borrowed to store data for signal generation during the analog test phase.

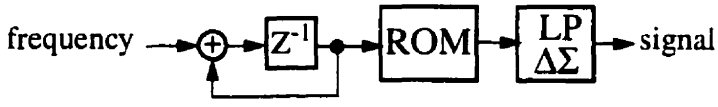

Figure A.9: ROM-based digital frequency synthesis.

A.5.2 Delta-Sigma Oscillator

Sinewaves may also be generated using digital resonators such as the one illustrated in Figure A.10 (a). It oscillates with a frequency set by the coefficient *k* and the amplitude se-

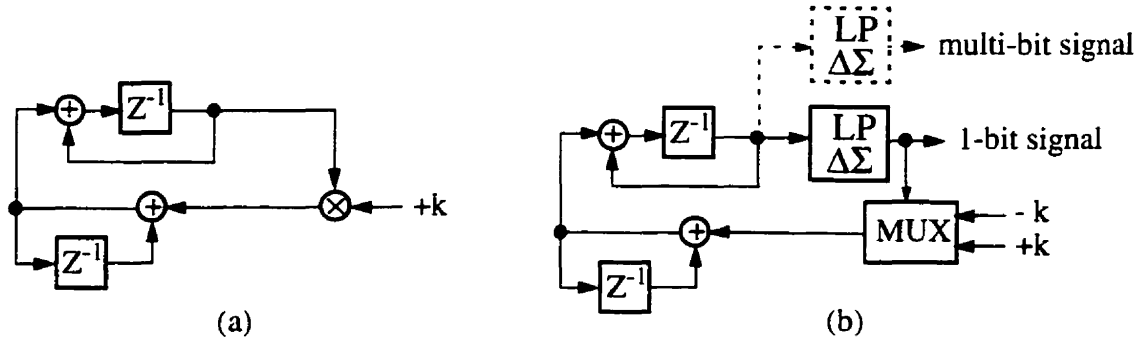

Figure A.10: (a) Digital resonator. (b) Delta-sigma oscillator.

lected by the initial conditions of the registers. Yet, the multiplication operation is expensive both in silicon area and in latency. For area critical applications, the lowpass ΔΣ oscillator was introduced [23]. In this circuit, depicted in Figure A.10 (b), the input to the multiplier is encoded on a single bit with the help of a ΔΣ modulator. The multiplication is then realized by a simple 2-to-1 multiplexer. For multi-bit signal generation, the size of the multiplexer could be extended. For example, a three-bit output would require an 8-to-1 multiplexer. Another option is to place a multi-bit ΔΣ modulator in parallel with the 1-bit ΔΣ modulator as illustrated with dashed lines in Figure A.10 (b). However, as both ΔΣ modulators are identical except for the quantizer, hardware can be time-shared. This principle has been used before in the generation of multi-tone signals [24]. Because it does not require a ROM or a multiplier, the silicon implementation of ΔΣ oscillators is very efficient. However the presence of a non-linear block in the feedback loop makes this device difficult to predict using a linear model. Off-line simulations are required to achieve maximum accuracy. Nevertheless, ΔΣ oscillators may be coded using a hardware description language. Thus they can be implemented quickly with a silicon compiler, for example on a field programmable gate array (FPGA). Also, available on-chip computing resources such as a DSP may also emulate this circuit as they are inactive during analog test time.

A.5.3 Fixed-Length Bit Streams

The last method consists in generating a data stream from a software sinewave generator and lowpass ΔΣ modulator and then selecting a subset of this stream [61]. This subset is stored in memory and the resulting data stream is then repeated as illustrated in

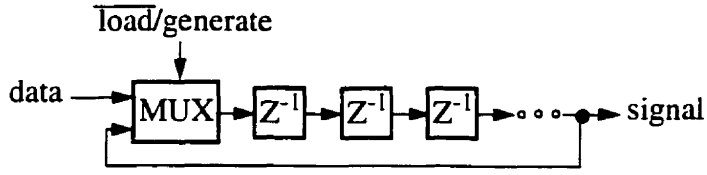

Figure A.11: Fixed-length periodic byte stream.

Figure A.11. One can view this approach as a special case of the ROM-based digital frequency synthesis scheme presented above [62]. This method is particularly useful for single-bit signals as they can be represented using a very small number of bits, on the order of a hundred. The downside is that a different data stream is required for each frequency and amplitude desired. However considerable speed can be achieved with this technique. Furthermore, this method is not restricted to sinusoids but is suitable for any periodic signal.

Yet, the output of a $\Delta\Sigma$ modulator is not periodic. Signal quality will thus vary widely with different subsets of the same length of a given $\Delta\Sigma$ modulator output. Some form of optimization is necessary to obtain good results. However, no method has been introduced so far beyond random sampling of the subsets. For the purpose of this work, two simple criteria were developed to reduce the search space. The first one is to restrict the search to subsets which have a zero DC component. This only works of course if the $\Delta\Sigma$ modulator has an NTF zero at this frequency. The second criterion draws on knowledge from windowing theory. Indeed, it has been shown that the best attenuation is obtained when the largest number of derivatives are continuous at the ends of the window. The proposed strategy is thus to compare the ends of the subset with the data at the boundary and look for the maximum

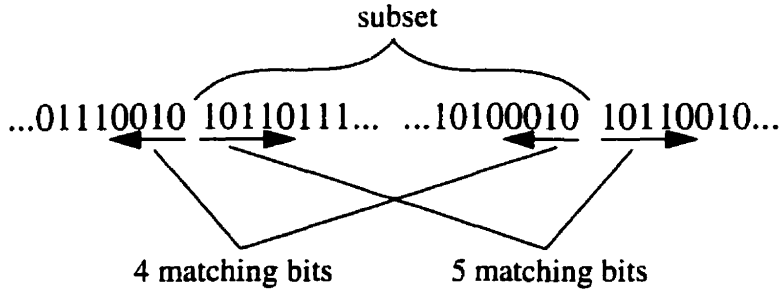

Figure A.12: Matching of byte stream ends.

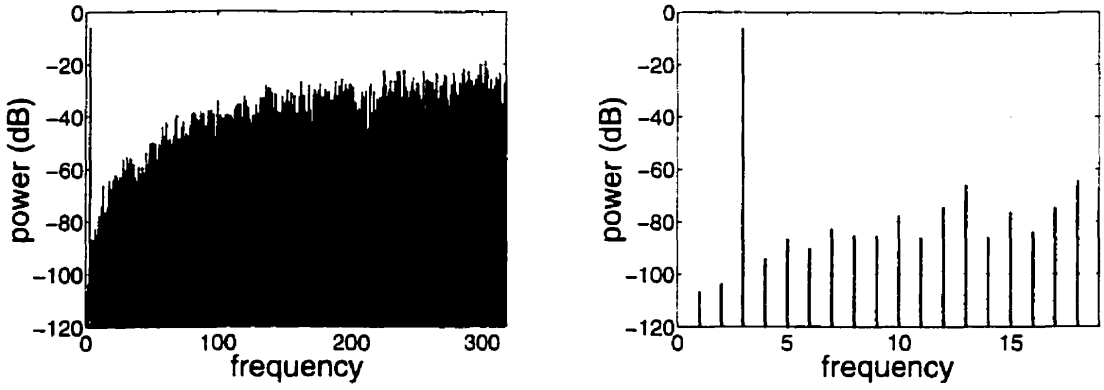

(a) (b)

Figure A.13: Spectral power density of finite length 3-bit stream:
(a) Nyquist interval. (b) Signal band.

number of symmetric symbols. An example is illustrated in Figure A.12. Here 5 bits and 4 bits match on the beginning and the end of the subset, respectively.

Figure A.13 illustrates the quality that can be obtained with finite byte streams. The signal displayed is a sine wave encoded on 640 three-bit words. Three cycles were encompassed in the byte stream. The spectrum is discrete because the signal is periodic.

Appendix B: Bandpass Signal Generation

As explained in Chapter 2, the implementation of built-in self-test (BIST) requires analog signal sources that do not demand external calibration. However, analog circuits are generally sensitive to process variations which occur during IC manufacturing. Their characteristics can not be predicted with accuracy before fabrication. Therefore, barring a catastrophic failure, the range of possible function must be described by a probability distribution. On the other hand, digital circuits are robust to local variations in integrated circuits. They may malfunction because of a catastrophic defect but since the behavior is then very different from the expected operation, this condition may be uncovered easily. Consequently, testing of digital circuits has matured to a point where solutions are automated with computer-aided software. A good strategy for on-chip measurement is therefore to rely as much as possible on digital circuits. Evidently, it is not possible to build an analog signal source or a measuring instrument without resorting to some type of analog circuit. Nevertheless, the dependency may be reduced such that all the failure modes of the outstanding analog circuits may be accounted for.

B.1 Bandpass Delta-Sigma Oscillator

Recently, a class of analog signal sources which reduce the analog portion to only a 1-bit digital-to-analog converter (DAC) have been introduced [63]. The lowpass delta-sigma ($\Delta\Sigma$) oscillator illustrated in Figure A.10 (b), in Appendix A, was the first member to be presented. It is suitable for testing devices which have a lowpass characteristic. For bandpass type devices such as a wireless receiver, a derivative of the previous circuit, labelled bandpass (BP) $\Delta\Sigma$ oscillator, was introduced [64]. This circuit is illustrated in Figure B.1

(a). It is composed of a digital resonator on the left side and a combination of a BP ΔΣ modulator and a multiplexer on the right side. These last two blocks actually implement a multiplier as was the case in the lowpass ΔΣ oscillator. The frequency of oscillation is a function of the constants K_c and k_f while the amplitude is set by these two parameters and the initial values of the registers. Figure B.1 (b) shows the spectral power density of the 1-bit output signal for a 1 MHz clock frequency as measured on a spectrum analyzer with a 73 Hz resolution bandwidth. The frequency range displayed is from DC to the Nyquist frequency (half the clock frequency). The signal is composed of a very high quality sinewave with a frequency near a quarter of the clock frequency and quantization noise at both low frequencies and high frequencies.

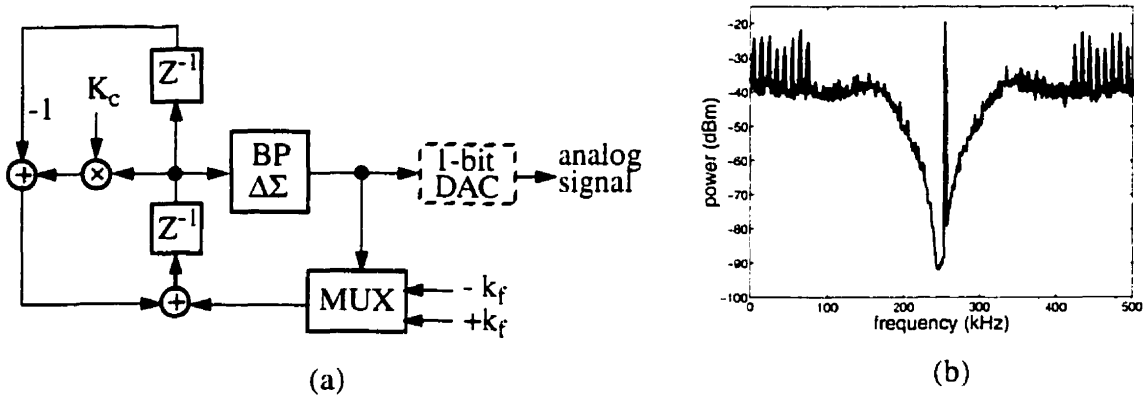

Figure B.1: Bandpass delta-sigma oscillator: (a) Circuit. (b) Spectrum of the output.

In mixed-signal test applications, the device-under-test (DUT) is expected to filter this noise. However, in practice, the attenuation may not be sufficient. Also, the analog front-end of the device may be unable to handle the 1-bit signal because of the large voltage swings. A possible solution to reduce this noise and therefore smooth the signal could be to place a bandpass filter between the bandpass ΔΣ oscillator and the DUT. Yet, this piece of analog circuitry must be tested for correct operation before it is used in the test and this task is not trivial. Increasing the number of quantization levels is an alternative as it reduces the noise power. However, a multi-bit DAC would be required. The behavior of this analog circuit can vary because of process variations and the signal quality may thus be lower than

expected. A technique for realizing multi-bit DAC, called dynamic element matching, uses $\Delta\Sigma$ modulation techniques for mitigating the effect of process variations [65]. It has become very popular recently in data converter applications. Nevertheless, the number of bits is limited by the exponential complexity of the method. The attenuation of the quantization noise might thus not be sufficient in many applications.

B.2 Bandpass Signal Generation Using a PLL

The problem of reducing the out-of-band quantization noise may be solved by employing a charge-pump PLL as a filter. While this circuit is analog, Chapters 4 and 5 have demonstrated how it can be verified using digital methods. Following its verification, a charge-pump PLL may then reliably attenuate the out-of-band noise in a digitally phase modulated signal which is obtained from methods described in chapter 4. The bandpass signal generation method is illustrated in Figure B.2. A baseband phase signal is first compressed with a $\Delta\Sigma$ modulator. The result is then encoded differentially. These operations may be performed in real-time or off-chip. This signal is used to digitally phase modulate a carrier as explained in Sections 4.1.3 and 4.1.4. The resulting bandpass signal is surrounded by phase noise which is mostly removed by a previously verified PLL. Note that generally the characteristics of the PLL need not be very accurate. A large tolerance can be provided for the parameters of the PLL.

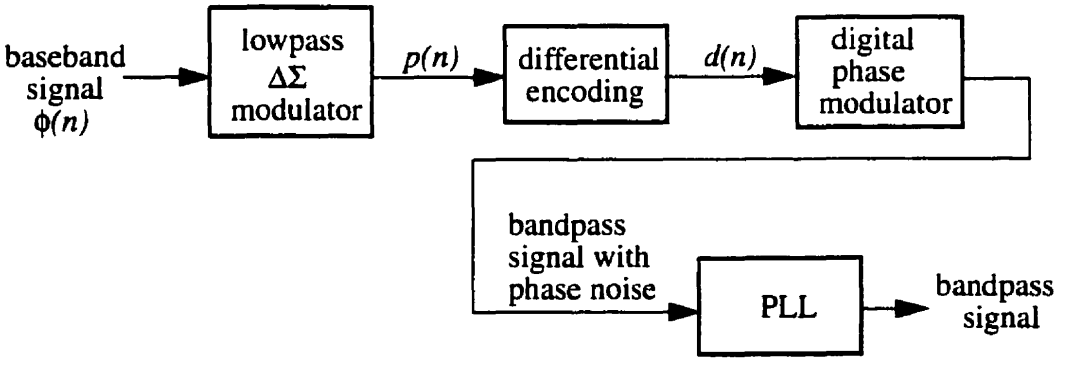

Figure B.2: Bandpass signal generation with a phase-locked loop.

B.3 Experimental Results

The phase filtering principle was verified by performing experiments with a field programmable gate array (FPGA) and a discrete PLL. The test clock frequency was chosen to be 4 MHz. As a value of 4 is selected for the oversampling ratio D (refer to Section 4.1.3), the carrier frequency appears at 1 MHz. The periodic phase signals, $\phi(n)$, are first generated using a floating point software. With the help of a slope-limited $\Delta\Sigma$ modulator (Section A.4), the phase index signal, $p(n)$, is calculated and then encoded differentially to obtain $d(n)$. This non-periodic signal is then turned into a finite length byte stream. A subset is carefully selected to maximize the signal-to-noise ratio (SNR). A technique for this operation is described in length in Section A.5.3. With a digital module, this byte stream is sent periodically to a FIFO on an FPGA. On the other side of the FIFO, a digital phase modulation circuit similar to what is seen at the bottom of Figure 4.10 consumes values of $d(n)$ to modulate a digital carrier. The result is fed to a 74HC4046 PLL for filtering. Results are captured at the output of the PLL using a spectrum analyzer having a 73 kHz resolution bandwidth.

B.3.1 FM Signal Generation

First a frequency modulated signal is generated. This is the type of stimulus used in previous sections for the measurement of the jitter transfer function of PLL. The fixed data stream length, L, is 620 words with three bits representing each word. Because of the peri-

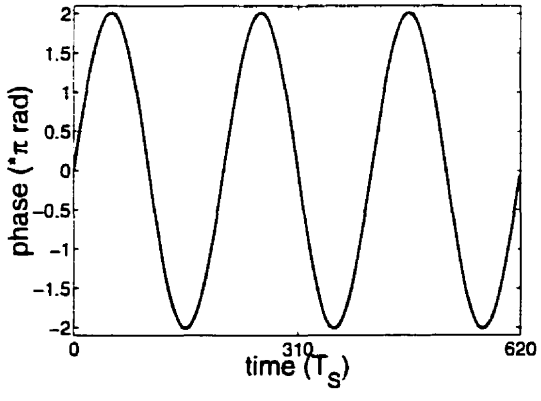

Figure B.3: Phase of one cycle of frequency modulated sequence.

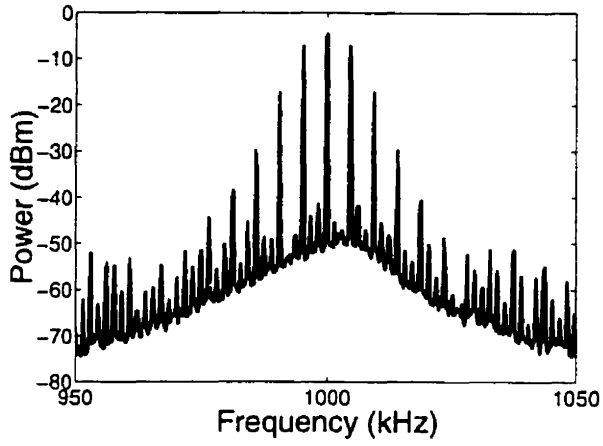

Figure B.4: Measured spectrum of FM signal.

odicity of the data stream, power pulses are expected to appear at intervals equal to the ratio of the carrier frequency, f_C, over the stream length, L, on the power spectrum. This data stream encompassed 3 periods of the sinusoid resulting in a continuous-time signal having a frequency of 4.7 kHz.

Figure B.4 shows the measured power spectrum of the digitally modulated square wave. The discrete power impulses are located 1.6 kHz apart. Sidebands appear at 4.7 kHz intervals as predicted. The phase noise becomes significant 30 kHz away from the carrier frequency.

B.3.2 Complex Bandpass Signal Generation

To illustrate the capability to generate complex signals, a Gaussian minimum shift keying (GMSK) [66] signal is generated using $\Delta\Sigma$ digital phase modulation and then filtered with a charge-pump PLL. GMSK is the modulation method used in GSM, the digital mobile cellular standard deployed in Europe and in parts of North America. It is part of a class of modulation schemes called minimal-shift keying (MSK). MSK modulation exhibits two very important properties: constant envelope and the amenability to coherent detection. A GMSK signal is defined as

$$s(t) = \cos\left(2\pi f_C t + \frac{\pi}{2T_b}\int_o^t \sum_n b_n h(\tau - nT_b)d\tau\right), \qquad \text{(B.1)}$$

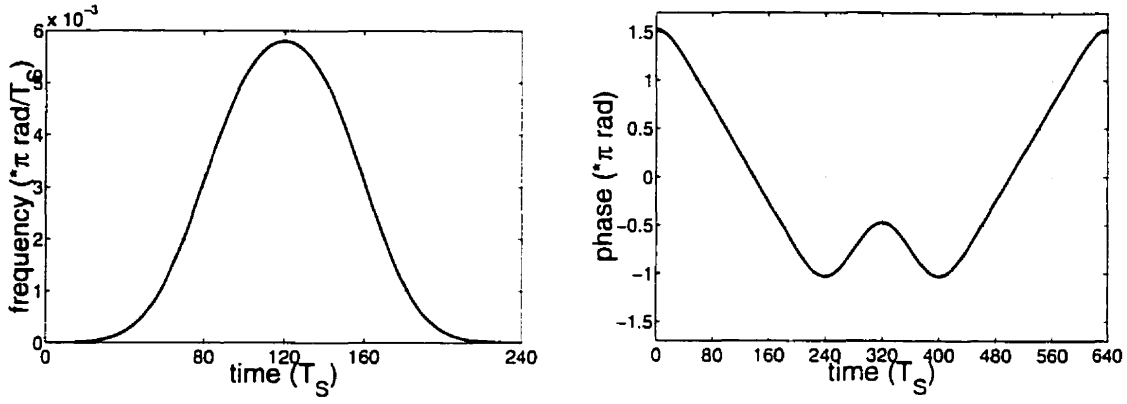

(a) (b)

Figure B.5: (a) GMSK frequency pulse. (b) Phase of one cycle of GMSK modulated sequence.

where f_C is the carrier frequency, T_b is the symbol interval, $\{b_n\}$ is the binary antipodal data sequence and $h(t)$ is a Gaussian filter. For the experiment, the clock frequency, f_S, is selected to be 4 MHz and the carrier frequency, f_C, is set to be 1 MHz. A bit rate, $1/T_b$, of 50 kHz is chosen, leading to a symbol interval consisting of 80 clock cycles. While the Gaussian filter theoretically extends over an unlimited number of symbol intervals, it decays quickly such that it can be well approximated using only three intervals. The resulting frequency pulse, $h(n)$, is shown in Figure B.5 (a). An 8-bit data sequence, {-1,-1,-1,+1,-1,+1,+1,+1}, is encoded resulting in a periodic signal extending over 640 clock cycles. One

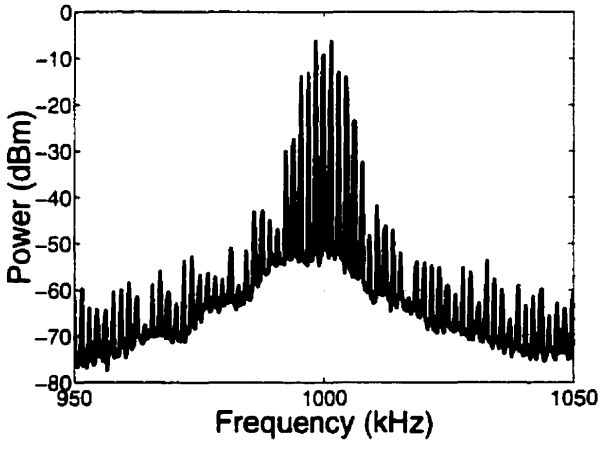

Figure B.6: Measured spectrum of GMSK signal.

period of this signal is shown in Figure B.5 (b). This periodic signal is then processed through a second-order lowpass $\Delta\Sigma$ modulator and a subset of the output is selected as described in Section A.5.3 of the Appendix and is used with the circuit of Figure 4.10. The measured spectral power is shown in Figure B.6. Again, the power density spectrum is discrete because of the periodic nature of the signal.

Another circuit for the generation of complex communication signals using a PLL was demonstrated a few years ago [67]. It uses $\Delta\Sigma$ fractional-N synthesis as explained in Section 4.2. This type of circuit modulates the output signal of a PLL using a counter in the feedback path. As argued in the above mentioned section, modulation of the input signal and of the output signal are quite similar except that the latter leads to more noise and distortion.

References

[1] B.M. Barber and M.N. Darrough, "Product reliability and firm value: the experience of american and japanese automakers, 1973-1992," *Journal of Political Economy*, vol. 104, no. 5, pp. 1084-1099, Oct. 1996.

[2] J.M. Miranda, "A BIST and boundary-scan economics framework," *IEEE Design and Test of Computers*, vol. 14, no. 3, pp. 17-23, July 1997.

[3] G. Singer, "The future of test and DFT," *IEEE Design and Test of Computers*, vol. 14, no. 3, pp. 11-14, July 1997.

[4] D.L. Wheater, P. Nigh, J.T. Mechler, L. Lacroix, "ASIC test cost/strategy tradeoff," *Proceedings of the 1994 IEEE International Test Conference*, pp. 93-102, 1994.

[5] G.W. Roberts, "Improving the testability of mixed-signal integrated circuits," *Proceedings of the IEEE 1997 Custom Integrated Circuits Conference*, pp. 214-221, Santa Clara, May 1997.

[6] B. Könemann, B. Bennetts, N. Jarwala and B. Nadeau-Dostie, "Built-in self-test: assuring system integrity," *Computer*, vol. 29, no. 11, pp. 39-45, Nov. 1996.

[7] F.M. Gardner, *Phaselock Techniques*, 2nd ed., New York, Wiley, 1979.

[8] L. DeVito, "A versatile clock recovery architecture and monolithic implementation," in *Monolithic phase-locked loops and clock recovery circuits: theory and design*, edited by B. Razavi, pp. 405-420, IEEE Press, New York, 1996.

[9] G. W. Roberts, "Metrics, techniques and recent developments in mixed-signal testing," *Proceedings of the 1996 IEEE/ACM International Conference on Computer Aided Design*, pp. 514-521, San Jose, California, Nov. 1996.

[10] C.-Y. Chao, H.J. Lin and L. Milor, "Optimal testing of VLSI analog circuits," *IEEE Transactions on CAD of Integrated Circuits and Systems*, vol. 16, no. 1, pp. 58-77, Jan. 1997.

[11] F.M. Gardner, "Charge-pump phase-lock loops," *IEEE Transactions on Communications*, vol. 28, no. 11, pp. 1849-1857, Nov. 1980.

[12] L. Bonet, J. Ganger, J. Girardeau, C. Greaves, M. Pendleton and D. Yatim, "Test features of the MC145472 ISDN U-transceiver," *Proceedings of the 1995 IEEE International Test Conference*, pp. 68-79, Washington, Oct. 1990.

[13] S.D. Huss and R.S. Gyurcsik, "Optimal ordering of analog integrated circuit tests to minimize test time," *Proceedings of the 1991 Design Automation Conference*, pp. 494-499, San Francisco, 1991.

[14] L. Milor and A. L. Sangiovanni-Vincentelli, "Minimizing production test time to detect faults in analog circuits," *IEEE Transactions on CAD of Integrated Circuits and Systems*, vol. 13, no. 6, pp. 796-813, June 1994.

[15] L. Milor and A. Sangiovanni-Vincentelli, "Optimal test set design for analog circuits," *Proceedings of the 1990 IEEE Integrated Circuits Computer-Aided Design Conference*, pp. 294-297, Santa Clara, 1990.

[16] J.P. Shen, W. Maly and F.J. Ferguson, "Inductive fault analysis of MOS integrated circuits," *IEEE Design and Test of Computers*, vol. 2, no. 6, pp. 13-26, Dec. 1985.

[17] D.B. Feltham and W. Maly, "Physical realistic fault models for analog CMOS neural networks," *IEEE Journal of Solid-State Circuits*, vol. 26, no. 9, pp. 1223-1229, Sep. 1991.

[18] M. Sachdev and B. Atzema, "Industrial relevance of analog IFA: a fact or a fiction," *Proceedings of the 1995 IEEE International Test Conference*, pp. 61-70, Washington D.C., Oct. 1995.

[19] F. Brglez, D. Pownall and R. Hum, "Accelerated ATPG and fault grading via testability analysis," *Proceedings of the 1985 IEEE International Symposium on Circuits and Systems*, pp. 693-698, 1985.

[20] B. Kaminska *et al.*, "Analog and mixed-signal benchmark circuits - first release," *Proceedings of the 1997 International Test Conference*, pp. 183-190, Washington D.C., Nov. 1997.

[21] M.J. Ohletz, "Hybrid built-in self-test (HBIST) for mixed analogue/digital integrated circuits," *Proceedings of the 1991 European Test Conference*, pp. 307-316, 1991.

[22] M.F. Toner and G.W. Roberts, "A BIST scheme for an SNR, gain tracking, and frequency response test of a sigma-delta ADC," *IEEE Trans. on Circuits and Systems-II*, Vol. 41, No. 12, pp. 1-15, Jan. 1995.

[23] A.K. Lu, G.W. Roberts and D.A. Johns, "A high-quality analog oscillator using oversampling D/A conversion techniques," *IEEE Transactions. on Circuits and Systems-II*, vol. 41, no. 7, pp. 437-444, July 1994.

[24] A.K. Lu and G.W. Roberts, "An analog multi-tone signal generator for built-in self-test applications," *Proceedings of the IEEE International Test Conference*, Washington, pp. 650-659, Oct. 1994.

[25] M.F. Toner and G.W. Roberts, "A BIST technique for a frequency response and intermodulation distortion test of a sigma-delta ADC," *Proc. IEEE VLSI Test Symposium*, Cherry Hill, pp. 60-65, Apr., 1994.

[26] H. de Bellescize, "La réception synchrone," *Onde Électrique*, vol. 11, pp. 230-240, June 1932.

[27] K.R. Wendt and G.L. Fredendall, "Automatic frequency and phase control of synchronization in television receivers," *Proceedings of the IRE*, vol. 31, pp. 7-15, Jan. 1943.

[28] C.A. Sharpe, "A 3-state phase detector can improve your next PLL design," *EDN Magazine*, pp. 55-59, Sep. 1976.

[29] B. Gilbert, "A versatile monolithic voltage-to-frequency converter," *IEEE Journal of Solid-State Circuits*, vol. 11, no. 6, pp. 852-854, Dec. 1976.

[30] A.A. Abidi and R.G. Meyer, "Noise in relaxation oscillators," *IEEE Journal of Solid-State Circuits*, vol. 18, no. 6, pp. 794-802, Dec. 1983.

[31] J.A. McNeill, "Jitter in ring oscillator," *IEEE International Symposium on Circuits and Systems*, vol. 6, pp. 201-204, London, May 94.

[32] T.C. Weigandt, B. Kim and P.R. Gray, "Analysis of timing jitter in CMOS ring oscillators," *IEEE International Symposium on Circuits and Systems*, vol. 4, pp. 27-30, May 94, London.

[33] J.P. Hein and J.W. Scott, "z-domain model for discrete-time PLLs," *IEEE Transactions on Circuits and Systems*, vol. 35, pp. 1393-1400, Nov. 1988.

[34] W. Goralski, *SONET: a guide to Synchronous Optical Network*, McGraw-Hill, New York NY, 1997.

[35] P. Trischitta and E. Varma, *Jitter in Digital Transmission Systems*, Artech House, Norwood MA, 1989.

[36] M. Lauterbach, "Getting more out of eye diagrams," *IEEE Spectrum*, v. 34, n. 3, pp. 61-63, March 1997.

[37] R.D. Gitlin, J.F. Hayes and S. B. Weinstein, *Data Communication Principles*, Plenum Press, New York, 1992.

[38] D-L Chen, "Designing on-chip clock generators," *Circuits and Devices Magazine*, pp. 32-36, July 1992.

[39] J. Alvarez, H. Sanchez, G. Gerosa and R. Countryman, "A wide-bandwidth low-voltage PLL for power PC microprocessors," *IEEE Journal of Solid-State Circuits*, vol. 30, pp. 383-391, Apr. 1995.

[40] E.H. Armstrong, "A method of reducing disturbances in radio signaling by a system of frequency modulation," *Proceedings of the IRE*, vol. 24, no. 5, pp. 689-740, May 1936.

[41] S. Kimura, M. Kimura, T. Nakatani and M. Sugai, "A new approach for PLL characterization on mixed-signal ATE," *Proceedings of the 1993 IEEE International Test Conference*, pp. 697-704, Oct. 1993.

[42] B.D. Kulp, "Testing and characterizing jitter in 100BASE-TA and 155.52 Mbits/s ATM devices with a 1 Gsamples/s AWG in an ATE system," *Proceedings of the 1996 IEEE International Test Conference*, pp. 104-111, Oct. 1996.

[43] G. Devarayanadurg, P. Goteti and M. Soma, "Hierarchy based statistical fault simulation of mixed-signal ICs," *Proceedings of the 1996 IEEE International Test Conference*, pp. 521-527, Oct. 1996.

[44] R.J.A. Harvey, A.M.D. Richardson, E.M.J.G. Bruls and K. Baker, "Analog fault simulation based on layout dependent fault models," *Proceedings of the 1994 IEEE International Test Conference*, pp. 641-649, 1994.

[45] P. Goteti, G. Devarayanadurg and M. Soma, "DFT for embedded charge-pump PLL systems incorporating IEEE 1149.1," *Proceedings of the IEEE 1997 Custom Integrated Circuits Conference*, pp. 210-213, Santa Clara, May 1997.

[46] Private communication with S. Sunter, November 1997.

[47] M.W. Hauser, "Principles of oversampling A/D conversion," *Journal of the Audio Engineering Society*, Vol. 39, No. 1/2, pp.3-26, Jan./Feb. 1991.

[48] B.R. Veillette and G.W. Roberts, "On-chip measurement of the jitter transfer function of charge-pump phase-locked loops," *Proceedings of the 1997 IEEE International Test Conference*, pp. 776-785, Washington D.C., Oct. 1997.

[49] P.J. Hurst and B.C. Rothenberg, "A programmable clock generator that uses noise shaping and its application in switched-capacitor filters," *IEEE Journal of Solid-State Circuits*, vol. 30, no. 4, pp. 403-411, Apr. 1995.

[50] T.A.D. Riley, M.A. Copeland and T.A. Kwasniewski, "Delta-sigma modulation in fractional-*N* frequency synthesis," *IEEE Journal of Solid-State Circuits*, vol. 28, no. 5, pp. 553-559, May 1993.

[51] R.K. Feldman, "A novel instrument for accurate time measurement in automatic calibration of test systems," *Proceedings of the 1993 International Test Conference*, pp. 544-551, Oct. 1993.

[52] T.E. Rahkonen and J.T. Kostamovaara, "The use of stabilized CMOS delay lines for the digitization of short time intervals," *IEEE Journal of Solid-State Circuits*, vol. 28, no. 8, pp. 887-894, Aug. 1993.

[53] J. Chapman, J. Currin and S. Payne, "A low-cost high-performance CMOS timing vernier for ATE", *Proceedings of the IEEE International Test Conference*, pp. 459-468, Oct. 1995.

[54] B.R. Veillette and G.W. Roberts, "Self-calibration of digital phase-locked loops," *Proceedings of the 1997 IEEE Custom Integrated Circuits Conference*, pp. 49-52, Santa Clara, California, May 1997.

[55] H. Inose, Y. Yasuda and J. Marakami, "A telemetering system by code modulation - Δ–Σ modulation," *IRE Transactions on Space, Electronic and Telemetry*, vol. SET-8, pp. 204-209, Sep., 1962.

[56] H. Inose and Y. Yasuda, "A unity bit coding method by negative feedback," *Proceedings of the IEEE*, vol. 51, pp. 1524-1535, Nov., 1963.

[57] R. Schreier and M. Snelgrove, "Bandpass Sigma-Delta Modulation," *Electronics Letters*, vol. 25, no. 23, Nov., 1989, pp. 1560-1561.

[58] J.C. Candy, "A use of double integration in sigma delta-modulation," *IEEE Transactions on Communications*, vol. COM-33, pp. 249-258, March 1985.

[59] J. Tierney, C.M. Rader and B. Gold, "A digital frequency synthesizer," *IEEE Transactions on Audio Electroacoustic*, vol. 19, pp. 48-57, 1971.

[60] J.G Kenney and L.R. Carley, "Design of multi-bit noise shaping data converter," *Analog Integrated Circuits and Signal Processing Journal (Kluwer)*, vol. 3, no. 3, pp. 259-272, May 1993.

[61] E.M. Hawrysh and G.W. Roberts, "An integration of memory-based analog signal generation into current DFT architectures," *Proceedings of the 1996 IEEE International Test Conference*, pp. 528-537, Washington, Oct. 1996.

[62] B. Dufort and G.W. Roberts, "Signal generation using periodic single and multi bit sigma-delta modulated streams," *Proceedings of the IEEE 1997 International Test Conference*, pp. 396-405, Washington, Nov. 1997.

[63] B.R. Veillette and G.W. Roberts, "Delta-sigma oscillators: versatile building blocks," *International Journal of Circuit Theory and Applications*, vol. 25, pp. 407-418, Sep.-Oct. 1997.

[64] B.R. Veillette and G.W. Roberts, "Bandpass signal generation using delta-sigma modulation techniques," *Proceedings of the 1995 IEEE International Symposium on Circuits and Systems*, pp. 637-640, Seattle, Washington, May 1995.

[65] I. Galton, "Spectral shaping of circuit errors in digital-to-analog converters," *IEEE Transactions on Circuits & Systems II*, v. 44, no. 10, pp. 808-817, Oct. 1997.

[66] K. Murota and K. Hirade, "GMSK modulation for digital mobile radio telephony," *IEEE Transactions on Communications*, vol. COM-29, vol. 7, pp. 1044-1050, July 1981.

[67] R.A.D. Riley and M.A. Copeland, "A simplified continuous phase modulator technique," *IEEE Transactions on Circuits and Systems-II*, vol. 41, no. 5, pp. 321-328, May 1994.