Frequency-Domain Compatibility in Digital Filter BIST

Laurence Goodby
Design Technology Center
Hewlett-Packard Company
Palo Alto, CA 94304

Alex Orailoğlu
Dept. of Computer Science & Engineering
University of California, San Diego
La Jolla, CA 92093

Abstract

We examine frequency-domain issues in the design and selection of on-chip test generators for built-in self-test (BIST) of high-performance digital filters. Test-generator/circuit compatibility is identified as a significant factor in testing large filters. A fault-injection experiment is used to show that when an incompatible test generator is used, high fault coverage (over 99%) does not guarantee that all serious faults will be detected. The frequency-domain characteristics of some basic test generation schemes are examined, and guidelines for test generator selection are proposed. Analytical techniques for identifying frequency-related testability problems are discussed, and several test generation schemes are evaluated by fault simulating them against lowpass, bandpass, and highpass filters. A mixed test generation scheme is shown to reduce the number of untested faults by a factor of two to three over a standard linear-feedback shift-register (LFSR) based test scheme, at little added cost.

1 Introduction

Frequency-domain techniques, through measures like total harmonic distortion, have long been used to characterize and test analog and mixed signal devices [1]. While the applicability of frequency-domain techniques to on-chip testing is less obvious, it nevertheless plays a significant role in the design of self-testing digital signal processing systems. Specifically, frequency-domain considerations are an important factor in the selection of a test generator for high-performance digital filter designs. A significant result is that common test pattern generators used for built-in self-test (BIST) can be behaviorally incompatible with certain classes of filters. By considering frequency-domain compatibility in the design of a test generator, coverage of an important set of faults can be significantly improved.

In digital filter datapaths, the bulk of the lower-order bits tends to be easy to test, while a small kernel of hard-to-test faults are typically found at the upper bits. The hard-to-test faults fall into two broad categories: those that are likely to be activated during normal operation, and those that are not. While not redundant, the last set of faults are only activated when the filter input is overdriven by a highly distorted signal, and thus fall outside the normal operating conditions of the filter. We will refer to these faults as *near-redundant*. These faults are naturally of less concern than the hard-to-test faults that can be activated by signals that fall within the filter's operating parameters. The last set of faults will be referred to simply as *difficult*.

Even taking the near-redundant faults into account, the small number of faults missed by a standard LFSR test sequence may include faults that correspond to serious flaws in the device. In fact, the relative ease with which the vast majority of faults in these devices can be tested tends to obscure the real testing problems, and dilutes the meaningfulness of fault coverage as an absolute measure of test quality.

Despite these concerns, BIST remains a particularly attractive approach for testing these designs for several reasons. First, large portions of the design can be tested with basic pseudorandom techniques requiring little additional chip area. Second, the high data rates that these devices generally operate at can make at-speed testing difficult using off-chip methods. Furthermore, such systems require the use of extremely low overhead techniques; performance is usually of the utmost concern, to the extent that many of the methods traditionally used to improve test access to the circuit-under-test (CUT) must be restricted. Using frequency-domain techniques, we can extend the effective range of a test generator, enabling the use of very low overhead BIST approaches employing a single generator at the input to the filter. Finally, by using frequency-domain analysis, we can choose a test generation scheme that significantly reduces the likelihood that a serious fault will escape detection.

Frequency-domain techniques are not the complete answer to digital filter testing: redundant logic, correlation effects, and random-pattern test resistance are some of the many test problems that may be encountered [2, 3, 4, 5]. However, frequency-domain considerations do rank as one of the most important issues in digital filter BIST. To understand how frequency-domain behavior impacts testability, we will first lay some groundwork by modeling the most difficult-to-test faults. We will then see, through a fault-injection experiment, how neglecting frequency-domain considerations can result in difficult faults escaping detection despite deceptively high single-stuck-at fault coverage. We will analyze a number of different test generation schemes from the frequency-domain perspective, and discuss their compatibility with some of the basic filter types. We will then compare the actual performance of several test generation schemes on three filters representing the basic lowpass, bandpass, and highpass filter types. Finally, we will look at the advantages of mixed test generation schemes.

2 Notation & Conventions

Signals are assumed to be represented using two's-complement arithmetic, where the value of an N-bit signal $b_0, b_1, \ldots, b_{N-1}$ is given by $-b_0 + \sum_{i=1}^{N-1} b_i 2^{-i}$. To aid discussion, all signal values are expressed relative to the bit width available at that point in the circuit. For example, a 5 bit signal will be interpreted as a two's-complement number in the range -1 to 1, even though it might actually represent the lower 5 bits of a 6-bit signal that is restricted to the interval $[-0.5, 0.5)$. In the fault simulation results, we assume no aliasing in the response analyzer.

3 Application Domain

The circuits we will consider here are high-performance, reduced-complexity digital filters. The circuit elements to be tested include delay registers, adders, subtractors, and fixed-coefficient multiplication operators. The multiplication operations are implemented via hardwired shift-and-add structures. We will assume that the adders and subtractors are implemented via ripple-carry addition. Carry-save adder arrays are a higher-performance alternative that come at the cost of doubling the number of registers in the design. The frequency-domain analysis that is described here applies to circuits implemented using either ripple-carry or carry-save adders, although the analysis is more complex in the case of carry-save arrays [4]. Consequently, we will focus on ripple-carry adders here. At the

This work was supported in part by grants from the Semiconductor Research Corporation (contract DJ-538), the University of California MICRO Program, and Hughes.

DAC 97, Anaheim, California
©1997 ACM 0-89791-920-3/97/06..$3.50

design	adders	regs	widths			faults
			in	coef.	out	
LP	183	60	12	15	16	57148
BP	161	58	12	14	16	50650
HP	175	60	12	15	16	55042

Table 1: Design statistics.

register-transfer level, the designs can be represented as networks of registers, adders, subtractors, fixed-shift, and sign-extension operators. The filters are implemented using a cascade of "tap" structures, each of which corresponds to a constant-multiplication operation and a delay register. Designs typically consist of anywhere from a dozen adders and subtractors to several hundred. Further details on the architecture of the filters may be found elsewhere [6].

Three basic types of filters are examined: lowpass, bandpass, and highpass. Three representative designs were created using a canonic-signed digit representation to convert multiplier coefficients to a small number of add and subtract operations [6, 7, 8]. Statistics for the designs are summarized in Table 1. The complexity of each design is similar: the number of registers (corresponding to the number of filter tap structures) was kept at or near 60, and the number of adders in the most complex design is within 14% of the number of adders in the simplest design. The input signal width, maximum coefficient width, and the output datapath width are also similar in all three designs. The number of adder faults modeled in each design is shown in the rightmost column. Register faults are not considered, since they did not pose any obstacle to full testing of this architecture.

The use of scaling techniques to identify and remove redundant sign bits is the first step towards obtaining a testable design. Once this is accomplished, further optimizations can be performed on the upper bits of many adders to eliminate redundancies that are induced by signal constraints [2, 3].

4 Modeling the Difficult Faults

Out of the adder faults modeled in these designs, usually only 1% or less are classified as difficult—those that lie well beyond the knee of the fault simulation curve. In this section, we will look at a high-level model of these difficult faults, which will enable us to relate test signal variance to fault coverage. The two basic mechanisms that result in random-pattern test resistance in digital filters will be described, and we will discuss the problem of near-redundant faults.

Difficult faults will be addressed in terms of difficult *tests*; whether or not all difficult tests are required will depend on the gate-level fault model used. In the most conservative test model, all difficult tests are considered essential. In other commonly used models, not all difficult faults are essential; their faults can be handled by easier-to-apply tests [5].

The hardest faults to test in the digital filters described here typically lie in the carry logic of the bits closest to the MSB. There are two closely-related mechanisms at work: adder-input *variance-mismatch* and *excess headroom*. The first occurs in adders where one input (the secondary input) is usually much smaller in magnitude than the other input (the primary input); for example, when adding a 4-bit signal to a 16-bit signal. The excess headroom problem occurs when a signal does not use the full dynamic range available to it; this occurs, for example, when a small signal is added to a larger signal, requiring the datapath width to expand to hold the result. If the resulting signal rarely uses the added dynamic range, a test problem results. Excess headroom is often associated with conservative scaling techniques, where one or more upper bits effectively act as redundant sign bits [5], making it extremely difficult to activate overflow conditions at the next-to-MSB adder. Some of these faults have such a small probability of being excited during normal operation of the filter that testing them may naturally take a lower priority—they are considered near-redundant. Distinguishing between faults that are near-redundant and those that are merely difficult usually relies on some knowledge of the input signal's worst-case parameters.

Test	Input	Output
T1a	$0 \le A < 0.5$	$A + B \ge 0.5$
T1b	$A < -0.5$	$A + B \ge -0.5$
T2a	$0 \le A < 0.5$	$A + B < 0$
T2b	$A < -0.5$	$A + B \ge 0.5$ (ovf)
T5a	$-0.5 \le A < 0$	$A + B \ge 0$
T5b	$A \ge 0.5$	$A + B < -0.5$ (ovf)
T6a	$-0.5 \le A < 0$	$A + B < -0.5$
T6b	$A \ge 0.5$	$A + B < 0.5$

Table 2: The four difficult test equivalence classes for an adder's next-to-MSB carry logic. 'A' refers to the value of the high-variance primary input, while 'B' refers to the low-variance secondary input.

The hardest faults to test are generally associated with excess headroom; for random-pattern testing, faults due to variance-mismatch generally have expected test lengths of at most a few thousand vectors, while excess headroom can account for expected test lengths in the hundreds of thousands of vectors, or more. These two types of faults are closely related in that a series of variance-mismatched additions can lead to excess headroom when the datapath width expands, and excess headroom in turn becomes a problem when the sum is fed into a variance-mismatched adder. Through bit-level analysis of variance-mismatched adders, it is possible to identify the most difficult tests to apply to upper adder bits. By translating these gate-level tests into behavior-level conditions on the adder's inputs and output, we can more easily identify test problems through analysis of the filter's behavior. We will describe these conditions in the following section.

4.1 I/O test conditions

At each full adder cell, eight tests are possible; we will identify these as T*n*. The test number, *n*, corresponds to the value of the binary number *abc*, where *a* is the value of the primary input bit; *b*, the secondary input bit, and, *c*, the carry input. Of these tests, half are considered difficult in variance-mismatched adders. The input-output conditions required to assert each of the four difficult tests at the next-to-MSB adder are shown in Table 2 (the MSB logic is less of a test problem since it does not contain any carry logic). Each test can be asserted by two equivalent test classes, labeled *a* and *b*. Two tests (marked *ovf*) correspond to testing the overflow behavior of the adder. Even if overflow cannot occur at the adder in question, the other tests in the T2 and T6 equivalence classes may be required, depending on the gate-level fault model used.

Excess headroom corresponds to a very low probability of an adder output exceeding magnitude 0.5 (and possibly lower thresholds, as well). This in turn constrains the primary input, making tests T1 and T6 difficult. In some gate-level fault models, test T6 is non-essential, in which case test T1 poses the greatest problem, as it is generally considered essential [5]. An example of this will be shown in Figure 3. Even in the absence of a severe headroom problem, tests T1 and T6 remain susceptible to reduced signal variance, as will be seen in the next section.

4.2 Test zones

In variance-mismatched adders, the secondary input has much lower variance than the primary input. Taken in conjunction with the input-output conditions shown in Table 2, this places a tight constraint on the values the primary input can take in order to activate the tests. The location of these test zones is shown in Figure 1, where the width of the test zones is proportional to the variance of the secondary input. This test model shows the importance of achieving high test signal variance throughout the datapath if the upper bits are to be adequately tested, as tests T1 and T6 can only be activated by signals near amplitude 0.5. This conclusion, while intuitively obvious, is not always clear from fault simulation results, which can report high fault coverage even when important faults have been missed due to test signal attenuation. Tests T2 and T5, on the other hand, are less susceptible to the effects of excess headroom. If these tests are missed, it is usually due only to a variance-mismatch problem,

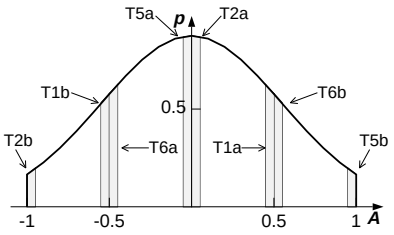

Figure 1: A hypothetical probability density function for the primary input to an adder. The shaded bars indicate the zones that the input must fall in for the difficult tests to be asserted.

which is more easily solved by extending the test length.

A similar analysis is possible at bits below the next-to-MSB. The test problem rapidly diminishes towards the lower bits, as the number of equivalent tests doubles for each step down from the MSB. Further details on the probability of asserting the difficult tests using random-pattern testing can be found elsewhere [5].

5 When 99% Isn't Enough

It is difficult to reach 100% single-stuck fault coverage on large filters using only a single test generator at the input to the filter and a compressor at the output. Despite very good observability of most signals, controllability can be much more limited due to the difficult and near-redundant faults described in the previous section. As discussed, the most difficult of these faults commonly result from the use of conservative design techniques, and can be considered effectively redundant in many cases. For example, assuming some knowledge of the input signal statistics, it may be possible to identify adder outputs at which the probability of the next-to-MSB output bit overflowing is less than 10^{-6} under worst-case conditions [4]. A much higher probability of overflow is often considered acceptable during the design of filter coefficients, but this information is not usually used to eliminate near-redundancies.

Assuming less than 100% coverage, how much coverage is enough to ensure that difficult faults will not escape detection? While this question cannot be completely answered without some knowledge of the input signal's worst-case statistics, consideration of the test model presented in the last section makes it clear that any faults that are missed because a test signal is attenuated with respect to the normal operating signal constitutes a serious test failure. A test signal can become attenuated if the signal generator does not place an adequate amount of power in the filter's passband. In the remainder of this section, we will show an example of a serious fault that is missed due to test signal attenuation despite high overall fault coverage.

A common BIST approach is to use an LFSR to generate a test signal that is applied to the filter's input. For the design in question, a 60-tap lowpass filter, this approach results in 99.1% fault coverage of the combinational logic; the fault coverage for the entire design would be somewhat higher due to the complete testing of the registers in the design. We might be tempted to stop testing at this point, assuming the remaining untested faults to be near-redundant due to conservative design techniques, i.e., the remaining faults can only be activated by input signals that would never occur under normal operating conditions. However, as we will see later, the LFSR test signal becomes highly attenuated as it passes through the filter. This results in upper-bit missed faults that have a relatively high probability of being excited during normal operation of the filter.

Injecting a fault missed due to test signal attenuation reveals the severity of these faults; Figure 2 shows the output from the faulty 60-tap lowpass filter with a sine wave signal applied to the input. The fault effect is visible as a spike (actually, a closely-spaced pair of spikes) at the peak of the output sine wave. While somewhat sensitive to the amplitude and frequency of the sine wave, this fault should nonetheless be considered a serious flaw, as it can be excited by a wide range of signals that fall within the filter's normal operating parameters. Figure 3 shows the location of the fault in question.

6 Test Generator Characterization

Common BIST test pattern generators exhibit a wide variation in their frequency-domain characteristics. In this section, we will examine the frequency-domain behavior of several of the most common schemes, and look at their compatibility with different filter behaviors. We will also look at some variations that can be used to improve the compatibility of the test generator with the CUT.

The most common test generator that we will look at, the linear-feedback shift-register, falls into two basic implementation schemes: Type 1 with an external XOR-tree, and Type 2 with embedded XORs [9]. Variations on these LFSRs include the use of a decorrelator circuit to reduce the linear correlation between successive test vectors, and a maximum-variance version which uses one bit per test rather than the entire contents of the LFSR. An entirely different type of test generator, the Ramp generator, is based on counters rather than LFSRs [10]. In all cases, the output of the test generator is interpreted as a two's-complement number in the interval $[-1, 1)$.

Type 1 LFSR: These LFSRs exhibit reduced power at low frequencies due to negative correlation between successive words. The signal variance is 0.3333, giving an average power of -4.77 dB. This LFSR's power spectrum is not sensitive to the particular seed or polynomial used as long as the bit stream generated has reasonable properties, i.e., the probability of generating a 1 should be approximately equal to the probability of generating a 0, and successive bits should be uncorrelated. This is generally satisfied by choosing a primitive polynomial. The curve labeled "LFSR-1" in Figure 4 shows the spectrum of a 12-bit Type 1 LFSR. The shape of the spectrum is not particularly sensitive to the bit width of the LFSR, nor is it sensitive to the direction of shifting (whether MSB-to-LSB or LSB-to-MSB). It can be altered by some permutations of the output bits; an interconnection network can be used at the output of the LFSR to accomplish this.

Type 2 LFSR: Unlike the Type 1 LFSR, this LFSR's frequency-domain characteristics are dependent on the choice of polynomial and shift direction. Generally, Type 2 LFSRs will have less low-frequency rolloff than their Type 1 counterparts. Choosing a polynomial that puts an XOR gate near the MSB can help flatten the spectrum. The power spectrum of a 12-bit Type 2 LFSR with polynomial 12B9h (shift direction LSB-to-MSB) is shown as curve "LFSR-2" in Figure 4. In some cases, using the reciprocal polynomial will help the frequency-domain characteristics by moving an XOR gate closer to the MSB. Like the Type 1 LFSR, the signal variance is 0.3333.

Decorrelated LFSR: One way of flattening the Type 1 LFSR spectrum is to attach a decorrelator circuit to its output. This can take a number of forms—in our experiments we used an XOR network to invert all bits other than the LSB whenever the LSB takes a 1 value. In terms of gate count, this is not the most efficient way to flatten the spectrum (we can do this with fewer gates using a Type 2 LFSR), but it does provide other desirable properties: the correlation between all bits in two successive vectors is reduced, rather than just a few bits. This can help eliminate some test problems, such as structural dependencies [11, 2]. It retains some properties of maximal-length shift-register sequences that may be desirable from a testing perspective, such as no repeated vectors and a near-zero mean value. As with maximal-length sequences, the variance of the generated signal is approximately 0.3333. The curve labeled "LFSR-D" in Figure 4 corresponds to the power spectrum of a 12-bit decorrelated Type 1 LFSR. The signal provides essentially equal power to all frequency bands.

Maximum-variance LFSR: The above three test generators all have signal variances of 0.3333. A higher variance signal can be generated from an LFSR simply by using the bit stream to select between the most positive two's-complement number or the most negative, yielding a signal with variance 1 and a flat spectrum like the decorrelated LFSR. Due to its high-variance, this test generator

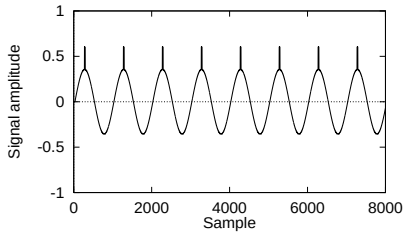

Figure 2: Despite reaching fault coverage over 99.1%, a standard LFSR-generated test sequence misses this fault. The output of a lowpass filter is shown corresponding to a sine-wave input signal. The fault effect is visible as a spike train superimposed on the output sine wave.

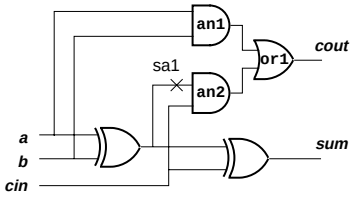

Figure 3: Location of the full-adder fault corresponding to Figure 2. The fault is located three bits down from the MSB of tap 20 of a 60-tap lowpass filter. This fault is only detected by the difficult test T1.

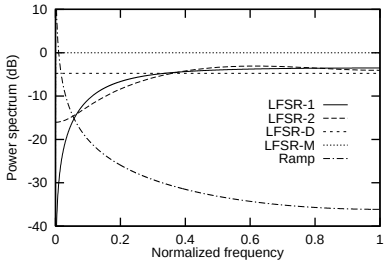

Figure 4: Power spectra of some common BIST test pattern generators.

	Lowpass	Bandpass	Highpass
LFSR-1	−	±	+
LFSR-2	±	±	+
LFSR-D	+	+	+
LFSR-M	+	+	+
Ramp	+	−	−

Table 3: Frequency-domain compatibility of test generators and filter types.

is relatively good at exercising the upper bits in the datapath, but does not effectively test lower bits in most adders due to the correlation between adjacent bits. In the lower bits, tests T1 and T5 are missed by this generator. The spectrum is shown as the curve labeled "LFSR-M" in Figure 4.

Ramp: Since counters are frequently available on-chip as part of a design, they are sometimes used as test signal generators [10]. In two's-complement notation, the signal produced is a ramp or sawtooth function. The frequency-domain behavior of this signal is substantially different from the LFSR-based signal generators: almost all of the signal's power is concentrated at very low frequencies. The spectrum of a 12-bit count-by-one ramp is shown in Figure 4 as the curve labeled "Ramp".

6.1 Frequency-domain compatibility

Based on the above frequency-domain characterization of test signal generators, we can judge the compatibility of each generator with the circuit to be tested. Following the discussion in Section 5, some of the most serious missed faults correspond to those that are due to the failure of the test generator to put adequate power into the filter's passband. Comparing the filter's transfer function with the test generator's spectrum thus gives a quick indication of their compatibility. Formally, we can estimate the output signal variance as $\sigma_y^2 = \frac{1}{L}\sum_{k=0}^{L-1} |G[k]|^2 |H[k]|^2$, where $G[k]$ is the discrete power spectrum of the pattern generator's signal, $H[k]$ is the Discrete Fourier Transform (DFT) of the filter's impulse response, and L is the length of the DFT. A mismatch between the shape of $G[k]$ and $H[k]$ results in reduced output signal variance, which impacts other taps as well, especially those closest to the output. Variance-based testability analysis will be revisited in Section 7.

The compatibility of each test generator with the three basic filter types is shown in Table 3, where '+' indicates good compatibility, '−' indicates poor compatibility, and '±' indicates that the compatibility is dependent on the specifics of the design. The LFSR-1 generator is compatible with bandpass filters as long as the passband is far enough above the LFSR-1's rolloff area. The frequency-domain characteristics of LFSR-2 generators place them between the LFSR-1 and LFSR-D generators, depending how flat the spectrum is, which in turn depends on the polynomial selected.

7 Analyzing LFSR-based Testing

There are a variety of techniques that can be used to identify potential test problems early in the design process, based on consideration of the filter's impulse response. In particular, it is possible (using signal variance analysis) to identify points in the design where the test signal's standard deviation is small compared to the bit width available. More advanced techniques are also available that are based on computing the signal probability distributions at each adder [5]. In this section we will show an example of how the linear model of an LFSR can be used to identify test signal attenuation problems through signal variance analysis. Simulation-based techniques will also be used to examine the test problem introduced in Section 5, and these results will be compared with a distribution-based approach.

7.1 Variance-based analysis

In a linear system, we can characterize the output of an adder by the impulse response corresponding to the subsystem that outputs at that adder. If the impulse response corresponding to the k-th adder is denoted by h_k, then the variance of the adder's output in response to a white-noise source of variance σ_x^2 is

$$\sigma_k^2 = \sigma_x^2 \sum_{i=0}^{M_k} h_k^2[i], \tag{1}$$

where M_k is the order of the subfilter that outputs at adder k, assuming a finite impulse response (or a finite approximation of an infinite impulse response). This basic result has been used extensively in analyzing roundoff noise in digital filters (see, for example, [12, Eqn. 6.107]), and can be directly applied to finding the variance at an adder's output due to a decorrelated LFSR source. For the LFSR-D generator, assuming perfect decorrelation, σ_x^2 is 0.3333, while for the maximum-variance LFSR (LFSR-M), σ_x^2 is 1, neglecting any overflow effects.

For other LFSRs, many correlation effects can be accounted for using a linear model of the LFSR-generated signal. For example, an MSB-to-LSB-shifting N-bit Type 1 LFSR can be modeled as a 0/1 white-noise source feeding a linear system characterized by the finite impulse response [5],

$$g[n] = \begin{cases} -1, & n = 0, \\ 2^{-n}, & n = 1, 2, \ldots, N-1, \\ 0, & \text{otherwise.} \end{cases}$$

When this model of the LFSR is cascaded with the CUT, the response due to an impulse at the LFSR model's input is given by the convolution of the model's impulse response with the subfilter's impulse response: $h_k'[n] = h_k[n] * g[n] = \sum_{i=0}^{M_k} h_k[i]g[n-i]$, for $n = 0, 1, \ldots, N + M_k - 1$ (zero otherwise). The variance of the adder's output signal is then given by Equation 1, replacing h_k with h_k' and setting σ_x^2 equal to 0.25, the variance of a 0/1 white-noise source with $P\{0\} = P\{1\} = 0.5$. A similar approach can be used for Type 2 LFSRs, although this involves computing a response due to each XOR gate embedded in the LFSR, and summing the variances due to each.

The impulse response model of LFSRs also determines their frequency-domain behavior; the power spectrum is the Discrete Fourier Transform (DFT) of the aperiodic autocorrelation function of the model's impulse response, which is given by $h_k[n] * h_k[-n]$. For Type 2 LFSRs, the power spectrum due to each XOR gate is added to compute the complete power spectrum.

To examine the effectiveness of the test generator, we can examine the variance at each adder's output; a low signal variance

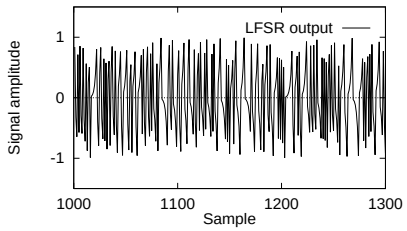

Figure 5: A segment of the test sequence generated by a 12-bit Type 1 LFSR. The standard deviation of the maximal-length sequence is 0.577.

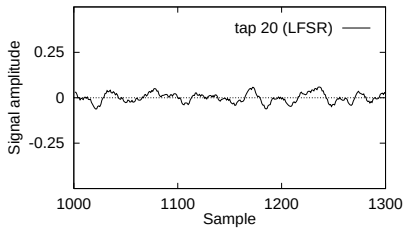

Figure 6: Reduced variance test signal at tap 20 of the 60-tap lowpass filter, corresponding to the LFSR-generated test sequence. The four bits below the MSB are not fully tested. Standard deviation: 0.036.

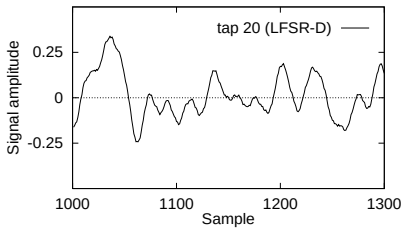

Figure 7: Test signal at tap 20 with a decorrelator attached to the LFSR. This time, only one bit below the MSB is not fully tested. Standard deviation: 0.121.

Des.	LFSR-1	LFSR-D	LFSR-M	Ramp
LP	519	331	1097	485
BP	201	193	1005	1230
HP	308	315	1030	1679

Table 4: Missed faults.

Des.	LFSR-1	LFSR-D	LFSR-M	Ramp
LP	2.84	1.81	5.99	2.65
BP	1.25	1.20	6.24	7.64
HP	1.76	1.80	5.89	9.59

Table 5: Normalized missed faults.

Des.	misses	normalized
LP	148	0.81
BP	74	0.46
HP	137	0.78

Table 6: Missed faults for a mixed LFSR-1/LFSR-M test (8k test length).

indicates the presence of a potential testing problem. This analysis permits testability problems to be identified early in the design process. Other approaches include simulation-based analysis, and analytical techniques based on probability distribution analysis. In the next section, we will take a look at how these techniques can be used to identify reduced test coverage points.

7.2 Simulation-based analysis

Continuing the example from Section 5, we will next use simulation-based techniques to look at how an LFSR-generated test sequence fails to detect the fault described. The output of a 12-bit LSB-to-MSB-shifting Type 1 LFSR is fed into the filter's input; interpreting this signal as a two's-complement number results in the waveform shown in Figure 5, a 300-sample segment of the test sequence. The short exponential segments are characteristic of the correlation properties of this LFSR. The filter is the 60-tap lowpass filter.

The response excited by the test signal at tap 20 is shown in Figure 6. Clearly, the test signal is severely attenuated at this point in the filter. Consequently, neither the adder in question nor the logic dependent on it will be fully tested. In all, the carry logic of the four consecutive bits below the MSB is not tested. This attenuation is due to the low-frequency rolloff of the LFSR acting in combination with the low cutoff frequency of the filter, and is predicted by the signal variance analysis technique described in Section 7.1.

The low-frequency rolloff of the LFSR-1 generator can be eliminated by placing a decorrelator at the LFSR's output. Although this test signal has the same variance as the original test signal, the standard deviation of the signal at tap 20 is now 3.4 times higher than in the previous case. As can be seen from Figure 7, the signal amplitude is much improved although still insufficient to test the next-to-MSB bit; this corresponds to one of the more difficult tests due to the conservative scaling used at this tap. This bit can be tested by a long maximum-variance LFSR sequence.

While signal variance is useful in analyzing this kind of attenuation problem, it is a very rough measure. More precise analysis of the test problem is possible if the shape of the signal's probability distribution is known. This can be obtained by histogram analysis of simulations, or can be found using analytical techniques when a suitable model of the test generator is available. The latter approach can be used for the LFSR-based test generators; an example is shown in Figure 8. This shows the predicted distribution at tap 20 for the LFSR-1 test signal, as well as the histogram estimate produced through simulations. Comparing this with the test zones shown in Figure 1 highlights the nature of the test problem.

A histogram corresponding to a decorrelated LFSR-1 generator (LFSR-D) is shown in Figure 9. This figure also shows the expected distribution for an idealized test generator, assuming statistically independent vectors. The LFSR-D histogram, while not matching as closely as the previous distribution, still matches fairly well, attesting to the efficacy of the decorrelating circuit. Details of distribution-based testability analysis can be found elsewhere [5].

8 Experimental Results

We will now look at the performance of four different test generators on the three designs described in Section 3, and compare their performance with that expected from frequency-domain considerations. Figures 10–12 show fault simulation results for the LFSR-1, LFSR-D, LFSR-M, and Ramp generators described previously. A 12-bit version of each generator was used; for the LFSR-based generators it is easy to extend the test length by using a larger LFSR. Table 4 shows the number of faults that remained undetected after 4k vectors; in Table 5 these numbers are normalized by the number of adders and subtractors used in the design.

In the highpass and bandpass designs, the LFSR-1 and LFSR-D generators give very similar performance. However, in the lowpass design, the LFSR-1 generator clearly lags behind. This is as expected, due to the low-frequency rolloff of the LFSR-1 generator which was shown in Figure 4. While the percentage difference in fault coverage is very small (0.3%), this consists of faults that we cannot afford to miss, for the reasons discussed in Section 5.

The maximum-variance LFSR (LFSR-M), lags all other test generators due to its poor coverage of faults in lower bits of adders. However, it is usually more effective at testing upper bits than other generators, and—due to its flat frequency response—gives similar performance in all three cases. We shall see in Section 9 how this test generator's ability to test upper bits can be combined with another test generation mode to reach very high coverage levels.

The Ramp generator is comparable to the best generators for the lowpass filter, but—as expected—performs poorly on the bandpass and highpass filters due to its extremely low power in the upper bands. For the wide-band test generators (which includes all but the Ramp generator), the bandpass filter is somewhat easier to test than the other two filters partly due to its wider passband: less test signal power is lost to stopband attenuation. The exception to this is the LFSR-M generator, which encounters a higher percentage of low-order bits that it is unable to test.

9 Mixed Test Generation Schemes

A low-cost pseudorandom BIST scheme that can offer good coverage for many designs consists of an LFSR that is switched between normal and maximum-variance mode. The use of maximum-variance mode not only helps exercise faults in upper bits through its increased signal variance, but also compensates for the Type 1 LFSR's low-frequency rolloff. While the maximum-variance test signal has good signal strength in all frequency bands, it has poor coverage of the lower bits in adders. Combining the two approaches

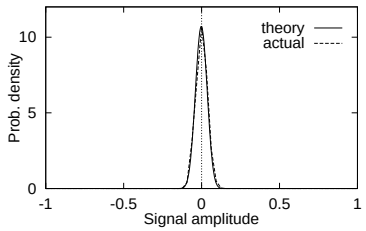

Figure 8: Signal amplitude distribution at tap 20 due to a Type 1 LFSR. The 'theory' curve uses the LFSR linear model to predict the distribution, while the 'actual' curve shows the histogram estimate obtained by simulation. The histogram estimate closely matches the curve predicted by theory.

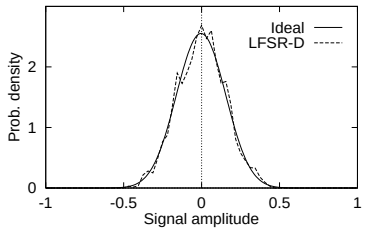

Figure 9: Signal amplitude distributions at tap 20, decorrelated tests. The theoretical curve assumes an idealized test generator producing statistically independent vectors, while the 'LFSR-D' curve shows a histogram corresponding to the use of a decorrelated LFSR.

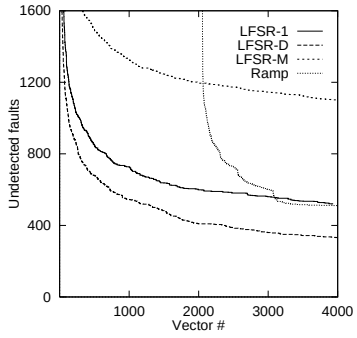

Figure 10: Fault simulation results of four test generators applied to the lowpass filter.

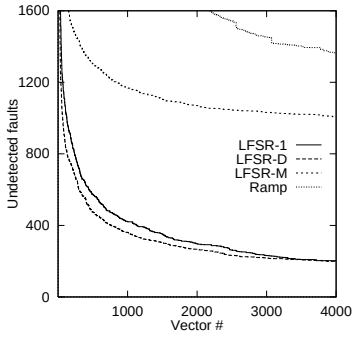

Figure 11: Fault simulation results for the bandpass filter.

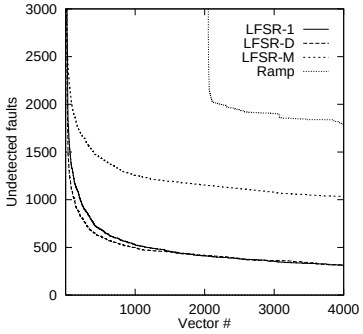

Figure 12: Fault simulation results for the highpass filter. Note that the vertical scale has been changed to accommodate the Ramp curve.

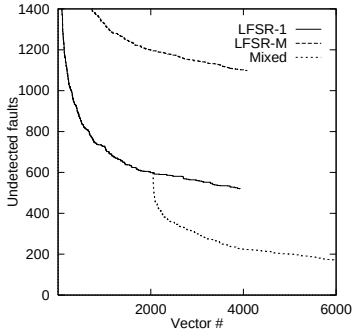

Figure 13: Fault simulation results showing the advantage of combining test generators for the lowpass filter. The curves correspond to a Type 1 LFSR, a maximum-variance LFSR, and the effect of switching the LFSR to maximum-variance mode after 2k vectors.

results in much better test coverage than can be obtained by either approach in isolation. This effect is shown in Figure 13, which shows fault simulation curves for the lowpass filter design. The LFSR is started in normal mode, and then switched to maximum variance mode after 2k vectors. The fault coverage results for a 4k normal-mode plus 4k maximum-variance mode Type 1 LFSR are shown in Table 6. The quality of this test is very similar to that of a mixed LFSR-D/LFSR-M test scheme, yet does not require a decorrelating circuit.

10 Conclusion

In testing high-performance digital filters, adequate testing cannot be guaranteed by high fault coverage alone; the significance of any untested fault depends on the likelihood of fault activation during normal operation of the filter. To be effective, a test generation scheme should put a substantial amount of energy in the filter's passband. Use of an incompatible test generator, such as a Type 1 LFSR with a narrow-band lowpass filter, or a ramp with a highpass filter, can lead to surprises, where important faults are missed despite deceptively high fault coverage.

Combining a CUT-compatible test generator with an incompatible test generator that has other important test properties can provide higher fault coverage than either generator in isolation. Often, this is enough to ensure good coverage of the most significant difficult faults, reducing the number of untested faults by a factor of 2.2 to 2.6 over the best single-mode approaches, and by as much as a factor of 3.5 over basic LFSR-based testing. To provide an even higher guarantee of coverage of difficult faults, some possible approaches include: restructuring the design to improve its random-pattern testability; use of longer test sequences (with larger LFSRs to avoid input cycling); use of more specialized test controllers to produce tests tailored to the specific filter (deterministic BIST); use of more aggressive scaling techniques, when appropriate; and identification of near-redundant faults, excluding them from the fault universe.

The intent of the last two points is to formalize the process of deciding which faults (if any) will not be tested, with the aim of reaching 100% coverage on the most critical faults. This typically requires some knowledge of the worst-case input signal statistics. Since this information is not always available, the test designer must often rely instead on very-high-coverage techniques based on the use of frequency-domain compatible test generators.

References

[1] J. A. Mielke, "Frequency domain testing of ADCs," *Design & Test of Computers*, vol. 13, no. 1, pp. 64–69, 1996.

[2] L. Goodby and A. Orailoğlu, "Towards 100% testable FIR digital filters," in *International Test Conference*, pp. 394–402, Oct. 1995.

[3] L. Goodby and A. Orailoğlu, "Synthesizing self-testable filters via scaling and redundant operator elimination," in *29th Asilomar Conf. Signals, Systems, and Computers*, vol. 1, pp. 132–136, Oct. 1995.

[4] L. Goodby and A. Orailoğlu, "Variance mismatch: Identifying random-test resistance in DSP datapaths," in *Intl. Conf. on Acoustics, Speech & Signal Processing*, vol. 6, pp. 3205–3208, May 1996.

[5] L. Goodby and A. Orailoğlu, "Pseudorandom-pattern test resistance in high-performance DSP datapaths," in *Design Automation Conference*, pp. 813–818, June 1996.

[6] R. Jain, P. T. Yang, and T. Yoshino, "FIRGEN: A computer-aided design system for high performance FIR filter integrated circuits," *IEEE Trans. Signal Processing*, vol. 39, pp. 1655–1668, July 1991.

[7] H. Samueli, "An improved search algorithm for the design of multiplierless FIR filters with powers-of-two coefficients," *IEEE Trans. on Circuits and Systems*, vol. 36, pp. 1044–1047, July 1989.

[8] T. Yoshino, R. Jain, P. T. Yang, H. Davis, W. Gass, and A. H. Shah, "A 100-MHz 64-tap FIR digital filter in 0.8-μm BiCMOS gate array," *IEEE J. of Solid-State Circuits*, vol. 25, pp. 1494–1501, Dec. 1990.

[9] M. Abramovici, M. A. Breuer, and A. D. Friedman, *Digital Systems Testing and Testable Design*. Computer Science Press, 1990.

[10] S. Gupta, J. Rajski, and J. Tyszer, "Test pattern generation based on arithmetic operations," in *Intl. Conf. on Computer-Aided Design*, pp. 117–124, 1994.

[11] P. H. Bardell, W. H. McAnney, and J. Savir, *Built-In Test for VLSI: Pseudorandom Techniques*. Wiley, 1987.

[12] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*. Prentice Hall, 1989.