

Test Quality and Fault Risk in Digital Filter Datapath BIST

Laurence Goodby
Design Technology Center
Agilent Technologies
Palo Alto, CA 94304
laurence.goodby@agilent.com

Alex Orailoğlu
Dept. of Computer Science & Engineering
University of California, San Diego
La Jolla, CA 92093
alex@cs.ucsd.edu

Abstract

An objective of DSP testing should be to ensure that any errors due to missed faults are infrequent compared to a circuit's intrinsic errors, such as overflow. A method is proposed for quantifying test quality for digital filters by measuring the risk associated with any untested faults. Techniques for finding upper bounds on fault activation rates under worst-case operating conditions are described. These techniques enable test designers to objectively discriminate significant missed faults from near-redundant faults, which are unlikely to be activated in normal operation of the device. This complements fault coverage as a measure of test quality, providing a means of locating high-risk missed faults even in very high coverage test regimes.

1 Introduction

Digital filter datapaths pose a number of unique challenges to testing. In these designs, full scan techniques suffer from high overhead due to the large number of registers inherent to this type of design, while partial scan or testing from the boundary of a filter usually means long test sequences due to the deep sequential nature of the datapath. In this case, built-in self-test based on pseudorandom pattern generation is an attractive approach since it offers inexpensive testing and usually provides high fault coverage, although it is rarely able to achieve 100% coverage.

In practice, it is not difficult to reach fault coverages in the high 90's using pseudorandom BIST on filter datapaths. However, it has been shown that these tests can be seriously flawed in the sense that some of the faults that are left untested can be relatively easy to trigger during normal operation of the device [1]. Such flaws can often be traced to some fundamental problem in the test generator, such as correlation properties that are incompatible with the device under test [1, 2]. In many cases, these deficiencies can be overcome to some extent with the use of decorrelating cir-

cuits or alternate test modes [1].

Even with these improvements to random pattern testing, the fundamental quality of a test will remain uncertain until the risk associated with any untested faults has been understood. Fortunately, linear digital systems offer an unusual opportunity to directly assess the risk of a fault—specifically, to bound the maximum fault activation rate expected under normal operating conditions. This information can be used to profile a test, providing information about which faults pose the greatest risk, and which are so unlikely to be activated that they are deemed “pathological”, or near-redundant.

Practical implementations of large fixed-point filters offer an interesting comparison point for test quality: it is usually possible to induce overflow errors during normal operation of a filter; such errors can be viewed as “designed-in” faults. Clearly, a fault with a maximum activation rate far below the activation rate of such intrinsic faults should not be a test priority. On the other hand, a large number of faults with potentially high activation rates can be a sign of a seriously flawed test, and may signal the need to design a new test, or to target the associated logic for test point insertion.

Test quality is an important factor in choosing a test strategy. For example, it is tempting to believe that a long white-noise test sequence will catch all faults that are likely to be activated during normal operation of the datapath. However, this is not always the case. For example, wide-band test signals can miss faults due to the *narrow-band test problem*: even in wide-band filters it is common to find portions of the circuit that exhibit a narrow-band response, leading to reduced test signal power when wide-band tests are used. High-variance wide-band tests may still detect the fault, but will typically require very long test sequences. In contrast, such a fault might be easily triggered during normal operation of the datapath if a narrow-band input signal is applied. Similarly, non-white LFSR-generated tests can provide high fault coverage, yet still miss significant faults due to low signal power in the filter's passband. An objective means of determining functional test quality is needed.

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

In order to place bounds on the activation rates of untested faults, we will need to model the normal operating behavior of the datapath. We will see how fault activation rates depend on internal signal variances and correlation, and ultimately how this can be linked back to the input signal power spectrum. The bounds on untested fault activation rates will be presented using *risk profiles*, which give insight into the quality of a test.

2 Design Faults

Overflow can be viewed as a designed-in fault; in this section we will look at the probability of activating this type of fault. This provides an introduction to the fault activation rate computations of the sections which follow.

We will model the input signal as a normally distributed, zero-mean random variable. In this case, all internal signals and the datapath output will also be normally distributed, by the Gaussian reproductive property (the sum of normally distributed random variables is again a normally distributed random variable). Signal variance is then sufficient to completely characterize the distribution of all derived signals, although it must be kept in mind that these derived signals may be correlated.

Let the idealized filter's output (in the absence of overflow) be modeled as a normally distributed, zero-mean random variable with variance σ_y^2 . The probability of overflowing the $[-L, L]$ interval is then

$$P_{\text{ovf}} = \text{erfc} \left(\frac{L}{\sqrt{2} \sigma_y} \right), \quad (1)$$

where erfc is the complementary error function, given by $\text{erfc}(z) = 1 - \text{erf}(z)$. The error function, $\text{erf}(z)$, is defined by

$$\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt.$$

L represents the absolute maximum value that the datapath output can support without overflow. We will usually interpret signals as two's-complement numbers on the interval $[-1, 1)$, corresponding to $L = 1$. The relation between P_{ovf} and σ_y for $L = 1$ is plotted in Figure 1.

In FIR filters, the designer's concern is to keep the overflow probability at the output within acceptable levels: internal overflow does not result in an output error as long as the internal width is at least as wide as the output width, and the final result using infinite-precision arithmetic would fall within the output range. In IIR filters, overflow should also be minimized at feedback nodes. Peak values of σ/L in the range of one-fourth to one-third are common.

3 Input Spectrum and Power Gain Factor

To compute fault activation probabilities, it will be necessary to compute the maximum signal variances under

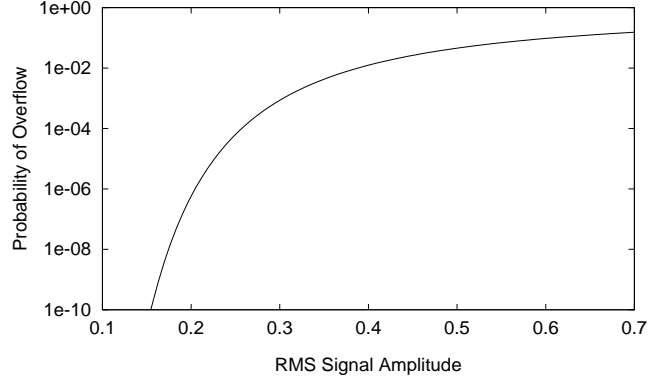


Figure 1. Probability of overflow vs. RMS signal amplitude.

worst-case operating conditions at each point in the datapath, where worst-case conditions are identified using information about the input signal's spectrum and maximum variance. The *power gain factor*, g , is defined as the ratio of output variance to input variance of a filter or subfilter, $g = \sigma_y^2 / \sigma_x^2$. Elements of scaling theory [3] can be used to find the maximum output variance for a stochastic model of the datapath's input signal. Maximizing variance requires concentrating all input signal power at the subfilter's peak gain frequencies. For a filter with frequency response $H(e^{j\omega})$, the maximum power gain is

$$g_{\text{max}} = \max_{\omega} |H(e^{j\omega})|^2.$$

As discussed in Section 2, systems are frequently designed with peak input/output σ values between 0.20 and 0.33, assuming a ± 1 signal range. Computing g_{max} for each signal in the datapath, we can determine the worst-case signal variances for a given peak input power, σ_x^2 . However, in some applications g_{max} will be too conservative. For example, if signal power in some bands is limited, or if it is not possible to concentrate all power at the subfilter's peak frequency response, g_{max} may overestimate output variance. In such cases the approach can be refined to account for the characteristics of the operational spectrum. Without knowing such specifics, we use g_{max} . Additional discussion of power gain factors can be found in a separate report [4].

4 The Bivariate Normal Model

In order to determine the worst-case fault activation rates in a datapath, we will need to characterize the behavior of each adder and subtracter. Neglecting truncation effects (which tend to diminish at the upper bits that are our chief concern), fault activation rates can be determined from the input variances (σ_A^2 , σ_B^2) and correlation coefficient (r). In this section, we will model the inputs to an adder as a bivariate normal distribution and examine how test activation rates are influenced by the parameters σ_A^2 , σ_B^2 , and r , which will in turn give an indication of the type of input spectrum that maximizes fault activation rates.

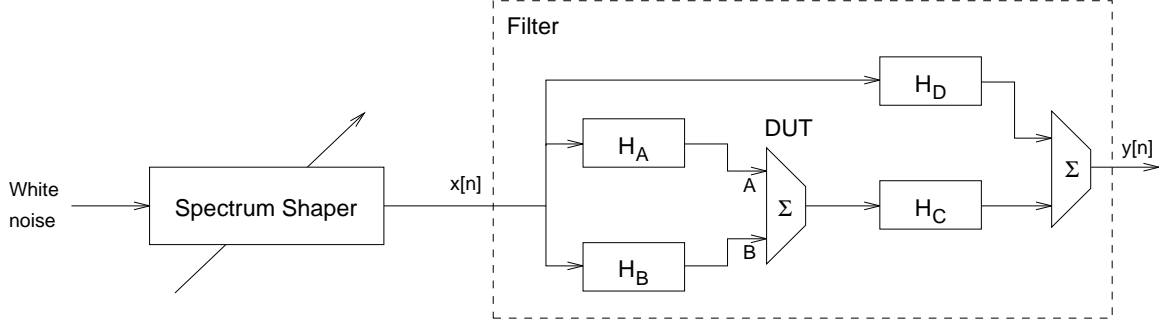


Figure 2. Block diagram used to examine effect of input spectrum on fault activation probabilities during normal filter operation.

Figure 2 provides the context for our discussion. The input to the datapath, denoted $x[n]$, has a spectrum modeled by a white noise source feeding a spectrum shaper. The adder being examined is fed by two paths; the A input is fed by a subfilter with response H_A , and the B input is fed by a subfilter with response H_B . When one input has a predominately larger input variance than the other, we use the convention that the A input is the high-variance “primary” input, and the B input is the lower-variance “secondary” input. Where there are testability problems, it is common to find $\sigma_B \ll \sigma_A$, although this relationship can vary with input spectrum, and even reverse at certain input frequencies, depending on H_A and H_B . The difference between input variances is referred to as the *variance gap* [2, 5].

The response of the datapath output to a signal injected at the adder’s output is given by H_C ; this can be used to model the transformation of fault effects as they travel from the adder’s output to the observation point $y[n]$. The remainder of the datapath functionality is modeled by H_D , where the overall filter response, H , is given by $H_D + H_C (H_A + H_B)$.

Fault activation during normal operation of the adder can be understood by modeling the adder’s inputs with a bivariate normal distribution with zero mean,

$$p(\mathbf{x}) = \frac{1}{2\pi|\Sigma|^{1/2}} \exp\left[-\frac{1}{2}\mathbf{x}'\Sigma^{-1}\mathbf{x}\right], \quad (2)$$

where $\mathbf{x} = [A \ B]'$, \mathbf{x}' denotes the matrix transpose of \mathbf{x} , and the 2-by-2 covariance matrix, Σ , is given by

$$\Sigma = \begin{bmatrix} \sigma_A^2 & \sigma_{AB} \\ \sigma_{AB} & \sigma_B^2 \end{bmatrix}.$$

The cross-correlation, σ_{AB} , is $\mathcal{E}\{AB\}$. The correlation coefficient, r , is defined as $r = \sigma_{AB}/\sigma_A\sigma_B$, $-1 \leq r \leq 1$.

In variance mismatched adders, we will find that triggering some difficult faults will require maximizing variances and input correlation, while other difficult faults will be less sensitive to total signal variance, but are dependent on the ratio of input variances and prefer uncorrelated inputs. Consequently, we need to know the range of values that σ_A^2 , σ_B^2 , and r can take.

The maximum values that σ_A^2 and σ_B^2 can take are determined by the maximum power gain factors of the H_A and H_B responses, respectively. As discussed in Section 3, this depends on what type of input spectra are possible during normal operation; if narrow-band signals at the subfilter’s peak response frequency are possible, the g_{\max} power gain factor should be used.

The correlation coefficient, r , can typically cover most of the $[-1, 1]$ range as input frequency is swept. In determining fault risk, we make no assumption restricting r ’s range.

5 Test Zones

In earlier work, the input conditions necessary for asserting “difficult” tests at the upper bits of an adder were derived [2]; these are repeated here in Table 1. These tests can be mapped to the corresponding numbered regions in Figure 3(a). Subtractor test zones are shown in Figure 3(b). These patterns, corresponding to tests of the next-to-MSB logic, are periodically extended to account for out-of-range values. For lower bits, the test zone size is scaled by 2^{-b+1} , where b is the distance from the MSB. The logic faults triggered by a signal can be determined by looking at the test zone that the adder’s inputs fall in.

To find how correlation and signal power relate to fault activation rates during normal operation of the design, we will combine these test activation regions with the bivariate normal model of the adder’s input distribution to compute test activation rates. This in turn maps directly to fault activation rates based on the logic fault model used.

Table 1. The four difficult test zones of the next-to-MSB adder.

| Test | Input | Output |
|------|-------------------|-------------------|
| T1p | $0 \leq A < 0.5$ | $A + B \geq 0.5$ |
| T1n | $A < -0.5$ | $A + B \geq -0.5$ |
| T2p | $0 \leq A < 0.5$ | $A + B < 0$ |
| T2n | $A < -0.5$ | $A + B \geq 0.5$ |
| T5p | $A \geq 0.5$ | $A + B < -0.5$ |
| T5n | $-0.5 \leq A < 0$ | $A + B \geq 0$ |
| T6p | $A \geq 0.5$ | $A + B < 0.5$ |
| T6n | $-0.5 \leq A < 0$ | $A + B < -0.5$ |

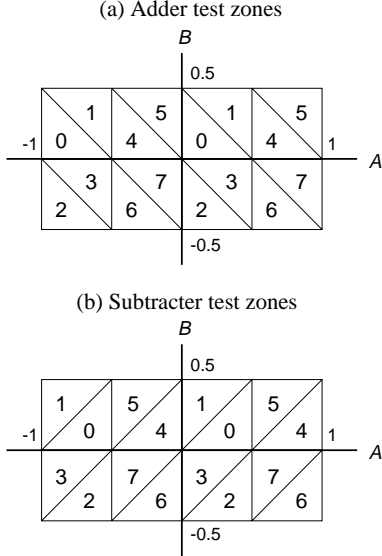


Figure 3. Adder test activation zones (a), and subtracter test activation zones (b) for the next-to-MSB bit. The pattern is periodically extended in all directions, and the size of the zones is halved for each progressively lower bit.

6 Fault Activation Rates

For a fault requiring a test T with test zones Z_{jk} , where j and k specify the horizontal and vertical indices of the test zones, the fault activation probability is

$$P_{\text{act}}(T) = \sum_{j,k} \iint_{Z_{jk}} p(\mathbf{x}) dA dB, \quad (3)$$

where $p(\mathbf{x})$ is the bivariate normal density defined in Equation 2. The activation probability is a function of σ_A , σ_B , and r . Figure 4 shows the test zones corresponding to the MSB-1 overflow tests T1 and T6, with the constant-density ellipses of a positively correlated ($r > 0$) normal distribution superimposed. Equation 3 is computed over the shaded regions.

While the A datapath width may only support values of A over the range $[-1, 1)$ (as is typically assumed for fixed-point two's-complement notation), we extend the regions beyond this limit. The reason for this is that two's-complement arithmetic permits overflow at internal nodes as long as the final (ideal) result is in-range. Consequently, we can treat the A input signal as though it covers a wider range, and extend the test zones to account for the portions that would be wrapped by overflow.

To demonstrate the effects of correlation on test activation rates, we examine a bivariate normal distribution with $\sigma_A = 0.2$, $\sigma_B = 0.01$. In this example, the test zones shown in Figure 3 are sufficient; the more distant test zones contribute little to the P_{act} computation. The integral in Equation 3 was computed for values of r ranging from -0.95 to 0.95 . The results are shown in Figures 5-8.




Figure 4. Computing activation probabilities for test T1, positively correlated A and B . The shaded areas indicate the test zones over which the integral in Equation 3 is computed. The ellipses indicate constant-density levels of a positively-correlated normal distribution.

7 Fault Characterization

We now characterize the basic fault classes based on the type of tests required to activate the fault in question. Faults are broadly classified as either easy or difficult faults. Difficult faults are further broken into central, overflow, and joint central-overflow faults. This characterization applies primarily to high-order datapath bits, where testability is most often a concern. At lower bits, test zones proliferate, leading to a more uniform distribution of tests. Thus, there is little to distinguish fault classes at the low-order bits.

Easy Faults: The so-called “easy faults” are faults that can be tested by any one of the four tests that are not sensitive to variance gap (T0, T3, T4, T7). For uniformly-distributed independent inputs to the adder, the probability of asserting any of these tests is near 0.125 in variance-matched adders, and rises to 0.25 as the variance gap is increased. Under certain conditions these tests may not be quite so easy to assert. For example, if the correlation between adder inputs approaches unity, T0 and T7 activation rates increase at the expense of T3 and T4. The tests are considered easy because they have high activation rates over a broad range of correlation values. With correlation, easy fault activation rates can reach 50%. Figures 5 and 8 illustrate the typical behavior of this class of faults as a function of r .

Difficult Faults: Faults that are not testable by the easy tests are classified as difficult faults. These are faults that are only testable by one or more of the tests T1, T2, T5, and T6. Difficult faults fall into three possible categories: central, overflow, and joint central-overflow faults. Of these, overflow faults tend to be the most difficult to test, while central faults tend to be relatively testable.

Central Faults: Difficult faults that are only testable by the central tests are termed central faults. At the next-to-




Figure 5. Test T0 and T7 activation probability as a function of correlation coefficient, $\sigma_A = 0.2$, $\sigma_B = 0.01$.




Figure 6. Test T1 and T6 activation probability as a function of correlation coefficient, $\sigma_A = 0.2$, $\sigma_B = 0.01$.




Figure 7. Test T2 and T5 activation probability as a function of correlation coefficient, $\sigma_A = 0.2$, $\sigma_B = 0.01$.




Figure 8. Test T3 and T4 activation probability as a function of correlation coefficient, $\sigma_A = 0.2$, $\sigma_B = 0.01$.

MSB bit, central tests are asserted when the primary input of the adder under test and its output have opposing signs. In adders, these are tests T2 and T5, while in subtractors the central tests are T1 and T6. The activation rates of central faults are typically maximized by uncorrelated input signals (zero correlation coefficient). In many filter architectures (e.g., the common transpose-form filter) this can be achieved by using a test generator that exhibits low sample-to-sample correlation.

Under some circumstances, central faults can reach high activation rates. This happens when the normal $\sigma_A > \sigma_B$ relation is inverted, as can happen when the test signal lies in the stopband of the subfilter feeding the A input, but is in B 's passband. In this case, activation rates comparable to those of the easy faults are possible. This provides some motivation for the test approach of applying narrow-band test signals that sweep the full input spectrum. However, we have not seen a strong need for this level of test generator complexity for testing central faults, given that wideband, uncorrelated test signals tend to give good activation rates for central faults in most architectures, and are cheaper in terms of chip area.

Maximizing central fault activation rates does not generally require the test signal's variance to be maximized. Low variance test signals can be just as effective as long as the adder inputs do not degenerate due to truncation effects. This property makes central faults particularly undesirable from a signal-to-noise perspective, as it is possible that fault effects will appear even when the input signal is at low power levels. Figure 7 gives an example of central fault activation dependence on correlation.

Overflow Faults: Overflow faults are those faults that are tested only by asserting an overflow condition at the adder slice in question: tests T1 and T6 for an adder or tests T2 and T5 for a subtractor. Note that this does not refer to overflow of the entire addition or subtraction operation, but only to overflow at a single bit ($carry_{in} \neq carry_{out}$ at a full-adder cell). Overflow faults at high-order bits form the class of faults whose behavior most closely resembles that of the normal, functional overflow faults. They tend to require high-variance input signals to activate, and their activation rates rapidly decrease as input power is reduced.

Since overflow faults require high signal variance, testing of these faults is usually most effective when the test signal has a substantial portion of its power in the passband of the subfilter outputting at the adder in question. Figure 6 illustrates overflow fault dependence on correlation.

Central-Overflow Faults: Difficult faults that are tested by either central or overflow tests are classified as *central-overflow faults*. In the commonly used adder models employed in this study, all difficult logic faults are either central or overflow faults; there are no central-overflow faults. Such faults do not pose any additional theoretical challenge since their activation rates are dominated by the activation rates of either the central or overflow tests.

8 Fault Activation Bounds

In gauging the quality of a test, we are interested in computing bounds on the fault activation rates of any untested faults. Of the fault classes described in Section 7, the most stubborn faults are typically overflow faults, followed by

central faults where there is a large variance gap. Any easy faults that go untested are penalized by assigning a high fault activation rate (e.g., 0.5, since under the right conditions this is possible). For any untested fault, the maximum fault rate attainable under worst-case operating conditions is a measure of the risk associated with the fault.

While, in general, computing fault activation rates requires evaluating double integrals over a wide range of power and correlation combinations, it is possible to efficiently compute bounds using the observations made earlier about the σ and r values that maximize fault rates. Based on the test zone model, we have seen that positive correlation maximizes overflow fault activation rates, while zero correlation maximizes central fault activation rates (although there is an exception to this, which will be discussed in Section 8.2). For central faults, low input signal variance can give good fault activation rates since the test zones are adjacent to the origin, while high signal variance is required to activate the more distant overflow fault regions.

These observations about fault behavior are based on the assumption that the test zone tiling period in Figure 3 is larger than the range of B ; if the tile size is on the order of a couple of σ_B or smaller, the fault characterization of Section 7 tends to break down as more test zones become involved in the P_{act} computation. In the extreme, as σ_A and σ_B are increased (or equivalently, the tile size is scaled down), the tests start to look more uniformly distributed. This is not a serious problem for computing fault activation bounds, as we are generally only interested in bounds on the most difficult faults, which lie at the upper bits of variance-mismatched adders. Untested faults that lie too far down from the MSB are assigned a high fault risk, just as for the easy faults at upper bits.

8.1 Overflow Faults

For overflow faults, we develop a bound on fault activation probabilities based on the assumption that the inputs are linearly dependent ($r = \pm 1$). This bound is conservative, as it requires simultaneously maximizing σ_A , σ_B , and r .

We will consider adders here; subtracters are handled analogously using the test zones of Figure 3(b). For overflow test T1, the first quadrant test zone gives the largest activation rates since the second quadrant test zone is slightly farther from the mean. This produces the slight asymmetry observed in Figure 6. Therefore, we assume positive correlation in this bound. (Similarly, we would choose positive correlation to maximize overflow test T6 since the third quadrant zone is closer to the mean than the fourth quadrant zone.) For subtracters, negative correlation is preferred.

The bound is based on the assumption that A and B are fully correlated, $B = \alpha A$, where $\alpha = \sigma_B/\sigma_A$. In this case, the 2-D integral of Equation 3 reduces to a line integral. Taking A as the independent variable, the T1p test constraints ($A < u_i$ and $A + B \geq u_i$) correspond to the




Figure 9. Overflow fault probability vs. σ_A , $\alpha = 0.05$.

integration ranges given by

$$\frac{u_i}{1 + \alpha} \leq A < u_i, \quad i = 0, 1, 2, \dots,$$

where $u_i = (1 + 2i) 2^{-b}$ for adder bit MSB- b . The width of the test zones is also limited by B_{max} , the maximum value that B can take. Accordingly, the fault activation bound is given by

$$P_{\text{act}} = \sum_{i=0}^{\infty} 0.5 \left[\text{erf} \left(\frac{u_i}{\sqrt{2}\sigma_A} \right) - \text{erf} \left(\frac{l_i}{\sqrt{2}\sigma_A} \right) \right]. \quad (4)$$

where the l_i are the lower edges of the test zones, $l_i = \max(u_i - B_{\text{max}}, u_i/(1 + \alpha))$.

Even though the outer test zones are more distant from the mean than the inner zones, negative correlation can conceivably result in a larger fault activation rate since the outer zones are wider than the inner zones (e.g., test T1n could have a higher activation rate than T1p). Therefore, when computing the bound we try both $r = 1$ and $r = -1$, taking the larger of the two P_{act} values as the bound.

The bound can be tightened in some cases by using the output standard deviation, σ_{A+B} , in place of the input standard deviation σ_A , with appropriate modifications to the above equations. Thus, we compute the bound using both approaches and select the tightest bound. The behavior of the bound as a function of input power is shown in Figure 9 for $\alpha = B_{\text{max}} = 0.05$. Applying the bound to the example in Section 8 ($\sigma_A = 0.2$, $\sigma_B = 0.01$), the bound gives $P_{\text{act}} = 0.0024$ for bit MSB-1, in close agreement with the right limit in Figure 6.

8.2 Central Faults

For central faults, fault activation rates tend to be maximized by uncorrelated signals ($r = 0$), since any correlation tends to pull more of the distribution into some of the easy test zones (T0 and T7 for adders with positively correlated inputs, or T3 and T4 for negatively correlated inputs). An exception occurs when $\sigma_B/\sigma_A > 1$, in which case the central test becomes an “easy” test, which is maximized by $r = -1$ for adders or $r = 1$ for subtracters. For fully correlated signals, the fault activation rates can reach as high as 50%. In variance mismatched adders where testing is

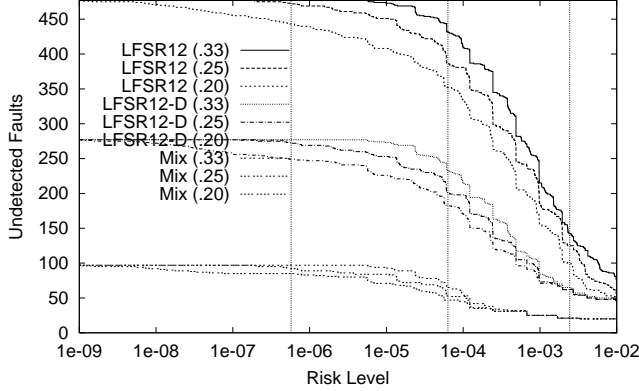


Figure 10. Fault risk profiles for three test sequences, 60-tap lowpass filter. Risk profiles are shown for maximum input standard deviations of 0.20, 0.25, and 0.33. The three test sequences were generated by a 12-bit LFSR (length 4k), a 12-bit decorrelated LFSR (length 4k), and a mix of a 12-bit LFSR (length 2k) followed by an 8k maximum-variance LFSR sequence (total length 10k). The vertical bars mark the probability of the fault-free filter generating an overflow effect for input standard deviations of (from left to right) 0.20, 0.25, and 0.33.

a problem, it is typical to find $\sigma_B \ll \sigma_A$. Even so, the easy/central inversion can still occur at frequencies that are blocked by H_A but passed by H_B .

For $r = 0$, we assume that σ_A and σ_B are small enough that the bulk of the distribution falls within range of the central test zones that lie at the origin. This gives the upper bound on central fault activation rates

$$P_{\text{act}} = \frac{1}{2\pi} \arctan\left(\frac{\sigma_B(\omega)}{\sigma_A(\omega)}\right), \quad (5)$$

where the ratio $\sigma_B(\omega)/\sigma_A(\omega)$ is maximized over the interval $[0, 2\pi]$. This result follows from a change of variables, $B' = \sigma_A B$, $A' = \sigma_B A$, which circularizes the distribution. The maximum P_{act} is attained by a narrow-band signal at the frequency where the ratio $|H_B(e^{j\omega})|/|H_A(e^{j\omega})|$ is greatest. If this ratio is greater than 1 for some value of ω , inversion can take place, and we accordingly penalize any associated faults by using $P_{\text{act}} = 0.5$. Continuing the example from Section 8, this bound gives $P_{\text{act}} = 0.008$, in agreement with the peak value from Figure 7.

9 Risk Profiles and Results

To examine the risk associated with untested faults in a design, fault activation rate bounds were computed for a number of different test sequences applied to a 60-tap lowpass filter. The results are presented using *risk profiles*, where the total number of faults with activation rates greater than a given risk level is plotted as a function of risk level. Since fault activation bounds can span a wide range, risk level is plotted on a log scale.

Figure 10 shows risk profiles for three test sequences: a 4k-long LFSR sequence (LFSR12), a 4k decorrelated LFSR (LFSR12-D), and a higher-quality mixed test consisting of




Figure 11. Fault risk profiles for the 60-tap lowpass filter, for two pairs of identical length sequences. Profiles for two 10k-long test sequences (from a 16-bit LFSR and the mixed sequence of Figure 10), and two 4k-length test sequences (from a 12-bit LFSR and a decorrelated 12-bit LFSR) are shown. The 4k decorrelated LFSR is generally superior to the much longer 10k LFSR sequence, although it lets a few more high activation rate faults escape, as indicated by the intersection of the two curves.

a 2k decorrelated LFSR sequence followed by an 8k maximum variance test [1]. Risk profiles are shown for three different maximum input signal levels (σ_x): 0.20 (representing a very conservative design), 0.25, and 0.33. The plot legend lists the value of σ_x assumed for each curve. It can be seen that the fault activation bounds for some faults rapidly decrease as the maximum input signal power is reduced.

For comparison, the normal operational fault rates given by Equation 1 at peak input power span the range from 5.7×10^{-7} for $\sigma_x = 0.20$ to 2.4×10^{-3} for $\sigma_x = 0.33$. For $\sigma_x = 0.25$, the fault rate is 6.3×10^{-5} . (For reference, these risk levels are marked with vertical lines in Figures 10–12.) For conservative design (e.g., $\sigma_x < 0.20$), many of the untested faults pose substantial risk compared to the normal operational faults, while for more aggressive design only a few of the untested faults are of greater risk than these intrinsic overflow faults. This ranking can be used by the test designer to focus test resources on the highest risk faults.

In Figure 10, it is interesting to note that the decorrelator is effective in reducing the number of untested faults with worst-case activation ranges between 10^{-5} and 10^{-2} , but does not significantly enhance detection of the highest risk missed faults, which is on the order of 50–80 faults for both the LFSR12 and LFSR12-D tests. This seems to be due to the fact that the decorrelator, by increasing test signal passband power, is better at reaching overflow faults. However, this effect does not help in detecting the toughest of the high-risk central faults. In fact, there are 49 undetected central faults after LFSR12 testing, but 73 after LFSR12-D (decorrelated) testing. The longer Mix test combining a decorrelated LFSR with maximum variance testing does a good job at proportionally reducing risk at all levels as compared to the LFSR12-D test. After applying this test, there

were 31 untested central faults.

Figure 11 compares two sets of similar length tests of differing quality for $\sigma_x = 0.25$. The 4k-long LFSR12 and LFSR12-D tests, and the 10k-long Mix test from Figure 10 are repeated here, and a 10k-long sequence from a 16-bit LFSR (LFSR16-10k) is added.

At a test length of 4k vectors, the decorrelated LFSR test (LFSR12-D) shows clear superiority to a standard LFSR (LFSR12). Although the number of very high risk undetected faults is similar in both cases, the decorrelated test cuts the number of undetected faults by roughly 2:1 at low-to-moderate risk levels. At a test length of 10k vectors, the Mix test provided an even greater advantage over the standard LFSR (LFSR16-10k), showing more than a 4:1 reduction in undetected faults at most risk levels.

Figure 11 also shows that increasing test length does not always offer much risk reduction if the test strategy is flawed, as can be seen by comparing the 4k LFSR12 test with a 10k LFSR16-10k test. Although more than twice as many vectors are applied by LFSR16-10k, little improvement is seen in the risk profile. Switching to the Mix approach offers much lower risk for the same 10k test length. In fact, the 4k decorrelated LFSR12-D test is superior in most respects to the 10k standard LFSR, although it lets a few more very high risk faults escape.

Figure 12 shows an example of a more dramatically flawed test, the 4k-long RAMP test (generated by a count-by-one circuit). This test leaves over 200 very high risk faults untested, compared to about 50 for both 4k-long LFSR-based tests. Interestingly, the standard LFSR provides almost exactly the same absolute fault coverage as the RAMP test, yet is clearly superior from a fault risk perspective. The number of undetected faults for the RAMP test was 480, as compared to 477 for the LFSR12 test. The total number of modeled faults was 57,124, giving a fault coverage of 99.2% for both tests. The design size was on the order of 25k gates. The RAMP test left many central faults untested (232), possibly a serious deficiency since this type of fault can generate large fault effects even when the filter is relatively quiescent.

10 Conclusion

In high-performance datapaths, there is often limited opportunity to insert test structures or to restructure logic for improved testability. Fortunately, many faults in DSP ASIC datapaths are highly testable, allowing efficient testing from the boundaries of the datapath. However, not all faults represent equal risk to the fault-free operation of a digital filter. Some faults have a high likelihood of being triggered, while others are almost impossible to activate with any operational input signal. Consequently, fault coverage by itself is not always a sufficient indicator of test quality. For example, a random pattern generator used in a BIST scheme might give

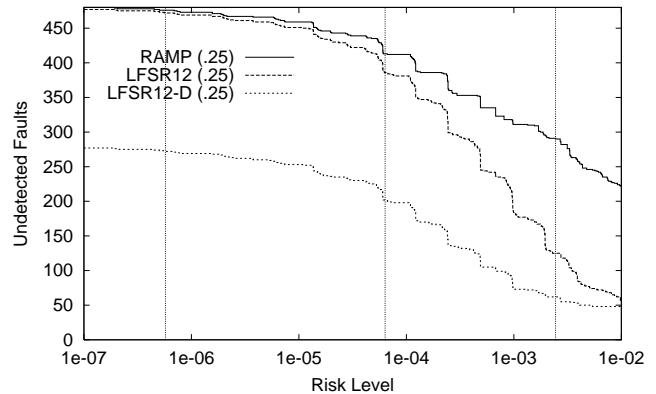


Figure 12. Fault coverage can be a poor indicator of test quality. Two identical length sequences (a Ramp and an LFSR-generated sequence) give the same fault coverage, but the Ramp sequence leaves over 200 faults with high activation probabilities untested, compared with roughly 60 for the LFSR sequence. The decorrelated LFSR (LFSR-D), while not significantly reducing the number of very-high risk faults below 60, does reduce the number of medium-risk faults (in the 10^{-2} to 10^{-4} range) as compared to the standard LFSR. Risk profiles are for $\sigma_x = 0.25$.

very high fault coverage, yet leave significant faults untested due to some spectral characteristic of the input signal that is lacking in the test signal. The seriousness of these untested faults is gauged by their likelihood of generating a fault effect during normal operation of the filter.

To analyze the quality of a test from a functional perspective, we introduced risk profiles based on worst-case operational fault activation rates. The risk associated with any untested fault is estimated from activation rate bounds using a model of the filter's input signal. The more information that is available about the input signal's spectrum and power, the more tightly we can bound the risk associated with any untested faults. However, even with fairly limited knowledge of the input signal, we can detect serious flaws in a test sequence. By identifying faults with potentially high activation rates, better tests can be developed, or test point insertion can selectively target the most critical faults.

References

- [1] L. Goodby and A. Orailoğlu, "Frequency domain compatibility in digital filter BIST," in *34th Design Automation Conference*, pp. 540–545, June 1997.
- [2] L. Goodby and A. Orailoğlu, "Pseudorandom-pattern test resistance in high-performance DSP datapaths," in *33rd Design Automation Conference*, pp. 813–818, June 1996.
- [3] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Prentice Hall, 1989.
- [4] L. Goodby, *Test Synthesis and Self-Test in High-Performance VLSI Digital Signal Processing*. PhD thesis, University of California at San Diego, 1997.
- [5] 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.