

ICID - Tracing Individual Die from Wafer Test through End-Of-Life

Keith Lofstrom^[1], David Castaneda^[2], Brian Graff^[2], Anthony Cabbibo^[2]

^[1] SiidTech, Beaverton, Oregon <http://www.siidtech.com>

^[2] LSI Logic, Gresham, Oregon <http://www.lsil.com>

Abstract

ICID (Integrated Circuit IDentification) is a small mixed-signal cell that can be added to the test logic on a CMOS integrated circuit. It provides a unique 224 bit identification number that can be accessed during die test. This identification can be used to correlate test information for individual die on the wafer, through package test, and into the field and back. The identification bits are produced from the fixed analog mismatch of an array of PFET pairs, and does not require process modifications or programming. LSI Logic is using ICID technology to trace individual die through test, and correlate test statistics from wafer test, package test, and failure analysis.

Introduction

No test is perfect. Incomplete tests can pass defects that cause failures in the final application, while excessively strict tests can fail parts that would be perfectly adequate in use. The cost of field failures is high, so most manufacturers err on the side of caution and test for extremes, discarding integrated circuit die that a customer would find acceptable. The result is lower yield and excessive test time, resulting in more expensive products.

The ability to trace die from manufacturing through end-of-life can reduce this expense [1-3]. A die returning from the field with an error can be retested to determine the cause of failure. But how does this re-test correlate to the original wafer and package test? Without data from the first tests, it is difficult to know how behavior may have changed while the part was in the field. This is especially troublesome for mixed-signal parts, and parts that are binned in multiple grades.

With ICID (Integrated Circuit IDentification) technology [4], manufacturers can now identify individual die at wafer test, package test, and during failure analysis, allowing data from all three tests to be correlated. Unlike other identification technologies [5], ICID uses unmodified digital CMOS processes and requires no special process modification, characterization, or programming steps.

How ICID Works

ICID is based on the matching behavior of minimum-sized FET device pairs. Even when a pair of field effect transistors are designed to be very well matched, there is inevitably some mismatch in threshold voltages due to the random nature of the ion implant process. Because the number and placement of channel dopant atoms is statistical, with variation around an average, the voltage necessary to turn the channel on and off will vary from transistor to transistor, in spite of the best efforts of process engineers to make the channel dose identical. The gate oxide capacitance of a minimum-sized FET in a 130 nm process is around 0.3 femtofarads; adding or subtracting a single dopant atom from the channel will move the voltage threshold by hundreds of microvolts. The threshold variance of the whole channel will be the statistical sum of thousands of dopant atoms, resulting in a Gaussian distribution of threshold voltages, with a one-sigma threshold mismatch of around 50 millivolts. This mismatch increases as processes are scaled down.

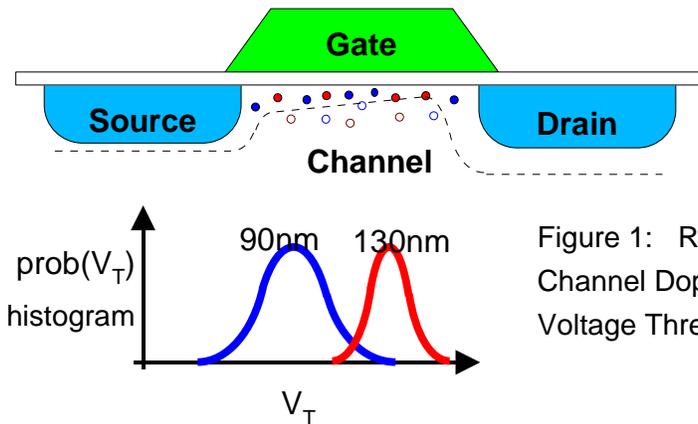


Figure 1: Random Variation of Channel Dopant Atoms Lead to Voltage Threshold Variaton

A pair of 0.13 μm FETs, properly biased, will exhibit the statistical sum of two mismatches, or about 70mV, which will appear as an input offset to the pair. If the pair is included in a differential amplifier stage, and the FET gate inputs are connected together, the output of the amplifier will be the mismatch voltage times the amplifier gain.

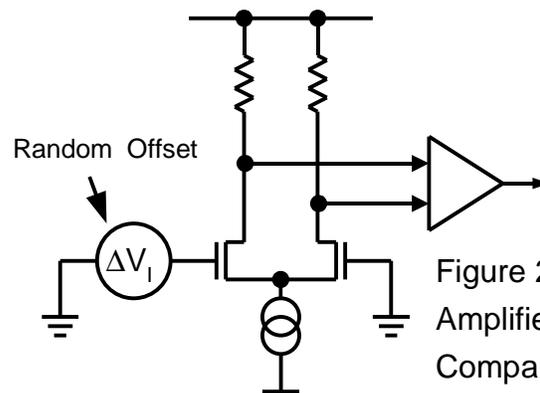


Figure 2: Differential Pair Amplifies Offset into Comparator

The bell curve of voltage mismatch can be detected with a low-offset comparator. For a properly designed circuit, the result will be a binary bit value with a 50% chance of zero and a 50% chance of one for each pair measured. The value of the bit produced by each pair is unpredictable in advance, since the dopant count in the FET channels is unpredictable. However, the dopant atoms are fixed over time; the binary bit value is also unchanging.

Well, almost; sometimes, the voltage mismatch will be near the center of the bell curve. The voltage threshold of a FET changes with temperature, current, and age. The comparator threshold and load mismatch will also vary. Thermal noise and power supply ripple will shift the comparator threshold by small amounts. Thus, measurements of bits near the center of the threshold will be noisy. The bits produced by this process will also have some noise and variance. We call this variance "bit drift", and it typically affects 1% to 5% of the bits produced by an ICID circuit.

The effects of this bit drift error can be reduced arbitrarily by using a lot of bits; while a few bits can drift, most won't, and a large number of bits in an ID are very unlikely to match a different ID.

An ICID circuit uses a two-dimensional array of 224 FET pairs, with switching circuitry to sequentially select one of the pairs to feed to an autozeroing comparator. The result is a 224 bit series of partially correlated ones and zeros. We read these IDs as a serial bit sequence, clocked out by a test clock. The measured bit sequence is almost entirely different between two chips.

When the same ID cell is measured repeatedly, there will be some changes in the sequence due to bit drift. Thermal noise effects can be averaged out by taking multiple measurements. Other errors are due to power supply variations, aging, and the changes in die stress associated with dicing and packaging. The largest component of bit drift comes from temperature changes, as much as one percent per 20C. All sources of bit drift add together in an RMS sum. If test temperature is constant, typically fewer than 5 bits change per 224 bit ID.

How do we deal with this bit change? When we want to find an ID in a large database, we cannot just look for an exact match. Instead, we look for the ID in the database with the smallest number of changed bits - the lowest "Hamming distance". Typically, any pair of 224 bit IDs will have an average of 112 bits that are different, while the same ID measured twice will have an average of 5 bits that are different. The actual number of bits for "same" and "different" will both follow bell curves that diminish rapidly for large variations from average.

Figure 3: Probability of Self Match and Others Match

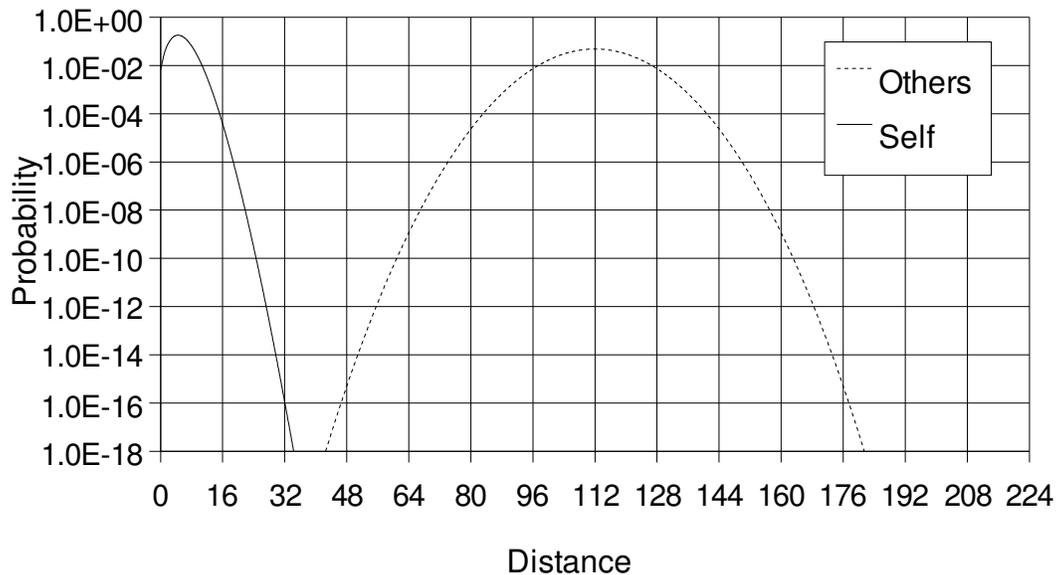
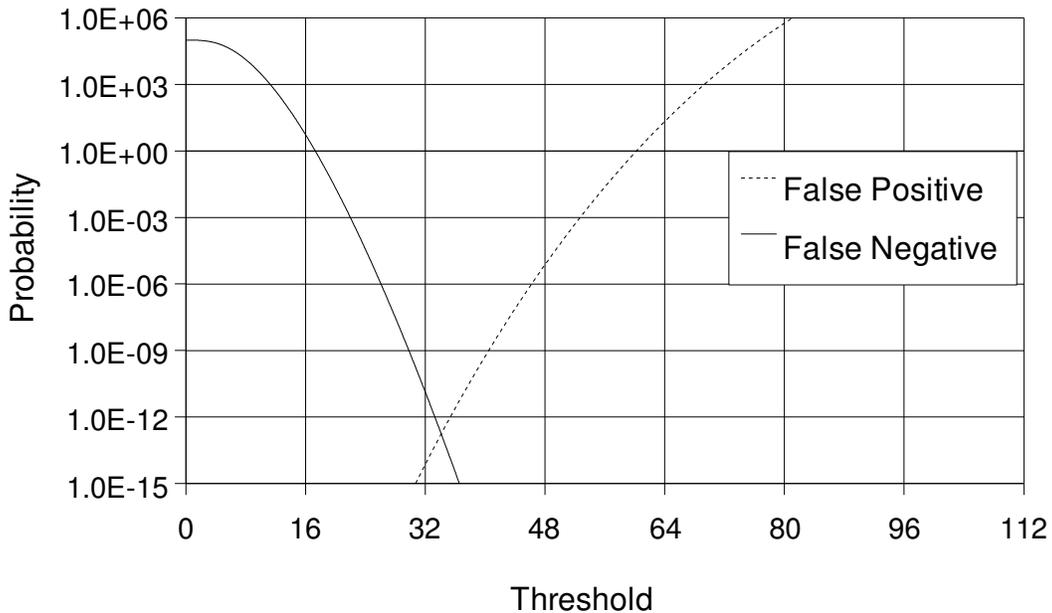


Figure 3 shows the "self" and "others" curve given an average 5 bits of drift in a 224 bit ID. The **self** curve measures the probability of a single ID cell showing a given amount of drift. The **others** curve shows the probability of a given distance to another ID. The self curve, multiplied by the number of die in a lot, computes the number of die per lot that will drift a given amount. The others curve, multiplied by the number of die **squared** (the total number of comparisons) computes the number of "false positive matches" that can be expected at a given distance.

These multiplied curves can be integrated to produce curves of the number of false positives and false negatives that could be expected above and below a given threshold. Based on these curves, we can pick a threshold number for the Hamming distance, say 30, and assume that any two measurements that are less than 30 bits different indicate a match, while two measurements that are greater than 30 bits indicate a mismatch.

Figure 4 shows the number of false positives and false negatives per lot that result as a function of the threshold chosen, for 224 bits with 5 bits of drift, and lot sizes of 100,000 die. If the ID cells are working properly, the chances of even a single ID failure are vanishingly small; a fab could produce millions of wafer lots without a single misidentification.

Figure 4: ID Error Probability, 100K devices per lot



Larger lot sizes will increase both curves, the false-positive curve by the square of the increase. If we wished to identify a device without knowing its lot number, the ID of every device ever made must be considered. For a device with a production history of 100 million components, the false negative curve in Figure 4 would be multiplied by 1000, and the false positive curve by 1,000,000, resulting in an intersection at an error probability of 10 parts per billion at a threshold distance of 32. This is still an extremely small possibility of a single identification failure for the entire production run.

Unfortunately, like any real circuit, the ID cell is subject to yield loss. SiidTech carefully designs the cells to minimize yield loss, and typically sees yields exceeding 99.98%. The remaining 200ppm defective cells can manifest themselves as stuck rows and columns, stuck comparator outputs (producing incorrect IDs that are all ones or all zeros), or clocking errors, resulting in serial ID patterns that shift unpredictably and cause false mismatches. Such IDs cannot be used for correlating test data.

So how do we detect these defective IDs? What constitutes a "failure" in a series of random bits? We detect the quality of random IDs with the addition of **typeID** bits and with the purposeful introduction of correlation in the bit pattern.

The typeID bits are 32 fixed, ROM-like bits added to the array. The array is laid out as 16 rows by 16 columns of selectable device pairs; the outer columns are typeID cells, forced by the mask to be a one or a zero, and these columns are selected to produce the first 32 output bits. This typeID bit sequence is the same on every die, and can be used to identify the chip type and the mask set revision - even the mask reticle position.

The bit sequence from an ICID cell is designed to sequentially wrap around. After power up reset, we will see 32 fixed typeID bits, 224 random bits, then the same 32 fixed typeID bits again, repeated over and over as long as we keep clocking. If there are any errors in reset or sequencing, we will not see the correct typeID bits in the ID pattern, either at bit positions 0 through 31 or 256 through 287.

Single, hard-failed rows and columns are not a major problem. This just reduces the number of effective bits. But if the entire array of supposedly random bits is stuck at one or zero, we will produce an erroneous sequence that will have a significant chance of matching the same erroneous sequence from a different defective chip. While we cannot distinguish these IDs, and conclusions based on data from the associated die are not trustworthy.

The ICID cell uses an autozero comparator to look at the voltage mismatch difference between sequential pairs. This introduces correlation; the value of a single bit pair affects two neighboring bit measurements. The correlation shows up in the output bit sequence as an increased probability of neighboring bits in the pattern being different (01 and 10) rather than the same (00 and 11). The chances of a long runlength of identical bits in a healthy array goes down as the factorial of the runlength, that is, the chances of 00 is 1/6, for 000 is 1/24, for 0000 is 1/120, and so forth. The chances of 16 sequential bits being all zero or all one is $1/17!$, or $2.8E-15$. Without correlation, a random binary sequence of 16 zeros will occur with a probability of $1.5E-5$.

We can compute an "ID quality" metric based on the sum of the logarithm of the factorial of these runlengths, divided by the number of bits. Large values of the quality factor indicate long runlengths, and suspect IDs.

The inter-bit correlation comes at the cost of a reduction of the "identifiability" of ICID sequences. It reduces the number of effective ID bits by about 9%, which means that about 18% more bits must be added to produce the same amount of identification information. However, even with the losses due to correlation, 224 bits is far more than is necessary for most production runs. The curves shown above in Figures 3 and 4 for self and others already incorporate this loss of effective bits.

The entire ICID cell typically has an area less than 100 by 50 microns in a 130nm process. It is made from short-channel "core" devices, has 3 logical interconnects (reset, clock, and serial output), and fits nicely into a scan chain or test logic. The best place for an ICID cell is underneath a power or ground bus near the test I/O. If extremely low identification failure rates are necessary, more than one ICID cell can be used per die, with the cells read out in parallel.

Extracting ID bits with ATE

ICID cells are typically incorporated into test circuitry, and are only activated on a tester, not during normal circuit operation. There are three reasons for this. First it reduces aging stress. Second, it eliminates standing currents, which waste power and confuse IDDQ measurements. Third, it insures that the ID circuit cannot be used to identify the circuit during normal operation. Such identification might violate security, or the privacy of the ultimate user of the equipment containing the cell. Nobody wants a crowd of privacy-minded picketers outside their fab.

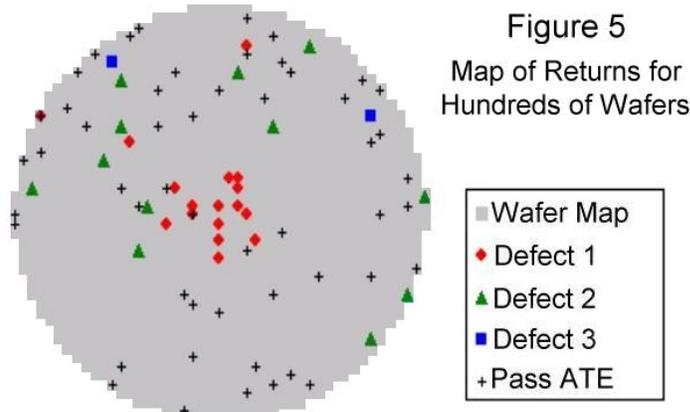
Most digital testers push input bits and expected output bits towards the test head, and expect only a single pass/fail bit in return. This makes extracting unpredictable bits from a random ICID bit more difficult. Typically, ICID bit readout involves stopping the tester and reading the output pin error buffer every bit time. This can make the whole ICID test take 10's or 100's of milliseconds. This is the major cost of using ICID in production. If an output capture memory (say, for IEEE 1149.1 scan output capture) is available, the test time can be reduced to a few hundred microseconds, as the array can be clocked faster than 10MHz.

ICID at LSI Logic

LSI Logic designs and produces custom mixed-signal integrated circuits on a number of deep submicron processes, and production ICID structures are being incorporated into most new designs. ICID technology has made possible a number of new test techniques, and is being used to increase yield and reliability of LSI processes, while reducing process monitoring cost and time to market. We will describe four interesting new measurements made possible by individual die trace using ICID. Others will be presented in another paper at the International Test Conference in October 2004 [6].

ICID technology can be used to retrieve test results and wafer position for a customer's post board assembly failures. One high volume circuit in production for a customer has had very low DPM (defects per million) at board test. However, this customer and LSI are striving for ever lower board failure rates, and have been using ICID to learn about the few defects that remain.

Figure 5 shows a map of the original wafer positions of 84 returned parts. Out of hundreds of wafers and hundreds of sites, very few show any defects at all; the map is a composite of all wafers with defects. Many of the returned parts showed no errors ("Pass ATE"); those parts will be used to develop tests with higher fault coverage. However, the rest of the returned parts showed three different types of defects that retesting *did*



catch. Defect types 2 and 3 are scattered randomly around the wafer map. However, 14 of the defect type 1 returns are clustered around the center of the wafer map. Even though these defective die are from different lots and wafers, there is some correlating influence in wafer manufacture that increases the probability of these rare defects in the center of the wafer. Analysis is going on right now to find the source of this defect, in design or in processing, and the result will be a lower defect rate for the customer.

Since ICID facilitates the observation of parametric variation of devices over time, it permits "dynamic testing" and "dynamic binning" of devices. This allows a new kind of test, called "Comet Hunting". Parametric variation of well designed devices are distributed in multidimensional clusters, looking much like an astronomical star cluster. Widely deviant devices, with parameters that exceed test limits, are like lone stars outside the cluster, and are identified and discarded by testing. However, the real concern for high reliability test is to identify the devices that may drift outside those bounds during assembly or use, even if they pass during test. The parametric variation of these devices resemble the movement of comets across photographic plates, and such variation is easily observed with ICID.

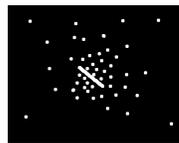


Figure 6: "Comet Hunting"
Detecting the parametric drift
that can lead to field failure

Figure 6 illustrates a comet, the streak representing a device that is near the center of the distribution but likely to fail later. LSI Logic will present techniques for comet hunting at the 2004 ITC [6].

Typically, LSI Logic strives for very high yields. There are obvious cost advantages, of course, but high yield parts have lower return rates and better DPM. Defects tend to correlate on wafers, so "good" die on wafers with lower-than-expected yields often have high field failure rates. Typically, wafers with low yield at wafer test are scrapped,

"good" die and bad alike.

However, some tests involve high currents, high speed, or other measurements that are difficult to make on an automated tester and with a probe card. These tests are deferred until package test. If a production lot of devices has a lower-than-expected yield at package test, this is often because one wafer out of the lot has a high defect rate. Without die identification, the only way to keep that hypothetical suspect wafer out of the production flow is to scrap the entire lot. With die identification, all the die from the suspect wafer can be identified and removed, after package test. This has happened a number of times at LSI, and the isolation of these suspect parts while retaining most of the lot has resulted in significant cost savings and defect reduction.

Figure 7 illustrates this discovery process, without and with ICID.

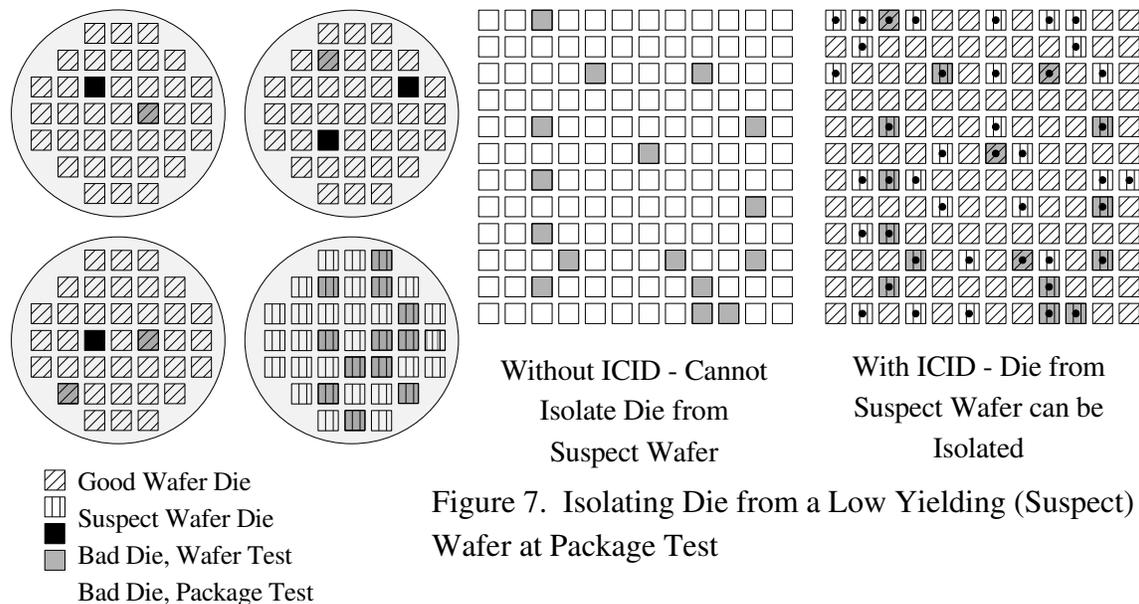


Figure 7. Isolating Die from a Low Yielding (Suspect) Wafer at Package Test

One unexpected result of ICID at LSI Logic is the detection of data handling errors during process and product qualification. LSI typically measures groups of 150 die to characterize processes and products. There are usually many groups tested, with groups skewed for different thresholds, polysilicon CDs, and other process variables. During one series of measurements, two groups of die were inadvertently swapped during measurement. The resulting data suggested an unacceptable process shift, which would have seriously impacted the introduction schedule for an important new product. However, the discrepancy was easily resolved with the accompanying ICID information. The data were restored to the proper place, and characterization and qualification continued on schedule.

Conclusion

The SiidTech ICID technology is helping LSI Logic reduce both cost and defect rate for new processes and designs, while reducing errors and time-to-market. The ICID technology uses existing digital CMOS processes without modification, and scales well into the deep submicron. The addition of ICID technology into the test flow is resulting in the development of new test techniques, which LSI will use to improve its product offerings and competitive advantage in the future.

References:

- [1] "Statistical Post-Processing at Wafersort - An Alternative to Burn-in and a Manufacturable Solution to Test Limit Setting for Sub-micron Technologies", R. Madge, M. Rehani, K. Cota, W. Daasch, 20th IEEE VLSI Test Symposium, pp 69-75.
- [2] "Screening MinVDD Outliers Using Feed-Forward Voltage Testing", R. Madge, B. Goh, V. Rajagopalan, C. Macchietto, R. Daasch, C. Schuermyer, C. Taylor, D. Turner, 2002 International Test Conference, pp. 673-682.
- [3] "Successful Implementation of Structured Testing", R. Richmond, 2000 International Test Conference, pp. 344-348.
- [4] "IC Identification Circuit using Device Mismatch", K. Lofstrom, W. Daasch, D. Taylor, 2000 IEEE International Solid-State Circuits Conference Digest of Technical Papers Volume 43 - IEEE Cat. No. OOCH37056. (<http://www.kl-ic.com/isscc2K.pdf>)
- [5] "A PROM Element Based on Salicide Agglomeration of Poly Fuses in a CMOS Logic Process", M. Alavi, M. Bohr, J. Hicks, M. Denham, A. Cassens, D. Douglas, M. Tsai, 1997 IEEE International Electron Devices Meeting, pp. 855-858.
- [6] "Feedforward Test Methodology Utilizing Device Identification", A. Cabbibo, M. Jacobs, J. Conder, submitted to the 2004 International Test Conference.