Transition Delay Fault Testing of 3D ICs with IR-Drop Study

Shreepad Panth and Sung Kyu Lim
School of Electrical and Computer Engineering
Georgia Institute of Technology
Email: {spanth, limsk}@ece.gatech.edu

***Abstract*—In order to ensure the correctness of 3D ICs, they need to be tested both before and after their individual dies are bonded. All previous works in the area of 3D IC testing consider only stuck-at fault testing. However, 3D ICs also need to be tested for delay defects. In this work, we present a transition delay test infrastructure that can be used to test a 3D IC both before and after bonding. Furthermore, we present a methodology to test the through silicon vias (TSVs) after bonding, without necessitating regeneration of test patterns. Results show that the overhead involved is negligible. In addition, at-speed testing of circuits can suffer from large IR drop problems. In this paper, we also study the IR drop of 3D ICs during transition delay fault testing. We study how different configurations of probe pads affect the pre-bond IR drop. We also study how this IR drop changes from the pre-bond to the post-bond case.**

I. Introduction

3D ICs are manufactured by fabricating each die separately, thinning the dies containing TSVs, and stacking them all together. Due to the additional manufacturing steps of thinning and stacking, it is possible to introduce additional defects into the circuit. Therefore, these ICs need to be tested both before stacking (pre-bond test), and after stacking (post-bond test). Furthermore these vertical interconnections also need to be tested, to verify their operation.

In this paper, we present a DfT architecture that supports transition delay fault testing of 3D ICs. It supports both pre-bond, and post-bond transition testing. In addition, it supports transition testing of TSVs after bonding. Since transition patterns are applied at the rated frequency, there could be severe IR drop problems. For this reason, we also study the IR drop during transition testing of 3D ICs. We explore what parameters affect the IR drop, and how it changes from pre-bond test to post-bond test.

There has been some work done in stuck at fault testing of 3D ICs. One of the first works that talked about 3D IC testing was presented by [1], which was a preliminary attempt at applying something similar to IEEE 1500. This idea was then formalized in [2], [3]. There also exists literature [4], [5] supporting transition delay testing of 2D SoCs. Only one work has considered probe pad placement [6], but it did not explore design options, or consider a realistic test architecture.

One of the reasons 3D ICs are being explored is because they are expected to be faster than 2D ICs. What this means is that it is essential to test them for delay defects. However, there has been no work in developing a transition delay fault infrastructure for 3D ICs. Furthermore, there has been no study as to how the probe pads are to be added into the layout, and how their configuration affects the IR drop. This paper aims to address these issues.

II. Transition Delay Fault Testing of 3D ICs

A. Transition Fault Capable Wrapper Boundary Register

The application of a transition fault vector to a circuit requires two cycles. The first cycle triggers a transition (launch) at the location to be tested, and the second cycle (capture) captures the response to this transition. We cannot use the IEEE 1500 Wrapper Boundary Registers (WBR) specified in [2], since they support the application of only a single bit to a primary input, and two bits are required to launch a transition. Instead, we use a three flip flop IEEE 1500 WBR specified in [4]. Such a register is shown in Figure 1. This figure also explains abbreviations that will be used in the remainder of this paper. Each flip flop is sensitive to a different combination of IEEE 1500 control signals, which are indicated above the clock. In order to apply a transition test, we scan in one bit each into the SC and ST flip-flops, and apply them sequentially through the Update register.

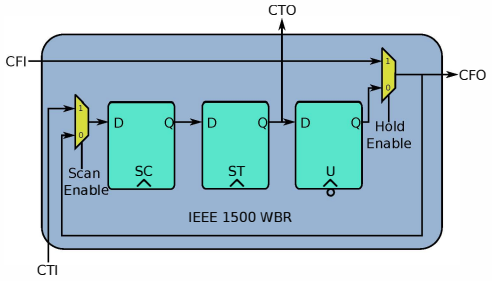

Fig. 1. An IEEE 1500 Wrapper Boundary Register capable of launching a transition on CFO. The abbreviations used are S:shiftWR, C:captureWR, T:transferWR, U:updateWR

B. Transition Fault DfT Architecture

Our transition fault DfT architecture is shown in Figure 2. In order to simplify the block diagram, we reduce its complexity by showing only the data path, and omitting the serial to parallel conversion. Parallel testing is essentially the same idea, but with a greater number of scan chains.

Each TSV is equipped with a WBR, so that values can be scanned into it during test. Once the values are scanned in, the launch and capture clocks are applied, and the responses are scanned out. Each die is tested independently of the other, during both the pre-bond and post-bond tests. Each unwrapped die is equipped with an internal bypass, so that the internal scan chains can be bypassed, if desired. In order to transport data to and from the top die, the bottom die is equipped with a multiplexer (elevator enable) to select the data from the top die. The various control signals are generated by the IEEE 1149.1 TAP controller, but it is not explained here due to space constraints.

This architecture is similar to that presented in [2], but with a few notable differences. The first one being that we use a transition fault capable wrapper boundary register. The second one is that our system has to support the transfer operation, in order to transfer data between the SC and ST registers. Therefore, an extra transfer signal is to be routed between the dies. In addition, the IEEE 1149.1 TAP controller does not natively support the application of delay tests, and certain modifications are warranted. Two approaches exist in the literature. The first one ([4]), uses the exit1-DR, exit2-DR and pause-DR states of the IEEE 1149.1 FSM to generate update, transfer and capture signals, while in delay test mode. The second approach ([5]), utilizes an additional TMS bit to change the state from update-DR to capture-DR within a single clock cycle. We choose the first approach, because additional package pins are undesirable.

978-1-4673-1074-1/12/$31.00

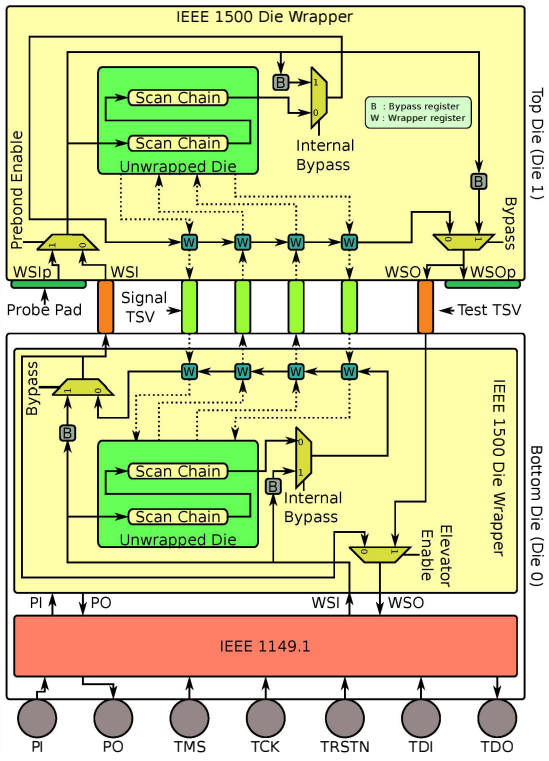

Fig. 2. Our DfT Architecture for Transition Delay Fault Testing of 3D ICs, showing only the data path and serial operation

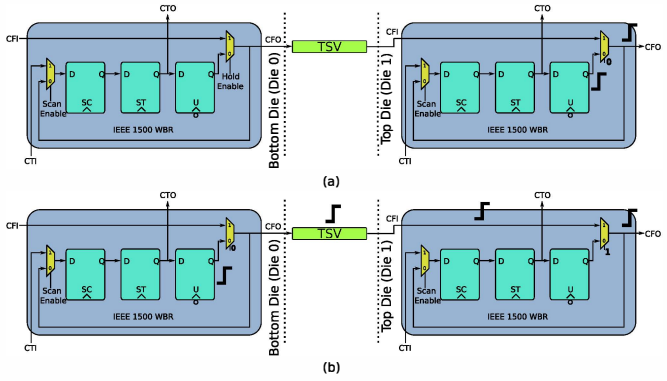

Fig. 3. (a) A 0 to 1 Transition launched from WBR on Top Die (no TSV testing), (b) An equivalent 0 to 1 Transition launched from WBR on Bottom Die (with TSV testing)

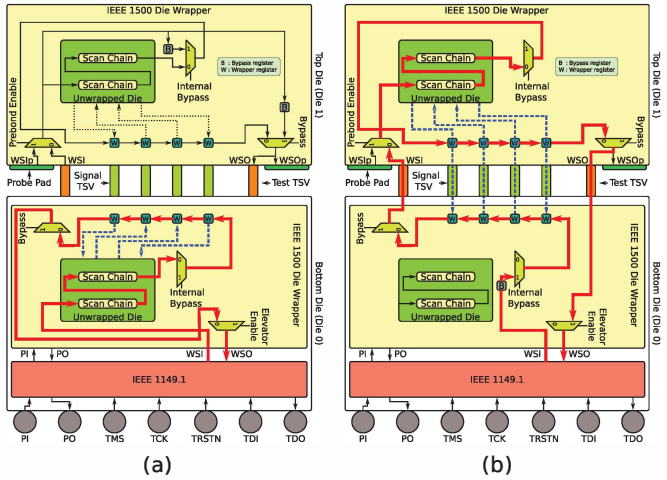

Fig. 4. (a) Post-bond test of bottom die, (b) Post-bond test of top die with TSV test. Solid red lines indicate flow of scanned data, and dashed blue lines indicate flow of data to and from WBRs in the launch-capture window

C. TSV Testing

If we are concerned only with stuck-at testing, TSV testing is trivial. Each TSV has a WBR on either side, and TSVs can be tested by placing both dies in their respective extest modes. However, for transition testing, the time between the launch and capture pulses has to be of the order of the TSV delay. This is a few tens of picoseconds, and it is unreasonable to assume that the clock can be applied with such a high speed.

We propose an alternate approach to test the TSVs after the dies have been bonded. Consider Figure 3(a). This represents the post-bond testing of the top die, with a transition launched from the WBR on the top die. Figure 3(b) shows the identical transition on the top die, but launched from the WBR on the bottom die. This transition would also occur on the TSV, and would hence test the TSV also. This implies that the a test vector generated for the top die, but launched from the bottom die will also test TSVs. If we perform testing of the top die exclusively through the WBRs of the bottom die, no additional patterns will be required, and all TSVs between the top and bottom die will be tested.

In order to support TSV test, we need an additional mode of operation that configures the WBRs as shown in Figure 3. The default modes presented by [2] are serial/parallel , pre-bond/post-bond, intest/extest/bypass and turn/elevator. We add another mode called TSVtest. If a die is placed into TSVtest, all WBRs facing the bottom die are made transparent. TSV test can then be performed by placing the bottom die into extest, and the top die in the intest_TSVtest mode.

Two example modes of operation are shown in Figure 4. Figure 4(a) Shows the post-bond test of the bottom die. The instruction used is post-bond_intest_serial_turn. Figure 4(b) shows the post-bond testing of the top die with TSV test. Here the bottom die is programmed with post-bond_extest_serial_elevator, and the top die with post-bond_intest_serial_turn_TSVtest. The solid red lines show the flow of data scanned in, and the dashed blue lines show the data flow to and from the WBRs in the launch-capture window.

D. Probe Pad Placement

In the case of pre-bond test, data needs to be provided to the bare die. For the bottom die, this is not a problem as it has bumps that can be directly probed. But the top die needs to be provided with probe pads for test access. Each IEEE 1500 data and control signal needs to be provided with its own probe pad. In addition to these, the die needs to be powered during test. This means that we need to provide power and ground probe pads as well. In our study, we focus on circuits that have a regular power and ground TSV placement as shown in Figure 5(a). Since the power and ground TSVs form a regular array, the space in between them are candidate locations for probe pads. Our power/ground and signal probe pads can be placed in a subset of these candidate locations. An example is shown in Figure 5(b). We have two choices of connecting a power probe pad to a power TSV, either in a horizontal or a vertical configuration. This figure also shows two signal probe pads. To simplify the design process, and reduce the search space, we place power and ground probe pads in either the horizontal or the vertical configuration, but not both. Figure 5(c) shows how we choose to place 4 power probe pads in a 2×2 array, in a horizontal configuration, and Figure 5(d) shows the same for a vertical configuration.

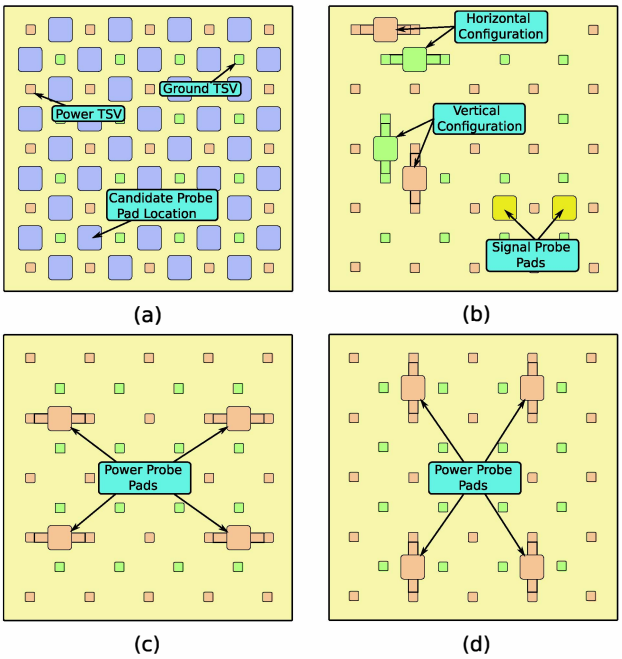

Fig. 5. (a) Candidate locations for probe pads, (b) Sample horizontal and vertical power/ground pads, as well as signal pads, (c) 4 power probe pads placed in a 2 ×2 horizontal configuration, and (d) in a vertical configuration

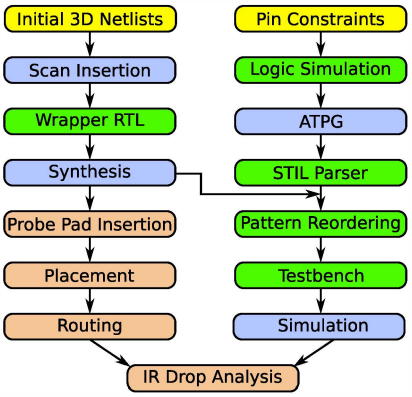

Fig. 6. Our Design Flow. Yellow indicates inputs to the flow, green boxes are custom scripts, blue indicates use of Synopsys tools, and red the use of Cadence tools

TABLE I
DESIGN STATISTICS FOR TWO DESIGNS, SPLIT BY DIE

	Jpeg		FFT	
	Bottom Die	Top Die	Bottom Die	Top Die
Gate Count	214,641	197,187	328,512	296,929
# Scan F.F	15,828	22,219	87,681	78,503
# Signal TSV	2,164	-	2,879	-
S-A Coverage (%)	99.77	99.61	99.99	99.99
S-A Patterns	2012	2217	12180	11610
Tr Coverage (%)	98.93	97.74	99.92	99.90
Tr Patterns	3892	5200	61,798	55,656

III. DESIGN FLOW

The design flow used in this paper is shown in Figure 6. It can be broadly divided into two categories. The left column represents physical design, and the right column is test related. Finally, IR drop analysis is performed. Each step is explained individually

With respect to physical design, we start with initial 3D gate level verilog netlists, generated by partitioning a 2D netlist. We then insert as many scan chains as required, using Synopsys Design Compiler. Our custom script takes this netlist with scan chains, and generates the RTL for the IEEE 1500 wrapper. This is then re-synthesized using Synopsys Design Compiler. We then insert probe pads into the layout, and treat these probe pads as locations where other TSVs cannot be placed. The design is then placed and routed using Cadence Encounter.

The test related steps starts with pin constraints, which are any pins that need to be constrained to a certain logic value during the test mode. We perform logic simulation on the bottom die, to get the pin constraints on the top die. Using this information, we perform ATPG on both dies using Synopsys Tetramax. We then parse the STIL files generated, and using the information about the wrapper chain ordering from the physical design stage, reorder the bits in the pattern using our custom script. We then generate the testbench, and simulate it using Synopsys VCS. Using the routed result, and the VCD file generated from the testbench, we perform IR drop analysis as described next.

For 2D IR drop analysis, as is the case with all pre-bond testing, we can simply use existing tools. However, 3D IR drop analysis is required to measure the post-bond voltage drop. We first perform power simulations die by die, using the switching activity from the VCD file, after annotating each die with TSV parasitics. We combine the DEF files from both the dies into a single DEF file, treating the TSV as a via. This tricks the tool into believing that we are dealing with a 2D design, but with a higher number of metal layers. We then use the power numbers generated earlier, to perform 3D IR drop analysis using Cadence encounter.

IV. EXPERIMENTAL RESULTS

All required scripts were implemented in C++. Our designs are synthesized using the nangate $45nm$ technology library. In this study, we assume the TSV diameter to be $4\mu m$, and its height to be $40\mu m$. The TSV landing pads size is assumed to be $7\mu m$, and the total TSV cell size including keep out zone is $8.4\mu m$. Power and ground TSVs are placed in a regular fashion, with a pitch of $130\mu m$. The TSV resistance, including contact resistance is considered to be $50m\Omega$. We assume that our probe pads have a size of $40\mu m \times 40\mu m$, and that the minimum pitch is $100\mu m$.

Figure 7 shows a sample testing waveform of a design with four scan chains. It also shows a pass/fail bit. During capture, the responses from the circuit are stored into the SC register, and the value of the ST is shown as a don't care. Only the first vector scanned out exhibits this don't care, and all subsequent vectors have a junk value in the ST register. Some sample layouts are shown in Figure 8.

We consider two designs, both from the OpenCores benchmark suite. We implement them in two dies due to limitations on the number of layers that encounter can handle while performing IR drop analysis. Design statistics are shown in Table I. This table splits up the statistics on a die by die basis. The top die does not have any TSVs, and hence its entry is blank. We also report the results of ATPG for both stuck-at faults, as well as transition faults.

In all the following experiments, each die is assumed to have five scan chains. Since the power consumption of stuck at tests can be controlled by reducing the frequency, all power numbers and IR drop results focus on transition tests. We also choose five transition test vectors to associate with each die. The test vectors of the bottom die are prefixed with "BD" , and those of the top die with "TD". Since ATPG runs in a greedy fashion, the first few vectors test a larger number of faults per vector than later vectors. We therefore choose five vectors at random out of the first few generated, to obtain patterns with high switching activity.

Since we test only a single die at a time in our experiments, the

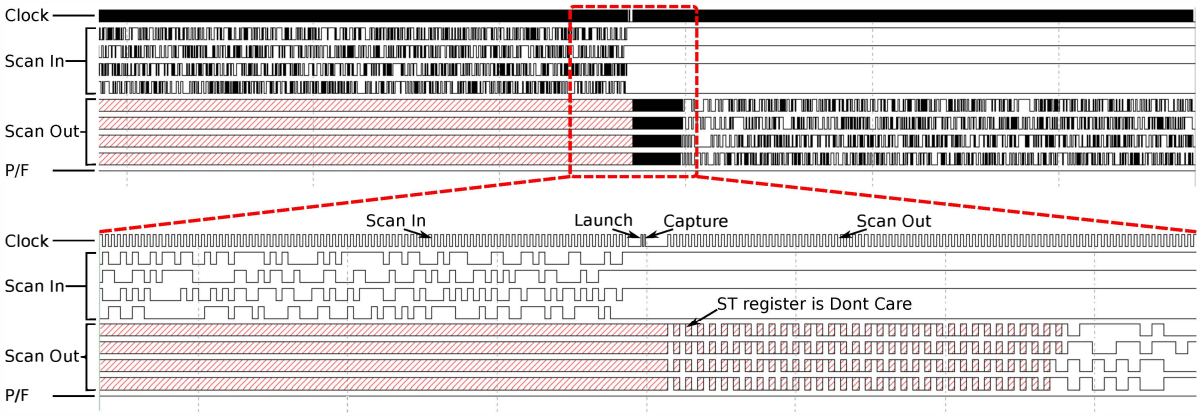

Fig. 7. A sample waveform obtained during testing, designed with four scan chains

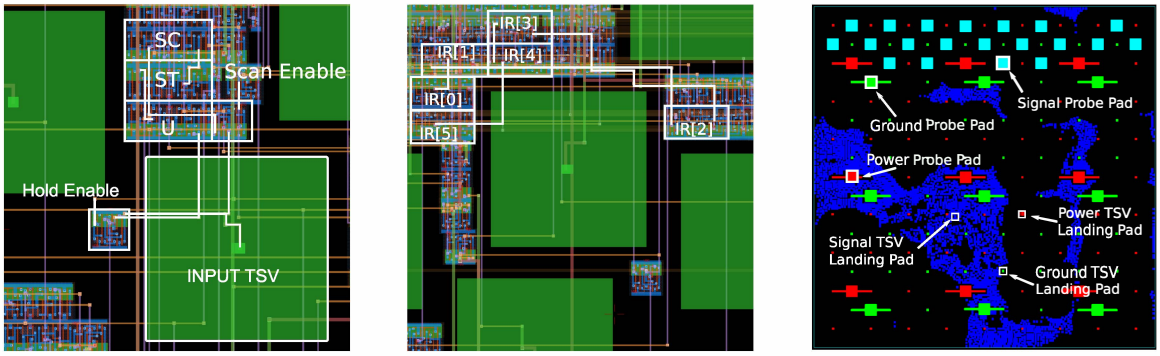

Fig. 8. GDSII images. (a) A close up of a TSV and its WBR, (b) IEEE 1500 Instruction Register Chain, (c) zoom out shot of the top metal layer of the top die, showing TSV landing pads and probe pads

TABLE II
POST-BOND TEST TIME RESULTS. ALL TEST TIMES ARE IN CYCLES

Design	Die	Stuck-at test			Transition test		
		[2]	Ours	% Increase	Without TSV test	With TSV test	% Increase
FFT	Bottom	220,646,633	227,662,889	3.17	1,155,085,107	-	-
	Top	189,003,857	195,703,404	3.53	938,154,390	1,002,271,254	6.83
Jpeg	Bottom	7,246,799	8,118,428	12.02	15,704,360	-	-
	Top	10,819,403	11,779,797	8.87	27,632,911	32,136,977	16.29

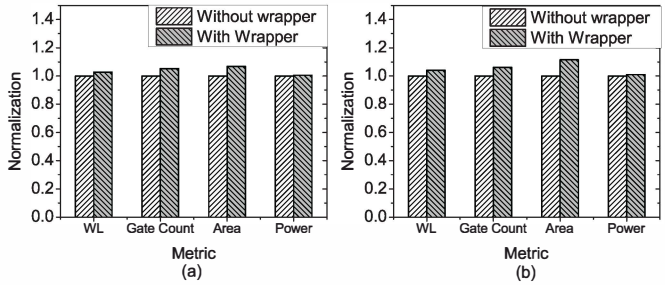

Fig. 9. Various overheads involved in adding wrappers for (a) FFT and (b) Jpeg

clocks to all the scan flip flops of the die not being tested are gated off. This helps reduce power consumption.

A. Overhead Study

In this section, we calculate the overhead involved in adding the IEEE 1500 wrappers to our design. We compute this overhead with respect to wirelength, gate count, area, and power. This is shown in Figure 9. From this graph, we can see that there is around a 10% increase in gate area for jpeg, but this reduces to 5% in the case of FFT. This is because this design has a smaller TSV to gate ratio. For both designs, the wirelength and gate count increase by less than 5%. We also see only a very small increase in the power consumption of both circuits. This because the test related elements do not switch during the normal operation, and any power increase comes only from the small increase in the wirelength, as well as increased leakage.

B. Test Time Study

Here, we observe the test time change for different configurations, and different types of test. These results are shown in Table II. We report the test time for post-bond test only, as the number of vectors is identical in the pre-bond case. The third and fourth column refers to the test time obtained by running stuck at tests only. We compare the test time for running stuck at tests with [2], which implements a stuck-at architecture only. Since in our flow, each WBR has one additional flip-flop, the test time is expected to increase. It is observed that this increase reduces with an increase in the circuit size.

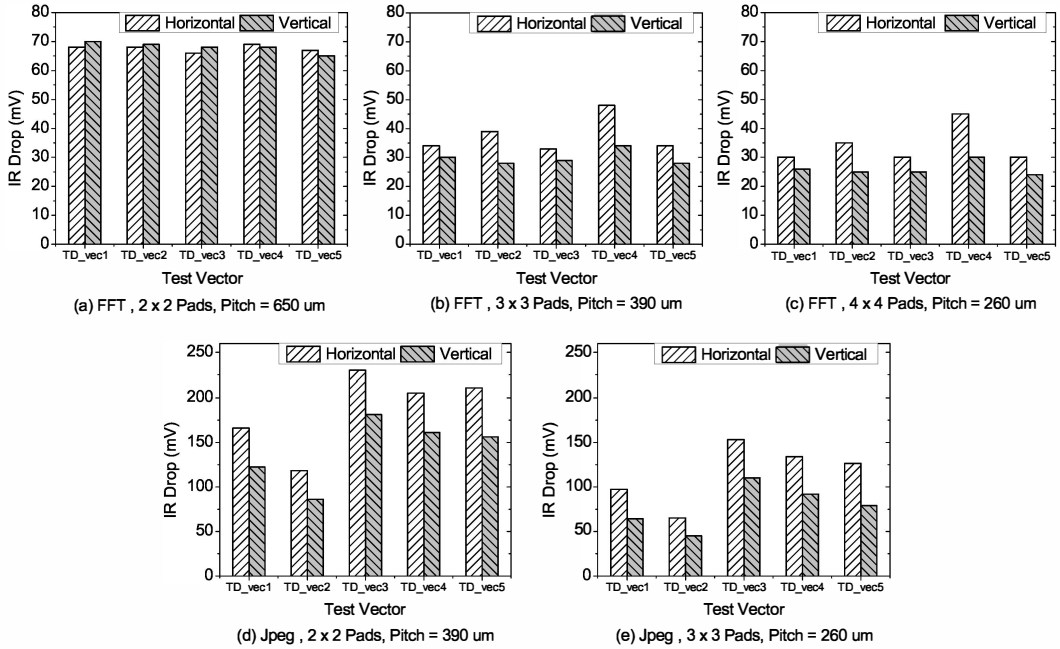

Fig. 10. Pre-bond IR drop under different probe pad configurations and test vectors for FFT (a, b, c) and Jpeg (d, e)

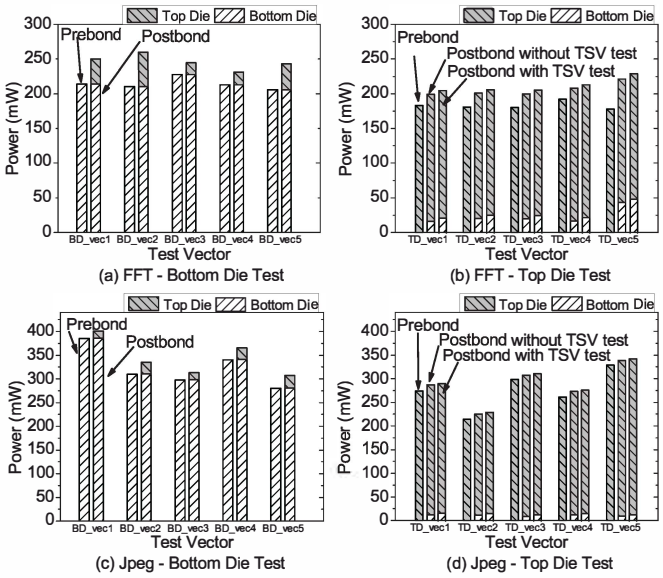

Fig. 11. Power comparison among (1) pre-bond, (2) post-bond without TSV test, and (3) post-bond with TSV test under five different test vectors. We show the total power consumption in each die.

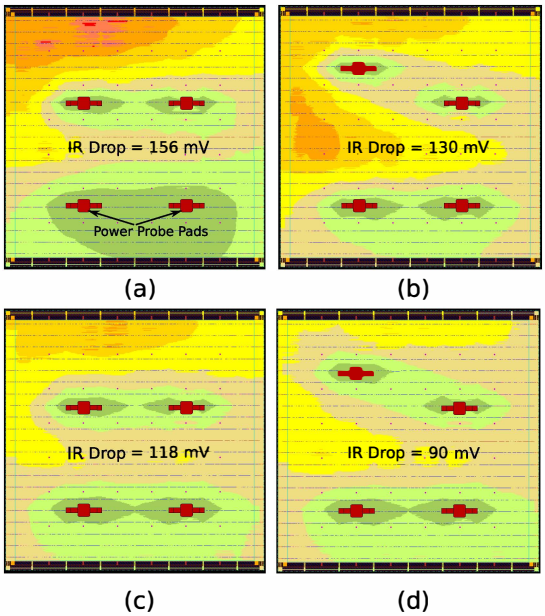

Fig. 12. IR-drop maps before (= a, c) and after (= b, d) probe pad optimization. We use TD_vec1 for (a) and (b) and TD_vec2 for (c) and (d).

Next, columns 6 and 7 compare test times of the top die, when tested through its own WBR, as opposed to through that of the bottom die. This corresponds to testing of the top die without, and with TSV test. Since the latter case has a longer chain length, the test time increases. Again, this increase is observed to be proportional to the circuit size. If this increase is found to be unacceptable, we can always bypass the WBR chain in the top die, incurring some additional area and wirelength costs due to extra multiplexers.

C. Power Study

In this section, we study how the power consumption changes with choice of pattern, as well as from pre-bond to post-bond test. In the case of the top die, we also compare post-bond without TSV test, and post-bond with TSV test. These results are plotted in Figure 11. The total power consumed in each case is split into the contribution by each die. From these graphs, we can see that the power consumed by a particular die changes very little when we move from pre-bond to post-bond test. However, the other die consumes some additional power due to leakage, and switching in the test circuitry, leading to an increase in the overall power. Furthermore, when the top die is tested in conjunction with TSVs, the power consumed by both dies increase, compared to the case when TSVs are not tested. This is because the logic driving TSVs in each die now consumes more

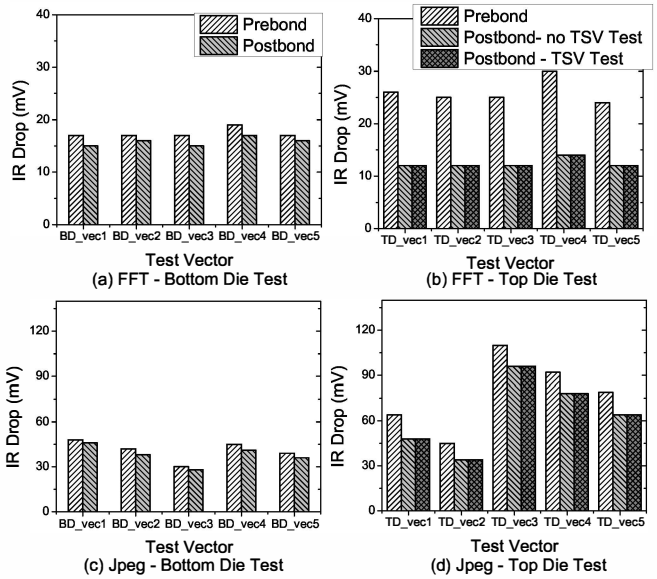

Fig. 13. Comparison between pre-bond and post-bond IR drop. (a) FFT, bottom die, (b) FFT, top die, (c) Jpeg bottom die, (d) Jpeg, top die

power.

D. Pre-bond IR drop

Here we study how different configurations of power probe pads affect the voltage drop during the pre-bond test. Since the bottom die receives power from solder bumps, it is of no interest to us in this study, and we focus on the top die alone. As mentioned earlier, we place the probe pads in a regular grid like fashion, at different pitches, trying different configurations. The results are shown in Figure 10.

We see that by reducing the pitch, the IR drop goes down, as expected. What is interesting to note is that the vertical configuration almost always outperforms the horizontal configuration. This is because the standard cells receive power from horizontal metal stripes, and placing pads in a horizontal configuration would simply mean that the same stripes get power at two locations. However, in the vertical configuration, it is easy to see that more of these stripes will get a direct connection to power, and hence the IR drop reduces.

As observed for the 2×2 configuration of probe pads of the circuit jpeg, the IR drop can be quite high. One obvious solution would be to go back to ATPG, and constrain the power budget. This would increase the total number of vectors, and hence the test time. Here, we take another approach, and investigate whether any improvement in the IR drop can be achieved by cleverly placing probe pads. We pick this configuration (Jpeg , 2×2 , horizontal) since it has highest IR drop, and try to manually optimize it.

After some trial and error, we were able to reduce the IR drop. IR drop maps are shown in Figure 12. We also enumerate the worst case IR drop before and after optimization in Table III. This shows that with this optimized configuration, the maximum IR drop can be reduced for all test vectors considered. Therefore, a careful choice of probe pad locations can reduce the IR drop.

E. Pre-bond versus Post-bond IR drop

We now study how the voltage drop of a particular die changes, depending on the stage in the bonding process. These results are plotted in Figure 13. In the case of the top die, we plot the lowest pre-bond voltage drop achieved among all possible combinations. Not surprisingly, the post-bond IR drop of the top die is much lower than the pre-bond case. This is because in the post-bond case, the top die receives power through TSVs at a much finer pitch than the probe pads in the pre-bond case. The small increase in the power consumption, when tested with TSVs is not sufficient to cause any change in the IR drop.

TABLE III
PROBE PAD OPTIMIZATION FOR THE TOP DIE OF JPEG. WE USE FOUR POWER PROBE PADS IN A HORIZONTAL CONFIGURATION.

Vector	IR drop (mV)	
	Before	After
TD_vec1	156	130
TD_vec2	118	90
TD_vec3	230	188
TD_vec4	205	167
TD_vec5	211	166

It is interesting to note however, that the IR drop of the bottom die also reduces slightly during post-bond, even though it still receives power from the same locations, and has a slightly higher power consumption. This is because during the post-bond test of the bottom die, the top die consumes very little power, yet attaches its entire power grid in parallel to that of the bottom die. This reduces the equivalent resistance of the power grid, and hence the IR drop is lower.

F. Normal vs Test Mode

Since transition fault testing aims to switch as many nets as possible with one vector, we expect the IR drop during the test mode to be much higher than the IR drop during the normal mode. The normal mode IR drop of Jpeg was found to be $10mV$, and that of FFT was found to be $6mV$. When compared with the post-bond numbers from Figure 13, we see that test mode has much higher IR drop.

V. CONCLUSION

In this paper, we presented a transition delay fault architecture for 3D ICs. We showed that there is minimal overhead when running stuck-at tests on this modified architecture. In addition, we provided a means by which TSVs can be tested for delay defects after bonding, without the need for regeneration of new test patterns. Although there is some increase in the test time with this scheme, it can be mitigated by using additional multiplexers. We also studied the IR drop issue of 3D ICs during transition test. We observed that the power consumed during the test increases from the pre-bond to the post-bond case. There is a further increase in power for the top die, if we also perform TSV testing. We showed that vertical placement of power probe pads gives us a smaller IR drop than horizontal placement. We further demonstrated that it is possible to optimize the locations of probe pads to reduce the IR drop. Lastly, we observed that although the post-bond test has higher power consumption, it always gives us lower IR drop.

REFERENCES

[1] D. Lewis and H. H. S. Lee, "A Scan Island Based Design Enabling Pre-Bond Testability in Die Stacked Microprocessors," in *IEEE International Test Conference*, 2007.

[2] E. J. Marinissen, C.-C. Chi, J. Verbree, and M. Konijnenburg, "3D DfT architecture for pre-bond and post-bond testing," in *IEEE International 3D Systems Integration Conference*, nov. 2010, pp. 1 –8.

[3] E. Marinissen, J. Verbree, and M. Konijnenburg, "A structured and scalable test access architecture for TSV-based 3D stacked ICs," in *IEEE VLSI Test Symposium*, 2010, pp. 269 –274.

[4] C.-Y. Lo, C.-H. Wang, K.-L. Cheng, J.-R. Huang, C.-W. Wang, S.-M. Wang, and C.-W. Wu, "STEAC: A Platform for Automatic SOC Test Integration," *IEEE Trans. on VLSI Systems*, 2007.

[5] P.-L. Chen, J.-W. Lin, and T.-Y. Chang, "IEEE Standard 1500 Compatible Delay Test Framework," *IEEE Trans. on VLSI Systems*, 2009.

[6] S. Panth and S. K. Lim, "Scan chain and power delivery network synthesis for pre-bond test of 3D ICs," in *IEEE VLSI Test Symposium*, 2011.