System-on-chip Scheduling Benchmarks for New Technologies

Spencer Millican,
Auburn University
millican@auburn.edu

Kewal Saluja
University of Wisconsin - Madison
saluja@engr.wisc.edu

Abstract – The benchmarks presented in this article are a set of power traces and stacked die specifications for evaluating SoC test scheduling methods. These benchmarks are built on previous benchmarks with the intent of providing information needed for power and thermally constrained test scheduling which is not provided in previous benchmarks. These benchmarks are free to use for scheduling tests (and other tasks) on SoCs in an (optional) 3D-stacked circuit environment. These benchmarks allow researchers to evaluate their scheduling methods while forgoing the time and effort of developing unique benchmarks for their experiments.

I. Motivation & Background

Benchmark circuits are a necessary component to electronic design automation (EDA) research, otherwise high-quality and consistent results would not be obtainable. To evaluate the effectiveness of new EDA algorithms, benchmark circuits from various sources are needed. These benchmarks need to accurately represent circuits used in industry, otherwise research results may not be relevant to industry. Examples of current benchmark circuits include the ISCAS'85 [1], ISCAS'89 [2], and ITC'99 [3] benchmark circuits, which are still used to this day.

The first set of benchmarks created for the purpose of SoC Testing were the ITC'02 Benchmarks [4]. These benchmarks presented several SoCs with information required to evaluate hardware-constrained and pin-constrained test scheduling. However, these benchmarks lack information required for other test scheduling environments. For instance, most of the ITC'02 Benchmarks lack power information for tests, and those which have power information lack detailed power traces. This makes the ITC'02 Benchmarks inadequate for power-constrained test scheduling or for other constrains that rely on power, such as temperature-constrained test scheduling. Researchers in these fields were required to generate their own benchmarks for each of their studies.

As new design constraints and design techniques are introduced with newer generations of technology, new benchmarks must be introduced so that these issues can be addressed by the EDA community. Notable examples of new issues include increasing power density resulting from the failure of Dennard Scaling [5], which in turn has increased circuit temperature density. Failing to confront these issues will result in inferior or unreliable circuits. New technologies, like 3D integrated circuits (3DICs), must also be addressed by the EDA community, as old algorithms which performed well for previous generations of technologies may not be directly applicable to newer generations. Unfortunately, older benchmarks cannot be directly used to address these issues.

The benchmarks provided here allow for researchers to forego the process of creating benchmarks, therefore allowing more time and effort to be spent on research quality. Since these benchmarks are directly based off of the ITC'02 Benchmarks, a diverse set of experiments can be performed using these benchmarks when multiple design constraints must be addressed.

II. Benchmark Creation

Each benchmark consists of one or SoC dies which creates a 3DIC (when more that 1 SoC die is specified). Each die is composed of one or more modules, and each module (core) has one or more tests (tasks) assigned to it, and each test has its own dynamic power trace.

A hierarchical floor plan layout is given for every SoC (die) in a benchmark. Each SoC (die) is based off of an ITC'02 Benchmark [4]. The area of each module (core) is proportional to the number of input/output pins of the module. Each SoC (die) has a 10mm by 10mm area, which can be scaled to fit a given target technology.

TABLE 1
BENCHMARK INFORMATION

Benchmark Name	# Dies	# Tests	Longest Test (power cycles)	Avg. Pow. Density	Highest Temperature
1-1-LLL	1	9	3049	1.09	29.07
1-2-MML	1	15	164	1.37	27.33
1-3-LLL	1	10	24	1.12	26.87
1-4-LLL	1	8	299	0.97	26.96
1-5-MHL	1	14	36	1.62	26.86
1-6-LML	1	4	842	1.57	28.97
1-7-LHL	1	4	5561	1.79	32.83
1-8-HHL	1	28	256	1.70	28.06
1-9-MML	1	19	1363	1.56	30.17
1-10-HML	1	32	287	14.33	27.83
1-11-HHL	1	31	13070	1.89	31.16
1-12-LHL	1	7	19353	1.89	34.5
2-1-MML	2	29	164	1.49	30.28
2-2-MLL	2	25	164	1.24	30.16
2-3-MML	2	24	36	1.37	26.87
2-4-LML	2	12	842	1.34	30.41
2-5-HHL	2	42	256	1.66	27.65
2-6-MMM	2	23	1363	1.45	36.37
2-7-HML	2	60	287	1.59	32.7
2-8-MHM	2	35	1307	1.84	44.28
2-9-HHM	2	63	13070	1.66	42.92
2-10-MHH	2	38	19353	1.89	46.62
3-1-MML	3	39	164	1.37	30.29
3-2-MMM	3	33	842	1.44	39.81
3-3-LML	3	22	5561	1.58	34.03
3-4-HMM	3	79	1363	1.56	37.53
3-5-MHH	3	54	19353	1.72	58.53
3-6-HHM	3	70	19353	1.74	41.95
4-1-MLL	4	47	299	1.27	29.1
4-2-HHH	4	83	5561	1.62	65.7
4-3-HHH	4	89	19353	1.69	48.23
5-1-MHH	5	64	19353	1.63	52.77

Dynamic power traces were created by random circuit simulation on ISCAS'85 [1], ISCAS'89 [2], and ITC'99 [3] benchmarks. These benchmarks were first matched to corresponding cores in the ITC'02 Benchmarks [4] by matching (as close as possible) the number of input/output pins of the cores to the number of input/output pins of the circuits. The number of random vectors (plus scan cycles) applied were then calculated based on core test information given in the ITC'02 benchmarks. Since the number of clock cycles can vary depending on how circuit pins are utilized, the assumption was made that a single scan-in and single scan-out pin was available for scan chain access and the corresponding power traces can be scaled appropriately to compensate for other hardware configurations. It was also presumed each die in these benchmarks had an average power density between 0.9 to 2.0 $\frac{watts}{mm^2}$. Each entry in a power trace corresponds to the average power of 400 clock cycles, which is determined by tracking the average switching activity in the circuit over 400 clock cycles.

The original intent of these benchmarks was for use in temperature-constrained test-scheduling [6].

Temperature traces can be obtained by performing thermal simulation under an appropriate thermal environment. These benchmarks are also effective in power-constrained test-scheduling [7]. If need be, one could use the benchmarks solely for the purpose of hardware-constrained test scheduling with added 3DIC geometry information.

III. Benchmark Naming

The naming scheme used for benchmarks is intended to represent the difficulty of scheduling tests on the benchmark. A template for each benchmark name is as follows:

{# of dies in benchmark}_{Benchmark # with given number of dies}_{Relative number of tests per die}{Power difficulty}{Temperature difficulty}

The first item of the name is straightforward: how many dies are in the given benchmark. If this is "1", this implies that the benchmark is a single-die package and not a 3DIC. The second item distinguishes between benchmarks that have the same number of dies. The third, fourth, and fifth items provide a "high (H)", "medium (M)", or "low (L)" difficulty rating for the number of tests to schedule per die, the power density of the circuit, and the temperature impact of the tests when tests are simulated in a modest cooling environment. The specific information these classifications are based on are shown in Table 1. For example, the benchmark "2-7-HML" has two dies in it and is the seventh benchmark with two dies in it. It has a large number of tests relative to the number of dies it has, it has fair power density, but its tests have a minimal temperature impact in a modest cooling environment.

These benchmark names are meant to help users decide which benchmarks to use based on their research needs. However, power values and floor plans of benchmarks can be scaled to adjust the difficulty of benchmarks.

IV. Benchmark Format

The file formats of the benchmarks were originally intended to be used in the HotSpot thermal simulator [8], [9]. The file types of each benchmark are as follows.

1. **soc:** The original ITC'02 Benchmark file attached to each die.
2. **flp:** The floor plan file for each die.
3. **ptrace:** The power trace file for each test.
4. **comp:** An example compatibility file for use in studies where hardware compatibility is not the primary focus.

A modest thermal simulation environment is provided for each benchmark in a separate file. It is presumed that this and other files will be scaled to fit the needs of the researcher.

V. Conclusions

This article presents a set of benchmarks which includes information not available in previous benchmarks: 3DIC information and dynamic power traces. By combining these new benchmarks with the capabilities of previous benchmarks, more research efforts can be spent on developing new methods as opposed to developing study-specific benchmarks. It is the authors' hope that these benchmarks will be utilized by the wider EDA community beyond their original intended audience (SoC test scheduling) and current users [10]–[15], and that if needed, more information will be added to these benchmarks to better fulfill the needs to of the EDA community.

VI. References

[1] F. Brglez and H. Fujiwara, "A neutral netlist of 10 combinational benchmark circuits and a targeted translator in FORTRAN," 1985.

[2] F. Brglez, D. Bryan, and K. Kozminski, "Combinational profiles of sequential benchmark circuits," in *IEEE International Symposium on Circuits and Systems*, 1989, pp. 1929–1934 vol.3.

[3] S. Davidson, "ITC'99 Benchmark Circuits - Preliminary Results," in *International Test Conference 1999. Proceedings (IEEE Cat. No.99CH37034)*, 1999, pp. 1125–1125.

[4] E. J. Marinissen, V. Iyengar, and K. Chakrabarty, "A set of benchmarks for modular testing of SOCs," in *Proceedings. International Test Conference*, 2002, pp. 519–528.

[5] R. H. Dennard, F. H. Gaensslen, V. L. Rideout, E. Bassous, and A. R. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *IEEE Journal of Solid-State Circuits*, vol. 9, no. 5, pp. 256–268, Oct. 1974.

[6] S. K. Millican and K. K. Saluja, "Linear Programming Formulations for Thermal-Aware Test Scheduling of 3D-Stacked Integrated Circuits," in *2012 IEEE 21st Asian Test Symposium*, 2012, pp. 37–42.

[7] S. K. Millican and K. K. Saluja, "Formulating Optimal Test Scheduling Problem with Dynamic Voltage and Frequency Scaling," in *2013 22nd Asian Test Symposium*, 2013, pp. 165–170.

[8] K. Skadron, M. R. Stan, K. Sankaranarayanan, W. Huang, S. Velusamy, and D. Tarjan, "Temperature-aware Microarchitecture: Modeling and Implementation," *ACM Trans. Archit. Code Optim.*, vol. 1, no. 1, pp. 94–125, Mar. 2004.

[9] "HotSpot 6.0 Temperature Modeling Tool." [Online]. Available: http://lava.cs.virginia.edu/HotSpot/.

[10] Y. Zhang, L. Ling, J. Jiang, and J. Xiao, "Thermal-aware SoC Test Scheduling with Voltage/Frequency Scaling and Test Partition," *J Electron Test*, vol. 34, no. 4, pp. 447–460, Aug. 2018.

[11] S. Banerjee, S. Majumder, and B. B. Bhattacharya, "Test-Time Reduction for Power-Aware 3D-SoC," in *2018 31st International Conference on VLSI Design and 2018 17th International Conference on Embedded Systems (VLSID)*, 2018, pp. 103–108.

[12] V. Sheshadri, V. D. Agrawal, and P. Agrawal, "Power-Aware Optimization of SoC Test Schedules Using Voltage and Frequency Scaling," *J Electron Test*, vol. 33, no. 2, pp. 171–187, Apr. 2017.

[13] Y. Ko and S. Huang, "3D IC Memory BIST Controller Allocation for Test Time Minimization Under Power Constraints," in *2017 IEEE 26th Asian Test Symposium (ATS)*, 2017, pp. 260–265.

[14] H. Zhang and V. D. Agrawal, "SoC TAM Design to Minimize Test Application Time," in *2015 IEEE 24th North Atlantic Test Workshop*, 2015, pp. 55–60.

[15] L. Ling and J. Jiang, "Exploit Dynamic Voltage and Frequency Scaling for SoC Test Scheduling under Thermal Constraints," in *2014 IEEE 23rd Asian Test Symposium*, 2014, pp. 180–185.