

Disclaimer: "The contents of this document are scribe notes for The University of Texas at Austin EE382V Spring 2007, Computer Architecture: User System Interplay. The notes capture the class discussion and may contain erroneous and unverified information and comments.*

Defect Tolerance on the Teramac Custom Computer - Scribe Notes

Lecture #6: Wednesday, 07 February 2007
Lecturer: Professor Mattan Erez
Scribe: Kshama Pawar
Reviewer: Professor Mattan Erez, Min Kyu Jeong

The introductory paper on defect tolerance identifies issues hampering yield and reliability. It highlights the importance of Design for Manufacturability (DFM). The emphasis is on the issues faced by the Silicon design engineer and the steps s/he can take to incorporate DFM. The following are the scribe notes for the class discussion on the paper which talks about defect tolerance on the Teramac computer.

1 Problem being Solved

The following are the problems being solved with respect to Teramac:

- Reconfigurable high performance hardware. (Though not explicitly stated.)
- Systems should work with a variety of abstract designs - Implementation of a variety of designs.
- Reduced cost - Utilization of defective parts.
- Hiding defect tolerance from Users - Automation of mapping of working parts with designs.

The following were some of the points regarding their methodology towards solving the above problems.

*Copyright 2007 Kshama Pawar and Mattan Erez, all rights reserved. This work may be reproduced and redistributed, in whole or in part, without prior written permission, provided all copies cite the original source of the document including the names of the copyright holders and "The University of Texas at Austin EE382V Spring 2007, Computer Architecture: User System Interplay".

- Use defective parts (FPGAs) to solve building a high performance design for large systems. The Teramac system runs at a speed of 1MHz for simulations of computer architecture explorations, which was significantly higher than its competitors.
- Remapping of designs are done post defect, they are not done during runtime and the failed simulations have to be run again.
- The software for defect-analysis can be run by the user on the faulting machine, which will then rectify itself by a revised mapping and can be used again.
- The paper does not deal with fault tolerance or fault detection. One cannot know that the system is not working as expected during runtime unless a fault is encountered. Discovering a fault has not been addressed by the paper.

2 Intended Users

- SOC manufacturers/IC manufacturers possessing defective parts.
- Consumers of cheap processors.
- Manufacturers wanting to advertize cost reduction for these defective parts.
- Custom computer designers.

The following are some of the points that came up in the context of the intended users for this paper.

- Frequency binning of designs is due to process variations (explained in the introductory paper). Process variations cause out/underperformance in terms of the expected tolerances.
- Including redundancy within the design implies adding an extra row of memory or cache. Traditional defect tolerance is done by adding rows. Teramac went with adding extra interconnect to work around faulty blocks at a smaller granularity.
- Teramac approach does not scale well to manufacturing because of the time it takes to find and map, around faults.
- The Teramac custom computer is about 100 times faster than regular computers of its time according to wall clock time.

3 Uniqueness

- Diagnostic tests isolate faulty design parts using mapping software, diagnostic software and fault database.
- No redundant dedicated regular modules. An attempt to reduce the redundancy in design and increase use of as many parts as possible.
- Column or cylindrical structure for routing faster. Hierarchical interconnect system.
- Utilize Rent's rule for modules and interconnects to map a variety of designs to the hardware; ensures different designs.
- Software systems to detect the exact location of the fault.
- Efficient use of resources, they narrow down as much as possible to isolate the defective parts that they will not use.
- Other fault tolerant designs have built-in redundancy for defects. Redundancy is built into the chip and the chip will work exactly as specified. This is what Teramac did not do. Teramac chips are not logically equivalent to one another, whereas traditional defect tolerant designs would be.
- Exposes all designs on-chip and maps the exact defects. A lot of commercial designs are identical logically, that is, they do not expose redundant parts of the design outside the test equipment. Teramac exposes the hardware and does not have explicitly redundant designs which cannot be used at all once they are hardwired.

There was an elaboration of Rent's rule by Professor Erez. A researcher at IBM observed a power law between the amount of functionality, in terms of logic gates, to the number of wires in the block. Rent's rule is roughly a square root law but generically is defined as a power law. CAD tools can make use of this phenomenon.

4 Evaluation

- The study lacked comprehensive quantitative evaluation.
- There is no comparison of the provided solution with any other design. The existing comparisons are qualitative. For example, number of FPGAs has been quoted but no quantitative details are given.
- Reduction of cost is indicated by examples like three-fourths of the parts being free - A strong hint that a lot of parts were obtained free.
- There were ratios expected for the goals set at the beginning of the paper. Even the goals which were set were not numeric but only qualitative. Aimed to making a functionally correct "working system".

5 Evaluation in line with the Stated Requirements?

- There has not been a convincing argument that they met user requirements.
- Speeds of resources have not been measured/evaluated. (They have mentioned it qualitatively as “excellent”.)
- No quantitative mention of how effective their defect tolerance scheme is in terms of numbers.
- Critical defective resources are thrown away, but they do not mention how many they throw away.
- Ribbon cables is not a viable solution today for interconnection networks. Ribbons are not really high-speed interconnects, and may not work for modern signalling (not entirely sure about this by the way).
- Teramac overprovisions interconnect to utilize more logic, this is the opposite of current VLSI trends that are interconnect limited. In today’s technology, the most expensive parts of multichip modules are the interconnection networks.
- They aim to build a very large and complex design with Rent’s rule, but no size or partitioning of the design has been developed in the paper.

There was a discussion of redundancy in designs. FPGAs have regular repeating simple blocks. Each block could have redundancy. The Teramac team did not want to put in redundancy. Instead “Put in more FPGAs and throw away little in the end.” Hence, they adopted non-redundant designs with additional routing capabilities. There was however redundancy in terms of wiring. Otherwise, because of Rent’s rule, with logic redundancy and no routing capacity, blocks of designs have to be thrown away.

6 Technology a Factor?

- The authors were able to use FPGAs for their design. The assumption is that interconnection resources are cheaper than logic blocks. However, today on-chip resources are all wire bound and areas are getting smaller. Hence, not projecting to the future technologies.
- They mention interconnect hierarchy but do not elaborate on the topology or other details related to it.

7 How does it affect other Users?

- FPGAs are known to expend higher power and this design is not very power friendly. However, the main target is reliability and defect isolation.
- How does the solution of fixing the defective parts compare against not using them at all? Wear outs have a strong correlation with the number of defects, though it is not necessarily a function of the time when it was sold.
- Connectors are the least reliable parts in a system. Number of defects are crucial to how fault tolerant a system is.
- They have not suggested that it is a commercial product. The design is a large simulation environment rather than a general purpose processor.
- This is not intended as an ASIC replacement. One of the reasons could be because switch fabrics are not very viable to be used in a production design.
- The time to find a defect and fix it has to be compared against throwing away the defective part and getting a new one in much lesser time. It does not make economical sense to mark a part defective, spend money on testing it, repackaging it and selling it again. This is a solution more suitable for custom computers. It is not suitable as a commercial product in terms of the time period for fixing and remapping the design.

8 Connecting User/System Interplay

- Software and Hardware combined to solve the defect isolation and tolerance problem.
- They build a system using defective parts and pay less for it. Identify a problem and work around it. Model the system based on the restriction of system cost according to user requirements.
- They are more creative and are not merely adding custom specific parts or redundancy but do an intelligent isolation for manufacturability/defect tolerance.
- The evaluation is not thorough and systematic but more on the lines of “We built it, it works.”. Having any one kind of these two evaluations is not sufficient, but an in-between approach is required. It is not sufficient to just show that it works, but it is important to include some quantitative evaluation as well.

9 Closing Comments

- Today the process of Optical Correction involves manufacturing below the wavelength. Statistical models, CAD tools and design rules are changing and design rules are getting stricter.
- Chips with many layers and metals in between result in etching, layer collapse.
- It has become absolutely critical to pay attention to process engineering, an example - Intel has a yield of high 90% (implies what they manufacture, they sell). They pay close attention to microarchitecture and process engineering. In contrast, IBM had initial Cell processor yields at 30%. Thus, process engineering, yield and reliability directly impact the profit margin.