Defect Tolerance on the Teramac Custom Computer

Bruce Culbertson, Rick Amerson, Richard J. Carter, Philip Kuekes, Greg Snider

Year of publication: *1997*
Area: *Architecture and Technology*

Resource type	Number	Number defective	Percent defective
Programmable logic cell	221,000	23,000	10.4%
Interchip signal	145,000	13,800	9.5%
Crossbar line	4,880,000	146,000	3.0%
Crossbar buffer	2,420,000	37,000	1.5%
Total	7,670,000	220,000	2.9%

This paper was groundbreaking in its use of component-specific mapping for defect tolerance. They built a full system with individual resources (logic elements, MCM traces, and board-to-board cable connections) that failed in the few percent range (see table) and developed the tools to make it all work!

Defect tolerance in regular structures such as RAM arrays was well established in the 1980s. Altera PLDs and FPGAs extended RAM-like techniques for tolerating defects. However, they both used an approach that tried to repair the IC behind the scenes to make it look like a perfect component to the user and tools. These techniques work reasonably well when the defect rate is low, but have prohibitive costs at the level of defects suggested above.

In contrast, the Teramac design embraces the reconfigurability of the base FPGA-like substrate and exposes the defects in each particular chip, board, and connector to the placement and routing tools. With this knowledge, place-and-route can avoid the defective components during mapping, reasonably tolerating defects that are orders of magnitude higher than typical for RAMs. The resulting strategy is closer to how operating systems help tolerate defective sectors on hard drives, by mapping logical blocks to working physical sectors and avoiding the bad sectors. In addition to an over-provisioned interconnect and defect-aware-place-and-route, the paper describes novel group testing techniques to quickly identify the defects in a particular component or system and in-system delay testing to determine the speed of components.

At the time, these techniques allowed HP to aggressively use technologies that would not otherwise have been possible (large IC dies, large and complex MCMs, myriads of inexpensive ribbon cables) and build a system of a scale that was not otherwise feasible (440K 6-LUTs spread over 864 FPGAs on 8 boards). However, the real value in this work is that it provided a roadmap and demonstration that this size of system and level of defects could be conquered. The ideas developed here have been enabling for bottom up molecular logic designs, both from HP and from researchers around the world, and may be essential for all near-atomic-scale technologies including highly-scaled lithography.

André DeHon

DOI: http://dx.doi.org/10.1109/FPGA.1997.624611