

Pushing the Tools Envelope with a Rad Hard MegaChip

**Jon Ballast, AJ Kleinosowski, Bryan Buchanan, Tuan Dao,
Duncan Lam, Rod Mercer, Charles Neathery, and Michael Smith**

Solid-State Electronics Development
Boeing Research and Technology
Seattle, WA

Abstract: *This paper describes the noteworthy challenges of using a standard commercial tools flow for a highly complex, massively multi-core, radiation hardened microprocessor.*

Keywords: Maestro; radiation hardened; multi-core

Introduction

The On-board Processing Expandable Reconfigurable Architecture (OPERA) Program [1] launched in 2006 to develop next-generation processing capabilities for government space systems. At the heart of the OPERA program is the Maestro chip, a 49-core parallel processor. The Maestro chip uses the Tiler TLR26480 instruction set architecture and follows the commercial Tiler Tile64 high level design. The chip was fabricated in the IBM 9SF device technology (90nm CMOS) and was synthesized using the Boeing Radiation Hardened by Design standard cell library [2]. Maestro taped out in November 2009 using the largest 9SF die ever fabricated by IBM. This 30.6mm x 25.6mm die contained 660 million transistors. Those transistors were used in 18 million combinational gates, 3 million flip-flops, and 44 million SRAM bits. Figure 1 shows the Maestro high level design and Figure 2 shows the Maestro floorplan.

Maestro was written in Verilog, synthesized, placed, and routed using a fully commercial design automation tools flow. Table 1 lists all of the tools used in the Maestro flow. The following sections walk through the successes and challenges of using this tools flow for a massive chip which needed to meet strict radiation hardness requirements.

What Worked Well

The ICC and PrimeTime ILM flows worked fairly well with the large chip size and multitude of placeable objects. The native support for blackboxing and hierarchical processing was essential. The PrimeTime ILM flow provided a 40% decrease in memory footprint over the full-chip (non-ILM) flow. Native support for license queuing in all the Synopsys tools enabled the twelve-person design team to easily share tools and kick off multiple overnight runs in parallel. Distributed processing inherent to Calibre worked very well. The wall-clock time needed to run DRC and LVS checks decreased linearly with the number of CPUs. The Synopsys RM flows were very logical and followed the make-driven flow in ICC. PrimeTime's ECO

flow enabled easy changes to upsize gates for radiation hardness or downsize gates for power savings. The ability to accurately predict post-ECO timing in PrimeTime was critical to reducing the number of ICC ECO iterations. The 'e' language used by Specman / IUS made functional verification faster than traditional Verilog testbenches. The TetraMax debugger built into DCT (2008.09) was extensively used. The Conformal LINT checker ran without any complications and provided early identification of RTL issues before they created problems in synthesis or functional verification.

What Didn't Work

The large Maestro chip challenged the scalability of many commercial tools. Namely, the runtime for ICC was a significant bottleneck at 5-7 days for one post-synthesis to post-route run. The Virtuoso layout editor for viewing the full chip was nearly unusable. Each screen re-draw took 10 minutes. The chip-level IR drop analysis in ICC was essential for this chip, but required a 128Gb machine to run. Adding or removing signal regeneration buffers with an ECO in ICC worked well, but the buffer placement was poor and unpredictable. Buffers needed to be manually moved or given specific x/y coordinates at insertion time, requiring hours of tedious GUI work. The clock skew degraded significantly between clock tree synthesis and route optimization at the chip level, requiring several ECO iterations to re-balance the tree and close timing.

Tools Wish List

Many of the scalability challenges of the Maestro chip could be improved with enhanced blackboxing and hierarchical processing. These features were supported in some tools, but their support could be further improved. The physical design challenges for radiation hardness were met with post-processing and custom scripts. Adding another attribute for radiation hardness to the Synopsys/ICC standard cell views and then incorporating this attribute into a cost function could prevent power and area optimizations from inadvertently swapping rad-hard cells for their smaller, low-power logical equivalent.

Overall, the Maestro chip demonstrated that a fully commercial CAD tools flow can scale to massive chips and can be used for designs with unique physical design needs.

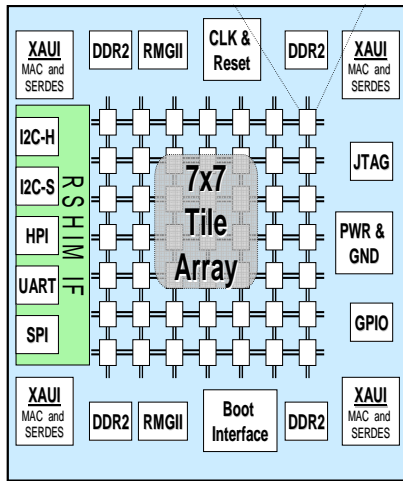


Figure 1. Maestro high level design.




Figure 2. Maestro floorplan (30.6mm x 25.6mm).

Table 1. Tools used in the Maestro flow.

CAD Tool	Version	Purpose/Usage
Cadence IUS	05.70-s019	Functional simulation (RTL)
Cadence Specman	5.1	Functional simulation (RTL)
Mentor Graphics Modelsim	6.3d (2007.11)	Functional simulation (RTL and gate-level)
Synopsys Design Compiler	A-2007.12-SP5	RTL synthesis
Synopsys DFT Compiler	A-2007.12-SP5	Scan insertion, compressor/decompressor insertion
Synopsys Power Compiler	A-2007.12-SP5	Clock-gating insertion, power estimation
Synopsys IC Compiler	A-2007.12-ICC-SP5-4	Floorplanning, placement, clock-tree synthesis, routing, DRC fixes, IR drop
Cadence Virtuoso	5.1.41	RDL Design, final chip integration, DRC fixes
Mentor Graphics MBIST	V8.2008_1.10	Memory BIST
Mentor Graphics Calibre	2007.2_34.24	LVS and DRC checking
Cadence Conformal	08.10-s480	Logical Equivalency Checking (RTL-to-gate, gate-to-gate)
Synopsys PrimeTime	C-2009.06-SP3	Static Timing Analysis
Synopsys TetraMax	2008.09-SP3	DFT DRC Checking, ATPG pattern and testbench generation
Synopsys VCS	2009.06-1	ATPG gate-level simulations
Synopsys STAR-RCXT	A-2007.12-SP1	Transistor-level parasitic extraction
Synopsys MilkyWay	A-2007.12-SP5	Standard cell library views
Synopsys HSPICE	Y-2006.03-SP1	Transistor-level functional and power simulations

Acknowledgement

This work is sponsored by the National Reconnaissance Office and The Boeing Company. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

References

1. Warren Snapp, AJ Kleinosowski, "On-board Processing Expandable Reconfigurable Architecture (OPERA) Program", Fault-Tolerant Spaceborne Computing, May 2009.
2. Warren Snapp, Tony Amort, Roger Brees, Jeremy Popp, Charles Neathery, Manuel Cabanas-Holmen, AJ Kleinosowski, Ethan Cannon, Bob Brogden, Steve Fischer, and Josh Mackler, "90nm Radiation Hard By Design SoC Demonstrations," GOMAC Tech, March 2009.