Best Practices in Design for FPGA or ASIC

Adam Sherer, Account Technical Executive, Cadence Design Systems November 1, 2021





2 © 2021 Cadence Design Systems, Inc. All rights reserved

Be Fast and Nimble with Low Latency Flexible Architectures The goal is to be out in front with both!

- Remove system latency
 - Speed RX/TX ethernet and memory subsystem but may require ASIC approach
- Add custom processor for flexibility
 - SW retargetable to better analytics or lower latency at the edge maintaining same HW
 - May need new IP and engineering resources
- Yet persistent bottlenecks can remain
 - Functionally correct but inefficiently coded
 - Efficiently coded but scalability/correctness occur



Memory Subsystem Server Communication

Coding Approaches for Low Latency

Point of Investment	Approach	Benefit
Architectural	 High-level, synthesizable design Early analysis for latency bottlenecks Effective digital twin for debug 	Detect systemic bottlenecks earlyDesign retargetable to FPGA or ASIC
Incremental	 Define coding styles for low-latency Refactor elements of design to eliminate bottlenecks Continuously analyze performance 	 Continuously reduce latency within existing architectures
Debug	 Profile design to identify bottlenecks Increase simulation and formal analysis use to resolve bottlenecks 	Achieve project confidence faster to get accelerator into market sooner

Identifying Potential Bottlenecks

- Lab: real-world data but limited visibility to internal bottlenecks
- Single simulation: high visibility into single tests
- Regression simulation: medium visibility across test suite
- Recommendation: Measure performance weekly

ᡖ Open 🖌 💿 Default View 🔻 📝 🤠 Go to 👻 🛬 🗸 Navigation History 📄 🗸				
Environment Hierarchy Instances O Types O Categories •	Code Blocks Calls			
Name Self Cumula Grade Grade □ Instances 0.0% 89.9% □ I apb_subs 2.6% 49.3% □ I apb_subs 2.6% 40.6% □ I gen 0.5% • ↓ □ 0.0% • ↓ □ • • ↓ □ • • ↓ □ • • ↓ □ • • ↓ □ • • ↓ □ • • ↓	Name All Self+ Block Kind (no filter) (no filter) (no filter) (no filter) Initial stmt (a) HDL Blocks 40.55% HDL Function (no filter) (a) HDL Blocks 2.24% Blocks (block_3 (a) Caller # # (a) block_18 Name: Initial stmt (block_23) (c) 0.09% (block_23) 0.09% 0.09% Blocks (a) block_23 0.09% 0.02% Blocks (a) block_16 0.01% 0.01% Blocks (b) block_16 0.01% 0.01% Blocks (a) block_16 0.01% 0.01% Blocks (b) block_16 0.01% 0.01% Blocks (a) block_16 0.01% 0.01% Blocks (b) block_16 0.01% 0.01% Blocks (b) block_16 0.01% 0.01% Blocks			
Children	Code Blocks Hierarchy Properties			
Name Self Cumulatin (no filter) (no filter) (no filter) I i_apb_subsyst 2.4% 46.6% I ahbi_m0 0.1% 0.1%	Source Attributes Name Self Self+ Initial stmt 0.0% 40.55% 317 end 318 319 always #5 tb_hclk = ~tb_hcl 320 321 initial begin 322 uvm_config_db#(virtual ua 324 uvm_config_db#(virtual ua 325 uvm_config_db#(virtual ua 325 uvm_config_db#(virtual ua			
Showing 2 items	X Q Name V V V V V V V V V V V V V V V V V V V			

There are Too Many Delays!

- Localize data
 - Place register files and/or memories physically close to analytics to reduce data access time
- Channelize compute
 - Replicate analytics to reduce buffering and localize interrupts
- Multiple independent clock domains
 - Reduce synchronization points to enable parallel computing

Note that increasing asynchronous compute creates risk of non-deterministic algorithmic execution

cādence

• We'll come back to this point when discussing formal verification methods

Apply Performance Best Practices

- Train engineers in common coding-for-performance approaches
- Eliminate loop invariants
 - Move constant computation ahead of loop
- Unravel loops
 - Ex: Implement as parallel computation with single summation
- Utilize incremental vs absolute calculations
 - Apply compute on just the data component used in the analytics

• Generally trade-off area for speed but keep an eye on physically long data paths

Multiple Verification Approaches Should be Applied

ation
e-metal ded SW p-cycle ster than
logic able code ehavioral a sim)
Emulation

Utilize Direct and Randomized Simulation

- Randomize packets and data loads
 - Models exchange traffic
- Direct stimulus
 - Tests specific corner cases
- Value: debug visibility and performance analysis of exchangesimilar traffic
- Challenge: incomplete analysis of overall design state-space



Formal Analysis

• Formal tests all possible stimulus, one cycle at a time

- All value combinations on inputs and undriven wires at every cycle
- All value combinations on <u>unitialized registers</u> at <u>first</u> cycle



cādence

•

Clock Domain Crossing Problem

Clock Domain Crossing (CDC) occurs when a signal crosses from one asynchronous clock • domain to another



Formal Analysis of Potential CDC Issues



Enforcing Design Practices with Linting



Comprehensive functional checks, violation debug, and waiver handling based on best-in-class formal analysis

Recent capabilities

- Deadcode debug rootcause analysis
- Improved waiver handling
- Scalability: performance and memory optimization
- FSM deadlock/livelock check performance improvements

Formal Analysis for Unreachable Coverage: Dead Code Removal



System Performance Challenges Latency Critical



System Performance Challenges **Real Time Critical**



Memory Subsystem Performance Challenges



Metric Driven Verification Created to Quantify Verification

- Early 2000s: Synthesis and packet-based design rapidly increased design size
 - Verification scaled with randomization (UVM precursor), coverage, self-checking testbench
- Metric Driven Verification (MDV) assesses how comprehensive verification is
 - All tests pass
 - All code exercised with passing tests
 - All critical state combinations are exercised
 - Illegal/undefined/unspecified states are not entered
- MDV scales for complex systems
 - Data collection from multiple verification engines



Next Steps to Being the Low-Latency Athlete

• Do your core work!

- Profile even if you don't see a performance bottleneck everyone has a need for speed
- Examine low-latency design approaches and how you can factor those into existing designs
- Maintain your performance health!
 - Increase use of simulation and formal methods to measure and correct any bottlenecks

- Break through your plateau!
 - Consider new low-latency design approaches when you implement a new architecture

Cadence is Your Design Partner

- Verification solution: apply objective analysis to improve FPGA and ASIC
 - Link to Tech Brief
- ASIC solution: broadly adopted in high-speed comms and mission-critical applications
- Tensilica IP: proven processor technology used in autonomous drive and other high-reliability apps
- High-perf IP: proven in leading comms systems
- Services: expert RTL to GDS design services

Digital Design to Implementation





cādence

© 2021 Cadence Design Systems, Inc. All rights reserved worldwide. Cadence, the Cadence logo, and the other Cadence marks found at https://www.cadence.com/go/trademarks are trademarks or registered trademarks of Cadence Design Systems, Inc. Accellera and SystemC are trademarks of Accellera Systems Initiative Inc. All Arm products are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All MIPI specifications are registered trademarks or trademarks or service marks owned by MIPI Alliance. All PCI-SIG specifications are registered trademarks or trademarks are the property of their respective owners.