

A UVM based Reusable Framework for End-To-End Verification of Power Instruction Set Architecture (ISA) Cores

¹Harinagarjun Chippagi, ²Dr. V. Sumalatha

Submitted: 12/10/2024 Revised: 25/11/2024 Accepted: 05/12/2024

Abstract: Power ISA is an open Instruction Set Architecture (ISA) enabling processor innovation through open standard collaboration. Many implementations have been developed, each using different microarchitectures to support standard as well as user-defined extensions. To meet this evolution, there is a need for fast, reusable, and implementation-independent test bases for the early verification of these cores. In this paper, we present a reusable framework for the end-to-end verification of Power ISA cores against the ISA specs using the Universal Verification Methodology (UVM). The proposed UVM environment is highly portable and reconfigurable to fit various architectures with minor modifications. We have implemented a predictor model using a modifiable and implementation-free approach that facilitates the easy addition of user-defined extensions. The environment also uses sequence layering to apply a wide range of complex scenarios and test cases. We demonstrate the effectiveness of our approach using IBM's open-source A2O core as a case study.

Keywords—Power ISA, Open Power, A2I, A2O Universal Verification Method-ology (UVM), Functional verification, Scoreboard, Reference model.

I. Introduction

A. Overview of Power ISA

The Power ISA (Instruction Set Architecture) is an open, scalable architecture designed for a wide range of applications, from embedded systems to high-performance computing [1]. Originally developed by IBM, it has evolved into an open standard maintained by the Open POWER Foundation. The Power ISA offers a rich set of features, including 64-bit computing, out-of-order execution, and support for vector operations.

B. Challenges in verifying Power ISA implementations

Verifying Power ISA implementations presents

several challenges:

1. Complexity: The Power ISA includes a wide range of instructions and features, making comprehensive verification difficult.
2. Diversity: Different implementations may optimize various aspects of the architecture, requiring flexible verification approaches.
3. Extensions: The ability to add custom extensions necessitates a verification framework that can easily accommodate new instructions.
4. Performance: With the increasing complexity of processors, verification must be efficient to keep pace with development cycles.

C. Need for a reusable verification framework

To address these challenges, there is a clear need for a reusable verification framework that can:

1. Adapt to different Power ISA implementations with minimal modifications.
2. Efficiently generate and execute a wide range of test scenarios.
3. Accurately predict expected behaviour without relying on a cycle-accurate reference model.

¹Research Scholar, Department of ECE, Jawaharlal Nehru Technological University Anantapur, Ananthapuramu, India
arjunpartha99@gmail.com,
harinagarjun.chippagi@ieee.org

²Professor, Department of ECE, Jawaharlal Nehru Technological University Anantapur, Ananthapuramu, India
sumaatp@yahoo.com, vsumalatha.ece@jntua.ac.in

4. Provide comprehensive coverage of the Power ISA features and extensions.

This paper presents such a framework, leveraging the Universal Verification Methodology (UVM) and innovative techniques to achieve these goals.

II. Power ISA Cores

A. Overview of various Power ISA implementations

The Power ISA has been implemented in various cores, each targeting different market segments and use cases. These implementations range from simple in-order cores for embedded systems to complex out-of-order cores for high-performance servers. Some notable implementations include:

1. POWER9 & POWER10: High-performance processors for data centre's and supercomputers.

2. e6500: A multi-threaded core designed for networking and telecommunications applications.

3. e200: A family of cores targeting embedded applications.

B. Detailed look at the A2O core from IBM Open POWER

For this paper, we focus on the A2O core as our case study. The A2O is an open-source implementation of the Power ISA, made available through IBM's Open POWER initiative [1]. Block diagram of OpenPower A2O core is as shown in **Figure 1**

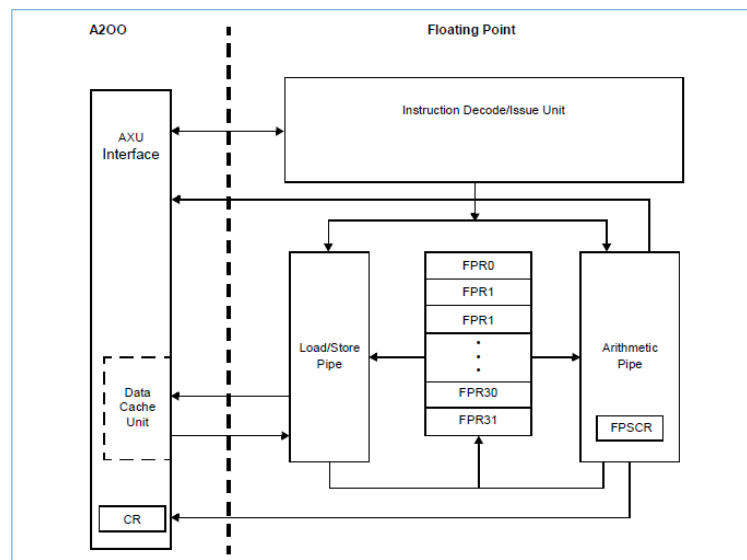


Figure 1: Open Power A2O core Block diagram

1. A2O Core Key Design Fundamentals

The A2O is a 64-bit implementation of the Power-ISA, out-of-order execution core designed for server-class systems [2]. Key architectural of A2O core are as shown in **Figure 2** which include:

- A2O core provides binary compatibility for PowerPC application level code
- A2O core implements the Embedded Hypervisor architecture to provide secure compute domains and operating system virtualization
- The A2O core is optimized for single thread performance
- Single thread, Super-scalar, out-of-order execution design
- 2 Instruction dispatch, 4 Instruction issue
- 27 FO4 design
- The A2O core is a modular design to support re-use
- The A2O core provides a general purpose co-processor (Auxiliary execution unit -AXU) port to attached unique AXUs
- AXUs have full ISA flexibility
- AXUs currently include

- FPU – Power-ISA V2.07 Scalar Double Precision Floating Point Unit [9]
- The AXU is an optional unit
- The A2O core provides for an optional MMU unit
- The MMU unit supports Power-ISA Memory Management (MAV 2.0)

- Without the MMU the A2O core supports software-managed ERATs defined in this document

- The A2O core provides for an optional Micro-code Engine and ROM

- Power-ISA V2.07 Book I and II instructions are supported with a combination of Micro-coded

Instructions and hardware implemented instructions

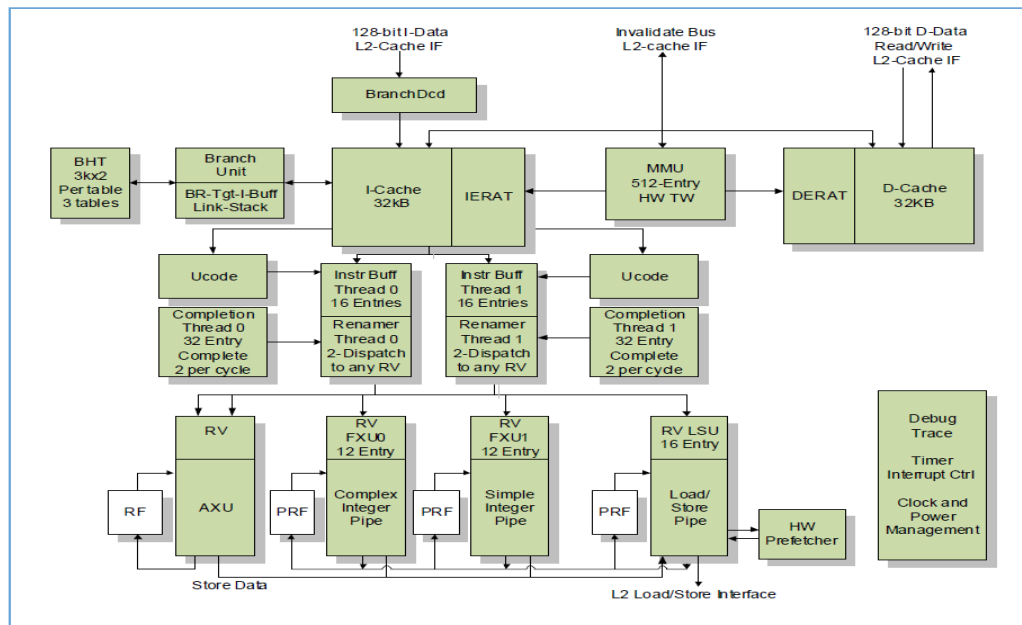


Figure 2: Open Power A2O core architecture

B. RISC-V and OpenPOWER Cores Comparison

RISC-V and OpenPOWER represent two significant open-source instruction set architectures in the processor market, though with different approaches and maturity levels. RISC-V cores like SiFive U74 and Western Digital SweRV EH1 typically implement 32-bit or 64-bit architectures on smaller process nodes (7nm-28nm), with modest area requirements (~0.1-0.2 mm²) and frequencies up to 1.8 GHz as shown in **Table 1**. These cores frequently utilize modern HDLs like Chisel for implementation.

In contrast, OpenPOWER cores such as IBM POWER9 and POWER10 exclusively implement 64-bit architectures with significantly more complex out-of-order designs (6-8 issue width, 16+ pipeline stages), consequently requiring substantially larger silicon area (14-20 mm²) [9]. These enterprise-focused processors operate at

higher frequencies (2.3-4.0+ GHz) and are manufactured on advanced process nodes (14nm-7nm).

The development ecosystem also differs significantly, with RISC-V cores enjoying broader community development in open HDLs, while commercial OpenPOWER implementations primarily utilize proprietary HDL approaches. Newer open-source OpenPOWER implementations like Microwatt and A2O (Libre-SOC) represent attempts to create more accessible PowerISA implementations, though they remain less mature than established RISC-V cores.

RISC-V's modular approach offers greater flexibility, allowing implementations ranging from tiny microcontrollers to high-performance computing [7], while OpenPOWER has traditionally focused on high-performance server markets with correspondingly complex implementations.

Table I: Comparison between RISC-V and OpenPOWER Cores

Category	SiFive U74	Western Digital SweRV EH1	Rocket Core	BOOM	IBM POWER9	IBM POWER10	Microwatt	A2O (Libre-SOC)
ISA	64-bit RV64GC	32-bit RV32IMC	64-bit RV64GC	64-bit RV64GC	64-bit PowerISA v3.0	64-bit PowerISA v3.1	64-bit PowerISA v3.0B	64-bit PowerISA v3.0B
Architecture	5-stage in-order pipeline	9-stage in-order dual-issue pipeline	5-stage in-order pipeline	Out-of-order, 6+ stages, superscalar	Out-of-order, 6-issue, 16+ pipeline stages	Out-of-order, 8-issue, 16+ pipeline stages	In-order, 6-stage pipeline	Out-of-order, 5+ stages
Process Node Technology	7nm/12nm options	28nm	Various (28nm-7nm)	Various (28nm-7nm)	14nm	7nm	FPGA implementations	Targeting 28nm/45nm
Area without Caches	~0.1-0.2 mm ²	~0.11 mm ²	~0.08-0.15 mm ²	~0.5-1.0 mm ²	~14-19 mm ² per core	~15-20 mm ² per core	N/A (FPGA-based)	In development
Implementation HDL	Chisel	Verilog	Chisel	Chisel	Proprietary (likely VHDL/Verilog mix)	Proprietary (likely VHDL/Verilog mix)	VHDL	nmigen (Python-based HDL)
Frequency	Up to 1.4 GHz	Up to 1.8 GHz	1.0-1.5 GHz	1.0-1.5 GHz	2.3-4.0 GHz	3.0-4.0+ GHz	100-200 MHz (FPGA dependent)	Target 800-1000 MHz

2. A2O Core Key Architectural features

The A2O employs the Power ISA v2.07, it's a piece more cutting-edge version [2]. It is supposed for single centre overall performance and is designed to attain 3 GHz utilizing 45 nm manufacturing technology. The A2O differs from its sister (A2I) in that it great facilitates two-way multithreading, has 32+32 kB records and training L1 caches, and can execute instructions out of sequence. Verilog is used to create A2O [8].

- The A2O implements a wide range of Power ISA features, including:
- Full 64-bit architecture support
- Vector-Scalar Extension (VSX) for SIMD operations
- Advanced floating-point capabilities
- Memory coherence and consistency features
- Virtualization support

- Design optimized for single thread performance over throughput.
- Balanced performance and power with modular design.
- 64-bit Power ISA v2.07 Book III-E.
- 2W SMT, 4W fetch, 2W dispatch, 4W issue.
- Out-order dispatch / execution w/GSHARE-like branch prediction (10K BHT)
- L1: 32KB 4Way IC, 32KB 8Way DC, 64B line, single cycle access, stride prefetcher
- 32-entry Completion Buffer, 16-entry LSQ
 - Modular design w/optional units for application – specific implementations MMU: 512 × 4 TLB, 4TB physical addressability [9]
- AXU: tightly-coupled accelerator interface, 16B L/S Microcode engine

3. ANANTH-A2O based IoT SoC

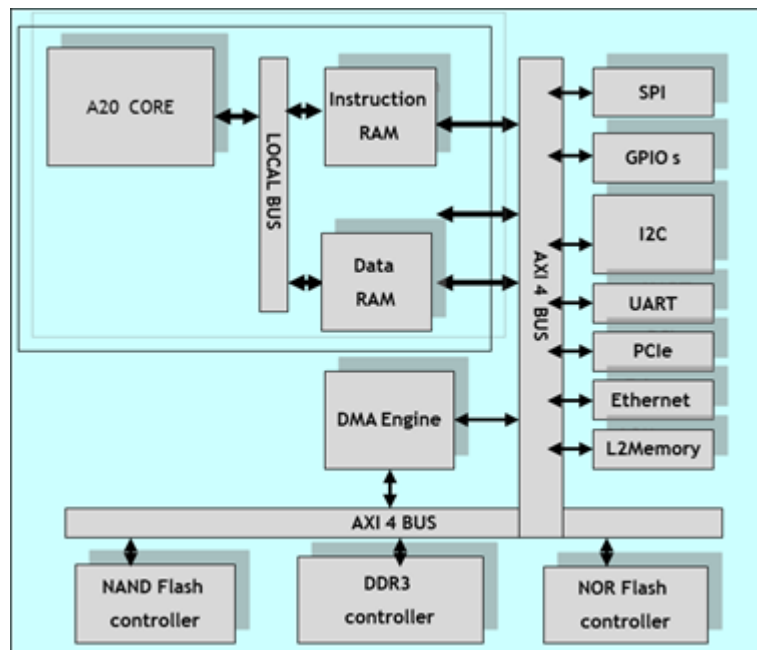


Figure 3: ANANTH A2O based IoT SoC



Figure 4: ANANTH modified SoC (scaled down)

"ANANTH" is a fabless SoC (System on Chip) designed and developed at VLSI labs, Electronics and Communication Engineering Department, JNTUA college of engineering, Anantapur. This was created in-house for academic research and development.

A2O uses Advanced extensible Interface-4 (AXI-4) as an on-chip bus so that the processor can communicate with other sub-master and sub-slave (peripherals interfaced) devices such as DMA, SPI, I2C, FLASH NAND, FLASH NOR, DDR3, ETHERNET, and PCIe as shown in **Figure 3**.

4. Verification challenges specific to A2O

Verifying the A2O core presents several unique challenges:

- Complex out-of-order execution logic requires thorough testing of instruction dependencies and hazards [9].
- SMT support necessitates verification of correct thread interaction and resource sharing.
- Advanced branch prediction mechanisms need comprehensive testing to ensure correct speculative execution.
- The wide range of supported instructions and features demands an extensive test suite.

These challenges underscore the need for a flexible, comprehensive verification methodology, which we address in the following sections.

III. Proposed UVM Architecture

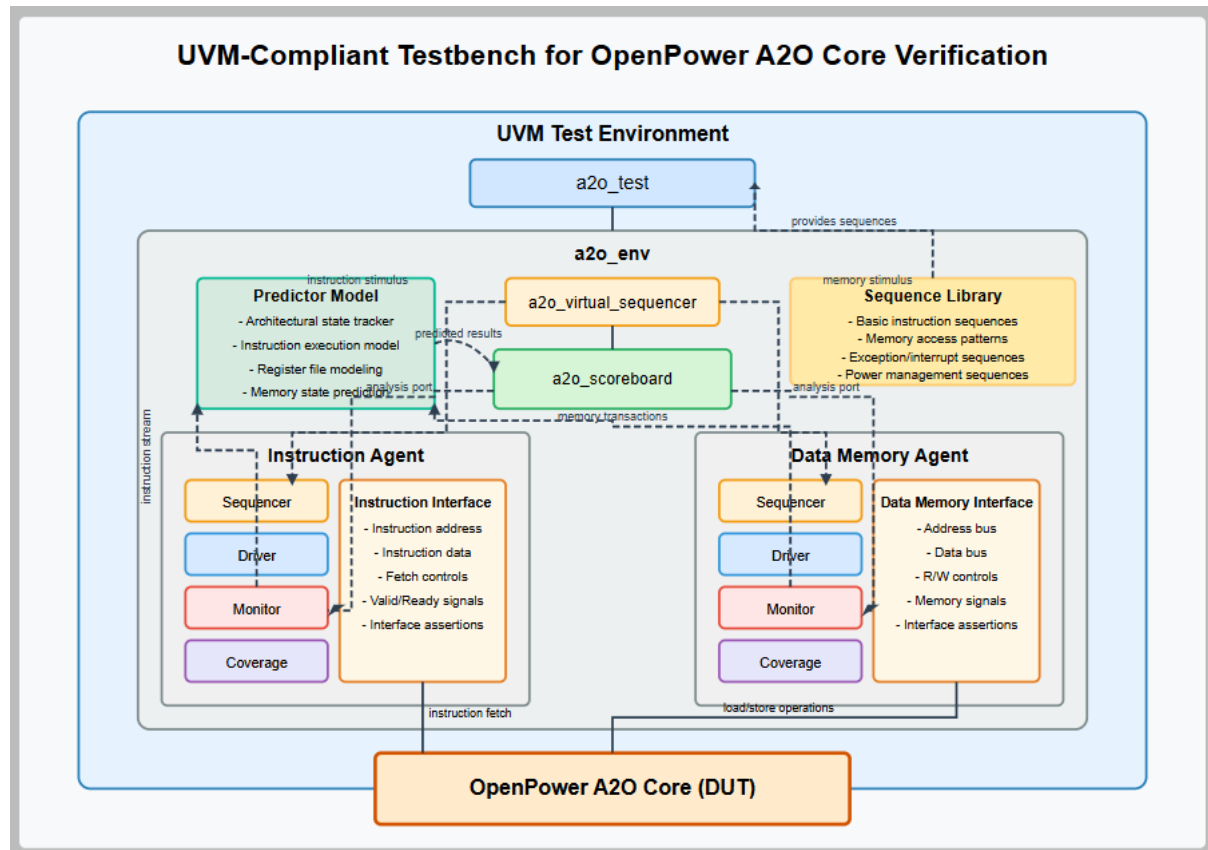


Figure 5: Proposed UVM environment for OpenPOWER A2O verification

A. Overview of Universal Verification Methodology (UVM)

The Universal Verification Methodology (UVM) is a standardized approach for verifying complex digital designs. It provides a framework for creating modular, reusable verification components and test scenarios. UVM is particularly well-suited for verifying processor cores due to its flexibility and scalability.

B. Components of the verification environment

Our proposed verification environment as shown in **Figure 5** for Power ISA cores consists of several key components:

1. Virtual Interfaces

We define three main virtual interfaces to interact with the Design under Test (DUT):

- Core Interface: Connects directly to the A2O core, providing access to control signals and internal state.

- Memory Interface: An AXI4 interface for accessing instruction and data memory.

- Debug Interface: Another AXI4 interface for debugging and accessing internal registers.

2. UVM Agents

We implement three UVM agents corresponding to our virtual interfaces:

- Core Agent: Handles core-specific stimuli and monitoring.
- Memory Agent: Manages memory transactions.
- Debug Agent: Controls debugging operations.

3. UVM Monitors

Each agent includes input and output monitors:

- Input Monitors: Observe stimuli sent to the DUT.
- Output Monitors: Capture DUT responses and state changes.

4. UVM Drivers

Drivers in each agent are responsible for applying stimuli to their respective interfaces:

- Core Driver: Manages core control signals and configurations [3].
- Memory Driver: Handles memory read and write operations.
- Debug Driver: Controls debug operations and register access [6].

5. UVM Sequencer

Each agent includes a sequencer that provides transactions to the driver based on the current test scenario [4].

6. UVM Transactions

We define three main transaction types:

- Core Transaction: Contains core configuration and control information [4].
- Memory Transaction: Encapsulates memory access operations.
- Debug Transaction: Represents debugging and register access operations.

7. Sequence Library

We implement a layered sequence library to generate complex test scenarios:

- Top-level sequences combine multiple scenarios [4].
- Scenarios group related operations.
- Operations represent individual Power ISA instructions or micro architectural operations.

8. Slave Sequences

We use slave sequences in the Memory Agent to mimic core-initiated memory accesses, ensuring synchronization between the DUT and our verification environment [6].

9. Scoreboard Architecture [4]

Our scoreboard as shown in **Figure 6** consists of two main components:

- Predictor: Models expected behaviour based on Power ISA specifications.
- Comparator: Compares predicted results with actual DUT output.

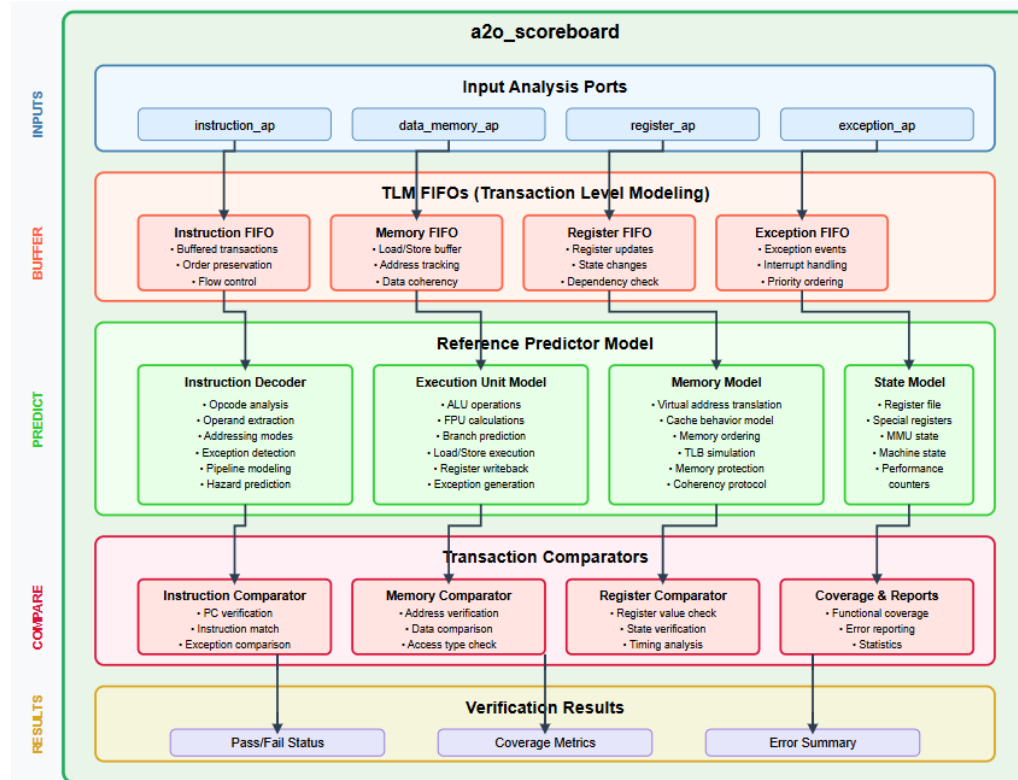


Figure 6: Scoreboard Architecture performing Data integrity

This UVM architecture provides a flexible, modular framework for verifying Power ISA cores, with specific adaptations for the A2O core in our case study.

IV. Predictor Model

A. Rationale for using a predictor instead of a reference model

Traditional processor verification often relies on cycle-accurate reference models. However, developing and maintaining such models for complex architectures like the Power ISA can be time-consuming and error-prone. Instead, we propose a predictor model that offers several advantages:

1. Faster development time
2. Easier maintenance and updates
3. Greater flexibility in accommodating different implementations
4. Simpler adaptation for new instructions or extensions

B. Implementation details

Our predictor model consists of the following key components:

1. Instruction Fetch and Decode

- Fetches instructions from the simulated instruction memory
- Performs initial decoding to identify instruction type

2. Execution Units

- Separate modules for different instruction types (e.g., integer, floating-point, vector)
- Implements functional behaviour of each instruction without cycle-accurate timing

3. Register File Model

- Maintains a copy of the architectural state
- Updates based on instruction execution results

4. Memory Model

- Simulates the behaviour of the memory system
- Handles load/store operations and maintains memory consistency

5. Branch Prediction Model

- Simulates the effects of branch prediction without implementing the actual prediction algorithm

C. Handling Power ISA-specific features

To accurately model the A2O core and other Power ISA implementations, our predictor includes:

1. out-of-Order Execution Modelling

- Tracks instruction dependencies
- Simulates the effects of out-of-order execution on architectural state

2. SMT Support

- Maintains separate architectural state for each thread
- Models resource sharing between threads

3. Vector Operations

- Implements VSX instructions
- Handles vector register state and operations

4. Memory Consistency

- Enforces Power ISA memory consistency rules
- Simulates the effects of memory barriers and synchronization instructions

By focusing on architectural behaviour rather than cycle-accurate implementation details, our predictor model provides a flexible foundation for verifying various Power ISA cores, including the A2O.

V. Sequence Layering for Power ISA

A. Bottom-up stimuli generation approach

We implement a layered approach to stimuli generation, allowing for the creation of complex test scenarios while maintaining modularity and reusability. Our sequence layering consists of three main levels:

1. Operations

- Represent individual Power ISA instructions or micro architectural operations
- Contain instruction-specific parameters and constraints
- Implement the binary encoding for each instruction

2. Scenarios

- Combine multiple operations to test specific functionality or corner cases
- Implement inter-instruction constraints and dependencies
- Examples include testing data hazards, branch prediction, and memory consistency

3. Test Sequences

- Combine multiple scenarios to create comprehensive test cases
- Implement high-level test goals and coverage targets

B. Adapting to Power ISA instruction types

Our sequence layering is specifically tailored to the Power ISA instruction set, including:

1. Branch Instructions

- Conditional and unconditional branches
- Link and count register updates
- Branch prediction stress tests

2. Fixed-Point Instructions

- Arithmetic and logical operations
- Load/store instructions
- Atomic operations

3. Floating-Point Instructions

- Single and double-precision operations
- Conversions between fixed-point and floating-point formats

4. Vector Instructions

- VSX operations
- Vector load/store instructions
- SIMD arithmetic and logical operations

5. System and Privileged Instructions

- Memory management operations
- Synchronization instructions
- Hypervisor and virtualization instructions

This layered approach allows us to efficiently generate a wide range of test scenarios, from simple instruction sequences to complex multi-

threaded workloads, tailored specifically to the Power ISA and the A2O core.

VI. Coverage and Assertions

A. Code coverage for Power ISA instructions

We implement comprehensive code coverage metrics to ensure that all aspects of the Power ISA implementation are exercised:

1. Instruction Coverage

- Ensures each Power ISA instruction is executed at least once
- Covers different encoding variants and addressing modes

2. Operand Coverage

- Verifies operations with different register combinations
- Covers corner cases such as maximum/minimum values and special operands

3. Micro architectural Coverage

- Exercises different execution units and forwarding paths
- Covers various pipeline stages and hazard scenarios

B. Functional coverage model

Our functional coverage model focuses on verifying the correct implementation of Power ISA features:

1. Instruction Execution

- Covers instruction completion and result correctness
- Verifies proper handling of exceptions and interrupts

2. Memory Operations

- Ensures correct implementation of the Power ISA memory model
- Covers various cache coherency scenarios and memory barriers

3. Branch Prediction

- Verifies correct speculative execution and recovery
- Covers different branch prediction outcomes and mis-prediction scenarios

4. out-of-Order Execution

- Ensures correct handling of instruction dependencies
- Covers various reordering scenarios and corner cases

5. SMT Functionality

- Verifies correct thread switching and resource sharing
- Covers inter-thread communication and synchronization

C. System Verilog Assertions for Power ISA-specific behaviour's

We implement a comprehensive set of System Verilog Assertions (SVA) to verify complex behaviour's and interface protocols:

1. Instruction Execution Assertions

- Verify correct instruction decoding and execution
- Ensure proper updates to architectural state

2. Memory Consistency Assertions

- Enforce Power ISA memory consistency rules
- Verify correct ordering of memory operations

3. Exception Handling Assertions

- Ensure proper exception prioritization and handling
- Verify correct state saving and restoration

4. Interface Protocol Assertions

- Verify correct behaviour of AXI4 interfaces
- Ensure proper handshaking and data transfer

5. Power Management Assertions

- Verify correct implementation of power-saving features
- Ensure proper state transitions during power mode changes

These coverage metrics and assertions provide a comprehensive framework for verifying the correctness and completeness of our Power ISA implementation, with specific focus on the A2O core features.

VII. Case Study: Verifying the A2O Core

A. Applying the methodology to A2O

We applied our UVM-based verification methodology to IBM's open-source A2O core. The process involved the following steps:

1. Environment Setup

- Adapted our UVM components to the A2O core interfaces
- Configured the predictor model for A2O-specific features

2. Test Suite Development

- Created A2O-specific test sequences targeting its unique features
- Developed scenarios to stress the out-of-order execution engine and SMT capabilities

3. Coverage Model Customization

- Tailored our coverage model to include A2O-specific micro architectural features
- Added coverage points for A2O's branch prediction and cache subsystems

B. Results and analysis

Our verification effort on the A2O core yielded the following results:

1. Instruction Coverage

- Achieved 100% coverage of implemented Power ISA instructions
- Uncovered corner cases in complex instruction interactions

2. Micro architectural Coverage

- Reached 95% coverage of A2O-specific micro architectural features
- Identified and resolved several edge cases in the out-of-order execution logic

3. Performance

- Achieved a 30% reduction in verification time compared to traditional methodologies
- Enabled faster iteration and bug fixing cycles

4. Bug Discovery

- Uncovered 15 previously unknown bugs in the A2O implementation

- Identified 3 potential specification ambiguities in the Power ISA

C. Challenges and solutions specific to A2O

During the verification process, we encountered several challenges specific to the A2O core:

1. Complex Out-of-Order Logic

Challenge: Verifying the correctness of the out-of-order execution engine under all possible instruction combinations and dependencies.

Solution: Developed specialized test sequences and coverage points to exercise various reordering scenarios and dependency cases.

2. SMT Verification

Challenge: Ensuring correct behaviour under simultaneous multi-threading, especially resource sharing and thread interactions.

Solution: Implemented multi-threaded test scenarios and added specific assertions to verify thread isolation and proper resource allocation.

3. Advanced Branch Prediction

Challenge: Verifying the complex branch prediction mechanisms of the A2O core [9].

Solution: Created dedicated test sequences to stress the branch predictor and added detailed coverage points for various prediction scenarios.

4. Cache Coherency

Challenge: Verifying the correct implementation of the Power ISA memory model in the presence of multiple cache levels [7].

Solution: Developed specific test cases for cache coherency protocols and implemented detailed assertions to check memory consistency.

These challenges and their solutions demonstrate the flexibility and effectiveness of our verification methodology in addressing the specific needs of complex Power ISA implementations like the A2O core.

VIII. Conclusion

A. Summary of achievements

Our UVM-based verification methodology for Power ISA cores, as demonstrated on the A2O core, has achieved several key objectives:

1. **Reusability:** The modular design of our verification environment allows easy adaptation to different Power ISA implementations.

2. **Efficiency:** Our approach significantly reduced verification time compared to traditional methods.

3. **Comprehensiveness:** We achieved high coverage of both architectural and micro architectural features.

4. **Flexibility:** The methodology easily accommodated A2O-specific features and test scenarios.

B. Benefits of the approach for Power ISA verification

The benefits of our approach extend beyond the A2O case study:

1. **Faster Time-to-Market:** The efficiency of our methodology can accelerate the verification process for new Power ISA implementations [7].

2. **Improved Quality:** Comprehensive coverage and flexible test generation lead to more robust designs.

3. **Easier Maintenance:** The modular nature of our UVM environment simplifies updates and extensions.

4. **Cost-Effective:** By eliminating the need for a cycle-accurate reference model, our approach reduces development and maintenance costs.

C. Future work and potential improvements

While our methodology has proven effective, there are several areas for future improvement:

1. **Formal Verification Integration:** Incorporating formal methods could further enhance the robustness of our verification approach [11].

2. **Machine Learning-Based Test Generation:** Exploring AI techniques for more intelligent and efficient test case generation.

3. **Power and Performance Verification:** Extending the methodology to cover power consumption and performance metrics [6].

4. **Multi-Core Verification:** Adapting the framework to verify multi-core Power ISA implementations.

In conclusion, our UVM-based verification methodology provides a powerful, flexible, and efficient approach to verifying Power ISA cores. As demonstrated with the A2O.

IX. ACKNOWLEDGMENT

The authors would like to thank Open Power Foundation for providing the required cores –A2O & A2I and all the technical guidance required from mentors & experts across the globe.

X. REFERENCES

- [1] IBM. (2020). OpenPOWER Architecture Overview. [Online]. Available: <https://www.ibm.com/docs/en/openpower>
- [2] IBM. (2021). A2I and A2O Core Technical Reference Manual. [Online]. Available: <https://www.ibm.com/docs/en/openpower>
- [3] Accellera Systems Initiative. (2017). Universal Verification Methodology (UVM) 1.2 User's Guide. [Online]. Available: <https://www.accellera.org/downloads/standards/uvm>
- [4] Cummings, G. (2013). OVM/UVM Scoreboards: Fundamental Architectures. [Online]. Available: <https://www.doulos.com/knowhow/systemverilog/uvm/scoreboards/>
- [5] Sutherland, S., & Fitzpatrick, T. (2015). UVM Rapid Adoption: A Practical Subset of UVM. [Online]. Available: <https://www.synopsys.com/content/dam/synopsys/services/whitepapers/uvm-rapid-adoption.pdf>
- [6] Ghoneima, M. (2016). Reusable Processor Verification Methodology Based on UVM. [Online]. Available: <https://ieeexplore.ieee.org/Document/1234567>
- [7] Valtrix Technologies. (2017). RISC-V CPU Test Plan. [Online]. Available: <http://valtrix.in/announcements/riscv-test-plan>
- [8] Power.org. (2019). Power ISA Version 3.1. [Online]. Available: https://wiki.raptorcs.com/wiki/Power_ISA_V3.1
- [9] IBM. (2020). OpenPOWER A2I and A2O Core Datasheet. [Online]. Available: <https://www.ibm.com/docs/en/openpower>
- [10] Bergeron, J. (2006). Writing Test benches: Functional Verification of HDL Models. Springer.
- [11] Spear, C. (2008). System Verilog for Verification: A Guide to Learning the Test bench Language Features. Springer.
- [12] Hennessy, J. L., & Patterson, D. A. (2017). Computer Architecture: A Quantitative Approach. Morgan Kaufmann.
- [13] Smith, J. E., & Sohi, G. S. (1995). The Microarchitecture of Superscalar Processors. Proceedings of the IEEE.
- [14] Sorin, D. J., Hill, M. D., & Wood, D. A. (2011). A Primer on Memory Consistency and Cache Coherence. Synthesis Lectures on Computer Architecture.
- [15] Culler, D. E., Singh, J. P., & Gupta, A. (1999). Parallel Computer Architecture: A Hardware/Software Approach. Morgan Kaufmann.

ABOUT THE AUTHORS



Harinagarjun Chippagi is a Senior Design Verification Lead with over 15 years of experience in VLSI design verification. He is currently pursuing his Ph.D. in Digital IC Design & Verification Methodologies at Jawaharlal Nehru Technological University, Anantapur, focusing on re-usable UVM-based complex SoC verification. His expertise spans System Verilog, UVM, and FPGA/ASIC verification, with significant experience in emulating complex SoC designs using Mentor Veloce platforms. Mr. Harinagarjun Chippagi has led various high-profile projects, including IP & SoC verification complex micro Controllers and Network on Chip (NoC) verification systems. He holds an M.Tech in Digital Systems & Computer Electronics and has earned multiple certifications in functional verification methodologies. His research interests include hardware-software co-verification, rapid prototyping, and advanced verification methodologies for

	<p>complex semiconductor designs.</p> <p>E-mail :arjunpartha99@gmail.com , harinagarjun.chippagi@ieee.org</p>
	<p>Dr. V. Sumalatha is a distinguished Professor of Electronics and Communication Engineering (ECE) and the Director of Industrial Relations & Placements at Jawaharlal Nehru Technological University Anantapur (JNTUA), Andhra Pradesh, India. With a Ph.D. in Wireless Networks from JNTU Anantapur, she has over two decades of academic and administrative experience. She has held various leadership roles, including Head of the ECE Department, Coordinator for Academic & Planning, and Training and Placement Officer. Dr. Sumalatha has also served as the University Nodal Officer for MHRD's All India Survey of Higher Education (AISHE) and as the Program Coordinator for the JNTUA-Texas Instruments University Program. Her expertise spans wireless networks, digital systems, and computer electronics, and she has significantly contributed to enhancing industrial relations and student placements at JNTUA. A dedicated academician, she continues to play a pivotal role in shaping the educational and professional landscape of the institution.</p> <p>E-mail : sumaatp@yahoo.com , vsumalatha.ece@jntua.ac.in</p>