

# The Evolution of RISC-V Processor Verification: Open Standards and Verification IP

Lee Moore, Imperas Software, [moore@imperas.com](mailto:moore@imperas.com)  
Aimee Sutton, Imperas Software, [aimees@imperas.com](mailto:aimees@imperas.com)  
Mike Thompson, OpenHW Group, [mike@openhwgroup.org](mailto:mike@openhwgroup.org)

## ABSTRACT

The OpenHW Group's [1] Verification task group has been a pioneer in the development of methodologies and verification collateral for RISC-V processor verification. Since 2019 the members have worked together to develop CORE-V-VERIF: a UVM environment for the verification of RISC-V processor cores. Over this period of time the CORE-V-VERIF environment has evolved as new processor verification projects introduced new challenges, and learnings from the previous projects led to the development of new approaches. With each generation the CORE-V-VERIF environment has improved to become more robust, more reusable, and ultimately better at finding RTL bugs. The current generation uses RISC-V processor verification IP enabled by the open standard RISC-V Verification Interface (RVVI) [2] to realize a comprehensive verification methodology that encompasses asynchronous peripheral events that occur randomly during program execution. This paper will describe the evolution of RISC-V processor verification methodology using CORE-V-VERIF as a case study. Readers will learn a proven approach to RISC-V processor verification that can be accessed through an open-source example. Readers who are already familiar with CORE-V-VERIF may choose to skip ahead to the description of the third-generation verification environment to learn about the latest developments.

## INTRODUCTION

The popularity of the RISC-V instruction set architecture (ISA) for microprocessors is growing, and with it is the need for, and interest in, processor-specific design verification techniques. One of the main contributors to the evolution of RISC-V processor DV is the OpenHW Group's Verification Task Group. The focus of this group is to ensure that the CORE-V family of open-source RISC-V processor cores is verified to industry standards, so that anyone who adopts these cores can feel confident that they are ready for silicon. In addition, for any developer that modifies or extends the open-source core IP, availability of the test bench infrastructure provides the basis for verification of the new enhancements. While industry-standard verification techniques such as constrained-random generation and functional coverage closure were applied, the verification task group has been a pioneer in developing RISC-V specific verification methodologies and best practices, and in turn making these available to the wider community.

The OpenHW Group and its member companies support an on-going effort to develop and maintain a portfolio of RISC-V compliant cores collectively known as the CORE-V family [3]. The verification environment used to verify the CORE-V cores is known as CORE-V-VERIF [4] and it too is an open-source artifact, available on GitHub. The reader is invited to review the CORE-V-VERIF Quick Start Guide [5] for information about running simulations on a CORE-V core under the CORE-V-VERIF environment.

The first and second generations of OpenHW Group's CORE-V-VERIF utilized a C reference model of the processor embedded in a UVM testbench. The reference model and the Device Under Test (DUT) were run in lock-step, each executing the same program, while the testbench compared the internal state of the two at the retirement of every instruction. This became known as the step-and-compare methodology. Used together with a random instruction stream generator and UVM agents providing peripheral stimulus (such as interrupts), it was an effective method of achieving verification closure for the CV32E40P RISC-V processor RTL.

Although successful, the time and human effort required to develop the first and second generation CORE-V-VERIF environments and achieve verification closure was high. The OpenHW Group was developing ambitious plans to achieve the same level of verification for at least five additional RISC-V cores, and it was clear that a more efficient methodology was required.

The experience with these verification projects led to the observation that some components should be common to all RISC-V processor testbenches, and common components should conform to standard interfaces. This led to the development of RVVI: the RISC-V Verification Interface [2]. In turn, the availability of standard interfaces made it possible to develop verification IP that implements these common components. These innovations have contributed to the third generation of the CORE-V-VERIF environment which is in use today with multiple RISC-V core developments.

The following sections will provide greater detail into the evolution of the CORE-V-VERIF environment and RISC-V verification methodology. The first and second generation CORE-V-VERIF environments will be described, along with the highlights and shortcomings of each. Next the RISC-V Verification Interface (RVVI) will be described in detail, with an explanation of how it starts to address the shortcomings mentioned previously. The evolution continues with the introduction of verification IP to promote reuse, improve checking, and gain further efficiency. Finally, the present day third generation CORE-V-VERIF environment will be shown to demonstrate the successful application of these techniques and provide context for future work.

### THE FIRST GENERATION CORE-V-VERIF ENVIRONMENT

The OpenHW Group develops open-source design and verification IP centered on the CORE-V family of RISC-V cores and related IP. The verification environments for the CORE-V cores are maintained as open-source artifacts in the [CORE-V-VERIF GitHub repository](#). The first-generation CORE-V-VERIF environment is illustrated in Figure 1 below.

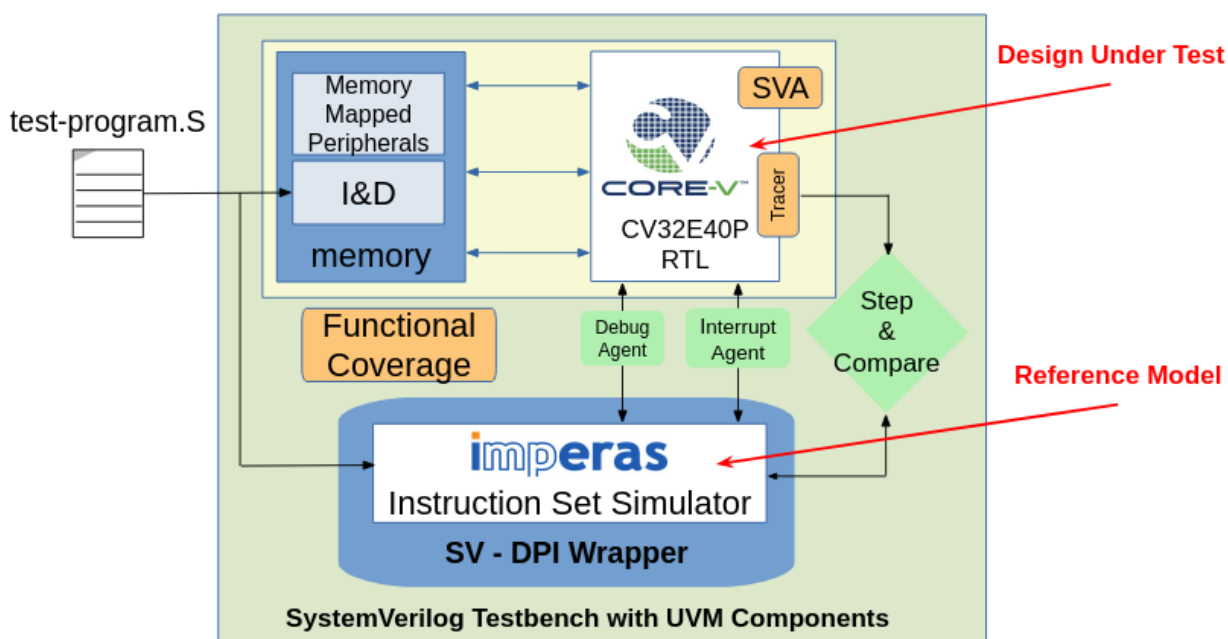


Figure 1: First Generation CORE-V-VERIF environment

## Reference Models in Verification

As illustrated in Figure 1, a common strategy in simulation-based verification is to integrate the device-under-test (DUT) and a reference model in a testbench and compare the state of the DUT and RM at specific times during the simulation. In processor core verification the state variables of interest will typically be the program counter (PC), general purpose registers (GPRs) plus the Control and Status Registers (CSRs). The most natural time to compare the state of the reference model and DUT is when an instruction has been *retired*. That is, execution of an instruction has completed and the values of the PC, GPRs and CSRs have been updated to reflect the effects of the instruction.

Conceptually, the theory of operation of such a testbench is as follows:

- Both the DUT and reference model execute the same program. Note that in the CORE-V-VERIF testbench, these programs can be manually written by a human or generated by a random instruction stream generator (not shown in Figure 1).
- Non-programmatic interrupts to normal program flow such as external interrupts and debug requests are sent to both the DUT and reference model. (Note: these “non-programmatic interrupts to normal program flow are termed “asynchronous events” in the remainder of this document.)
- As each instruction retires, the state variables of the DUT and reference model are compared. Mismatches are reported as errors. This method is known as step-and-compare.

In a typical processor core, the state of the PC, GPRs and CSRs are not accessible as top-level I/O ports. Additionally, due to the implementation of the processor’s internal pipeline, the values of these state variables may be updated at differing clock cycles as instructions are executed. Exposing the state of the DUT to the testbench is the job of a module called a **Tracer**<sup>1</sup>. An important consideration for a Tracer is ensuring that all aspects of the core’s state are presented in a way that allows the consumer of Tracer outputs to attribute the state to a specific instruction.

In the first generation testbench, the Tracer used was an ad-hoc behavioral model bound to the RTL exposing all processor states such as GPR, PC and CSR values. This Tracer lacked a documented interface and required frequent changes in both the Tracer itself and the Step-and-Compare testbench logic.

The design and implementation of the reference model is a key consideration for this testbench architecture. An Instruction Set Simulator (ISS) is often used as a reference model. Often developed and/or used by software development teams, an ISS can be custom-built in-house, a commercial product, or available as an open-source artifact. The CORE-V-VERIF environment used Imperas OVPsim[6,7] as the reference model for the embedded class cores. Using an ISS significantly reduces the development effort of the RTL testbench. This was the strategy employed by the first generation CORE-V-VERIF environment. However, as we will see, it came with significant challenges, and the remainder of this paper will discuss these and introduce strategies for addressing them.

The most obvious problem encountered when using an ISS as a reference model is caused by the different abstraction levels of the ISS and DUT. The majority of DUTs are simulated at the Register Transfer Level, which by definition is a clock-cycle accurate description of hardware logic in which time is modeled as clock events. This is in contrast to a transaction-level model in which time is modeled as transaction cycles. In our case, an ISS of a processor core models time in terms of instructions. While an RTL model of a core may take multiple clock cycles to execute an instruction, depending on instruction fetch bus timing, the specific instruction being executed or the state of the internal pipeline, an ISS will usually execute all instructions in a single instruction cycle because the cycle timing of the physical interfaces to memory and the operation of the core’s pipeline are abstracted away.

The abstraction inherent in an ISS is a powerful modeling technique, and it is highly effective in producing predictions of the core’s state when executing a series of instructions which are not subject to asynchronous events such as interrupts. Nevertheless, a method must be found to maintain synchronization between the ISS, which operates in terms of instruction cycles, and the DUT, which operates in terms of clock cycles. The method used by CORE-V-VERIF to maintain this synchronization is called “step-and-compare”.

---

<sup>1</sup> Unrelated to the Efficient Trace for RISC-V specification (<https://github.com/riscv-non-isa/riscv-trace-spec/releases>)

A second problem when using an ISS as a reference model is the timing of “side effects”. An instruction is said to have side effects if it updates one or more state variables which are not explicitly part of the instruction. For example, the CSR minstret is updated each time an instruction is retired. In most cases, side effects are easy to predict. In other cases, particularly in the event of an asynchronous interrupt, the timing of side effects in the ISS may differ from the RTL.

### *Step-and-Compare Architecture*

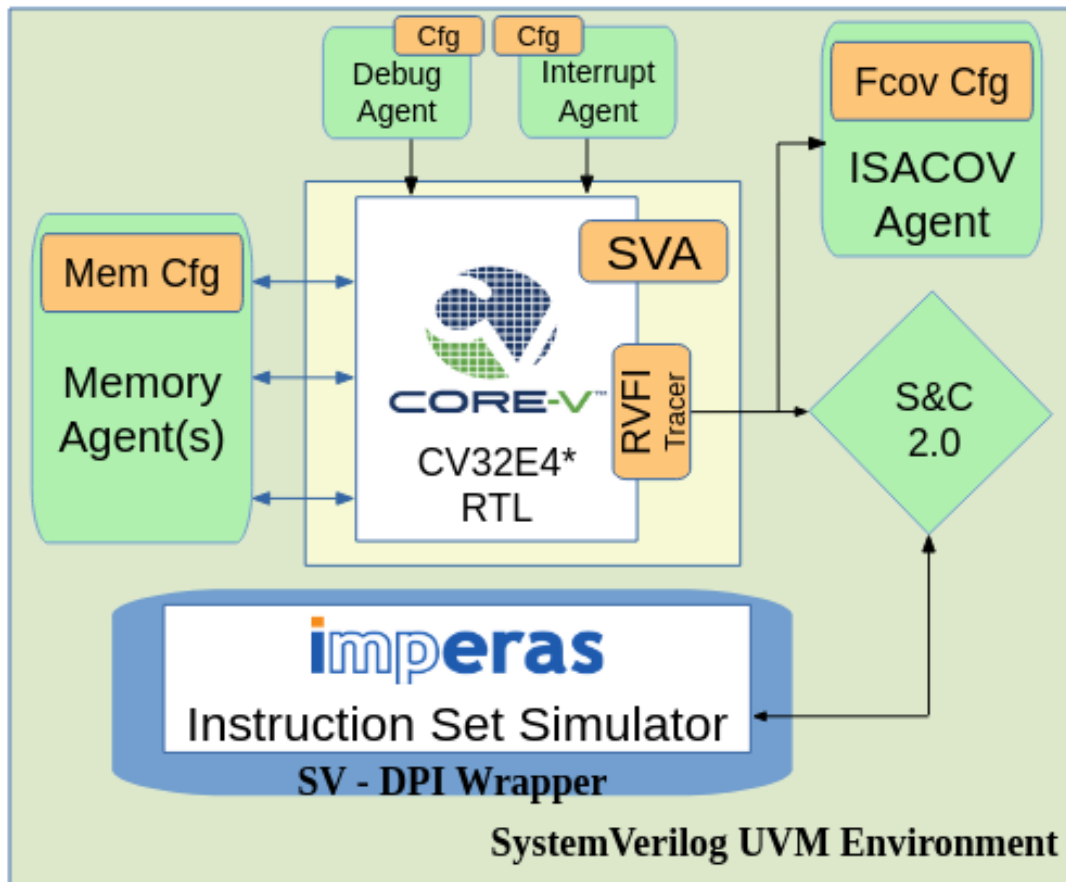
[8] documents the original step-and-compare method used in the first generation of CORE-V-VERIF testbenches. An ISS is used as a reference model and is kept in sync with the RTL by running the RTL clock until an instruction is retired, at which point the RTL clock is stopped (“gapped”) and the reference model is allowed to run until it retires an instruction. The testbench then compares the predicted processor state (from the reference model) to the actual processor state (from the RTL). This process repeats until the last instruction is retired and compared.

This technique works very well in the absence of external asynchronous events such as debug requests, interrupts, memory access delays and memory bus errors. In the first generation of CORE-V-VERIF, the ISS and RTL were connected to the same interrupt and debug request inputs. Due to the clocking system enforced by the step-and-compare logic, the ISS would always “see” these external events as being synchronous to the start of an instruction. Due to the internal micro-architecture of the RTL, these external events may or may not be “seen” as occurring before or after the start of the same instruction. This would inevitably result in differences in the predicted and actual processor states and a false-negative comparison would result.

To make matters worse, each asynchronous event needed to be handled as a unique event with core-specific code in the testbench to properly handle the behavior and side effects of these asynchronous events. This made the testbench “buggy” (prone to false-negatives) and difficult to maintain.

In response to this, a second generation of step-and-compare was developed to build upon and fix many issues with the first generation while maintaining the same verification effectiveness. For the purposes of this discussion the most purposeful improvements were to formalize the Tracer interface definition.

## THE SECOND GENERATION CORE-V-VERIF ENVIRONMENT



**Figure 2: Second Generation CORE-V-VERIF UVM Environment**

Figure 2 above, illustrates the changes implemented to create the second generation UVM environment:

- The ad-hoc Tracer used in the first generation is replaced by a Tracer conforming to an interface extended from the RVFI [9] specification. A new, simplified, second generation “step-and-compare 2.0” is developed.
- Asynchronous events such as interrupts and debug requests are connected to the core RTL (device under test), but not the ISS.

Each point from the above list is significant and will be discussed further below.

### *CORE-V Tracer using the RVFI + extensions*

In the first generation testbench, monitoring of processor activity was enabled by a specific-purpose Tracer. The Tracer was bound to the RTL exposing all processor states such as GPR, PC and CSR values. This Tracer was difficult to maintain and implemented with an ad-hoc interface, requiring customized frequent changes in the Step-and-Compare implementation.

To address the issues with the ad-hoc tracer, the RISC-V Formal Interface (RVFI) specification was adapted to define the requirements for the processor Tracer. As much as was practical the RVFI was followed verbatim, with updates for extra requirements introduced by its usage in Step-and-Compare. The details of the specific CORE-V RVFI implementation extensions used is captured in the User Manual [10] of the CV32E40X, an RV32 processor core targeting embedded applications.

The CORE-V Tracer probes the internal design of the core to perform the following functions:

- Unambiguously identify instruction retirement.
- Unambiguous reporting of both synchronous traps and asynchronous exceptions, interrupts and debug requests.
- Expose the state of the core (PC, GPRs, CSRs).

Depending on the complexity of the core’s pipeline, the logic required to implement the CORE-V Tracer ranges from trivial to complex. To simplify the CORE-V Tracer, it is typically implemented as a behavioral module that is bound (using SystemVerilog *bind*) into the RTL model. This simplifies the implementation because the coding is not constrained by synthesis requirements and can be maintained as one or more source files outside of the RTL sources.

A significant benefit of having a tracer interface specification is that it provides a well-defined blueprint of the requirements of a complete Tracer. This is more impactful than it at first appears. It can be difficult to fully specify the requirements of a Tracer apriori, without first having the experience of using a Tracer that does not fully fit your needs. It is often the case that such requirements can only be fully understood after several cycles of trial-and-error. This experience is what eventually led to the new open-standard RVVI specification.

#### *Handling Asynchronous Events with CORE-V Second Generation Tracer*

In the first generation of CORE-V-VERIF, the reference model “saw” asynchronous events such as interrupts and debug requests simultaneously with the RTL model. Due to differences in timing, this could mean that the RTL and reference model would “take” the interrupt on different instructions. This placed a burden on the step-and-compare logic in the testbench to ensure that the reference model would take the interrupt on the same instruction as the RTL.

In the second generation of CORE-V-VERIF, the reference model is not connected to asynchronous events. Thus, on its own, the reference model cannot determine when interrupts to normal program flow occur, such as external interrupts or external debug requests. The Tracer monitors and reports those events, and thus the “Step-and-Compare 2.0” logic can be used to *inform* the reference model to interrupt normal program flow, maintaining processor state lock-step with the DUT. This is possible because the tracer interface is explicitly defined to indicate when this information is presented to the reference model and all of the information required by the reference model is provided at the time the RTL retires an instruction.

For an example, consider what happens when a debug request is asserted. The CORE-V tracer interface defines two signals *rvfi\_dbg* and *rvfi\_dbg\_mode* that are valid when an instruction is retired. Together these two signals indicate whether the core executed the retired instruction in debug mode and for the first instruction after entering debug, *rvfi\_dbg* contains the debug cause. In addition to providing this information, the CORE-V Tracer specifies clear rules for how debug entry is recognized:

*Debug entry is seen by the tracer interface as happening between instructions. This means that neither the last instruction before debug entry nor the first instruction of the debug handler will signal any direct side-effects. The first instruction of the handler will however show the resulting state caused by these side-effects (e.g. the CSR rmask/rdata signals will show the updated values, pc\_rdata will be at the debug handler address, etc.).*

The CORE-V tracer interface has similarly comprehensive and rigorous definitions for how an interrupt is signaled and how side effects are modeled. This information greatly simplified the step-and-compare logic required to keep the RTL and ISS state in sync.

However, this methodology left a serious verification hole: the reference model was not able to provide independent verification of the DUT’s response to asynchronous events. For example, in a case where multiple interrupts are enabled and pending, it was unable to verify that the correct one (by priority) was taken. It simply mirrored the actions of the DUT. Checking this type of behaviour was left to other testbench components. Given the difficulty of validating responses to randomly generated asynchronous events, this checking was often incomplete. This serious deficiency was addressed in the third generation CORE-V-VERIF environment.

## *RVVI: The RISC-V Verification Interface*

The work of the OpenHW Group VTG on the first and second generation CORE-V-VERIF environments led to the observation that certain components should be common to all RISC-V processor verification environments, and common components should be accessible through standard interfaces. These interfaces have now been formalized in RVVI [2], an open and evolving standard for functional verification of RISC-V processors. Two major components of RVVI will be discussed below: RVVI-TRACE and RVVI-API.

### *RVVI-TRACE*

A common component that benefits from a standard interface is the previously-discussed tracer. Every RISC-V processor under test needs a tracer module in order to extract internal state information required for effective verification. While the implementation of the tracer is specific to each processor's microarchitecture, the requirements for the information that a tracer should provide are common, and are defined by the needs of the testbench. These requirements led to the creation of the RVVI-TRACE specification. RVVI-TRACE is specified as a SystemVerilog interface that connects the processor under test with the testbench.

When determining what the standard interface for tracers should look like it was natural to study formal interfaces being used in processor verification and see if they were appropriate for dynamic (simulation based) verification. It was clear that RVFI did not meet the needs of functional verification, as it was designed for formal verification, and heavily extended and modified for CORE-V-VERIF. The CORE-V extensions to it were insufficient for a standard interface as they were specific to the requirements of the CV32E4\* cores and did not anticipate the needs of the full set of RISC-V ISA variations. For example, RVFI did not have a method of signalling more than one register change (side effect) per instruction retirement, something which does happen in the RISC-V Zc extension. Today the RVVI-TRACE interface addresses this and other issues with RVFI, and can be seen as a natural evolution and extension into dynamic verification of RVFI.

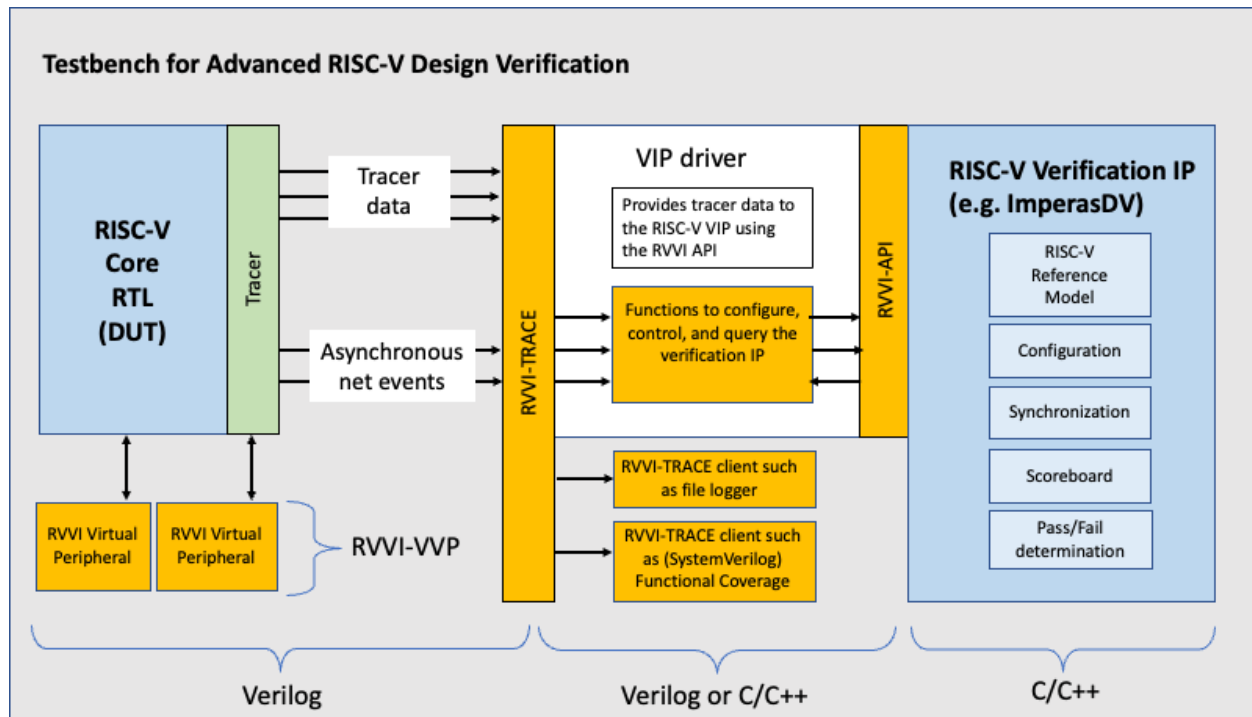
Another key feature of RVVI-TRACE is the mechanism for handling asynchronous inputs to the processor. RVVI-TRACE contains a SystemVerilog queue that is used to store multiple net changes that occur during the interval between instruction retirements, as well as when those changes occurred. This information is now available to the reference model or checker for validation of the processor's response to these events.

For the CORE-V-VERIF environment, the benefits of adopting RVVI-TRACE further extend the benefits realized by the development of the initial CORE-V tracer interface. There is still the existing benefit of knowing upfront what are the requirements for an effective tracer and being able to clearly communicate these to the design team. There is the additional benefit of having an interface that supports the full suite of RISC-V ISA subsets so there is no need to modify it between different processor projects. This, in turn, permits reuse of any component that is a client of the RVVI-TRACE interface.

### *RVVI-API*

Another component required for comprehensive RISC-V processor verification is a behavioural reference model of the processor. The reference model provides an independent representation of the processor's internal state. It is subject to the same configuration and initial conditions and executes the same program as the DUT. The model should have the ability to run in lock-step with the DUT so that the two states can be continuously compared and bugs can be identified at the time they occur. In addition to the reference model, a processor verification environment needs a component to perform the comparisons between the model and the DUT, to keep track of any mismatches in state. This set of requirements led to the development of RVVI-API: a set of functions that must be implemented in the processor verification IP and supporting testbench components in order to comprehensively check the behaviour of the RISC-V processor under test.

The diagram in Figure 3 below illustrates a canonical RISC-V processor verification environment using RVVI and RVVI-compliant processor verification IP.



**Figure 3: Testbench for Advanced RISC-V processor verification using RVVI**

The introduction of the RVVI-API into CORE-V-VERIF has addressed some of the shortcomings and issues encountered with the previous step-and-compare environments. One of these areas is the configuration of the RISC-V verification IP and processor reference model. RVVI-API specifies functions (Figure 4) to configure specific memory regions, registers, or register fields as volatile.

Volatile control and status registers (CSRs) and memory regions are those that change asynchronously to the program execution. They often require a cycle-accurate representation of the processor to model accurately. An example of a volatile CSR is a counter that increments every clock cycle. An example of a volatile memory location is the address space used by a memory-mapped peripheral. It is important to note that the reference models used for RISC-V processor verification are instruction-accurate, and not cycle-accurate. These are architectural models that model processor behaviour but do not model microarchitectural details such as the processor's pipeline.

In the first and second generation CORE-V-VERIF environments the verification IP and processor reference model had no notion of volatility. It was up to the step-and-compare logic to maintain a list of register comparisons to discard and to populate the correct register values in the reference model. This was necessary to ensure that the program running on the model and the processor core would exhibit the same behaviour.

To address this problem, RVVI-API specifies functions to mark a register, register field, or memory region as volatile. It is now the responsibility of the reference model to keep track of volatile addresses and to ensure that the contents of these regions stay consistent with the DUT. RVVI-API also specifies functions to inform the reference model (and/or the VIP that encapsulates it) about processor read/write activity. When a volatile region is accessed the data from that location can be propagated to the reference model's memory. This ensures that the test program will run as expected.

```
import "DPI-C" function int rvviRefCsrSetVolatile(
    input int hartId,
```



```

input int csrIndex);

import "DPI-C" function int rvviRefMemorySetVolatile(
    input longint addressLow,
    input longint addressHigh);

```

Figure 4: RVVI-API volatile functions

### THE THIRD GENERATION CORE-V-VERIF ENVIRONMENT

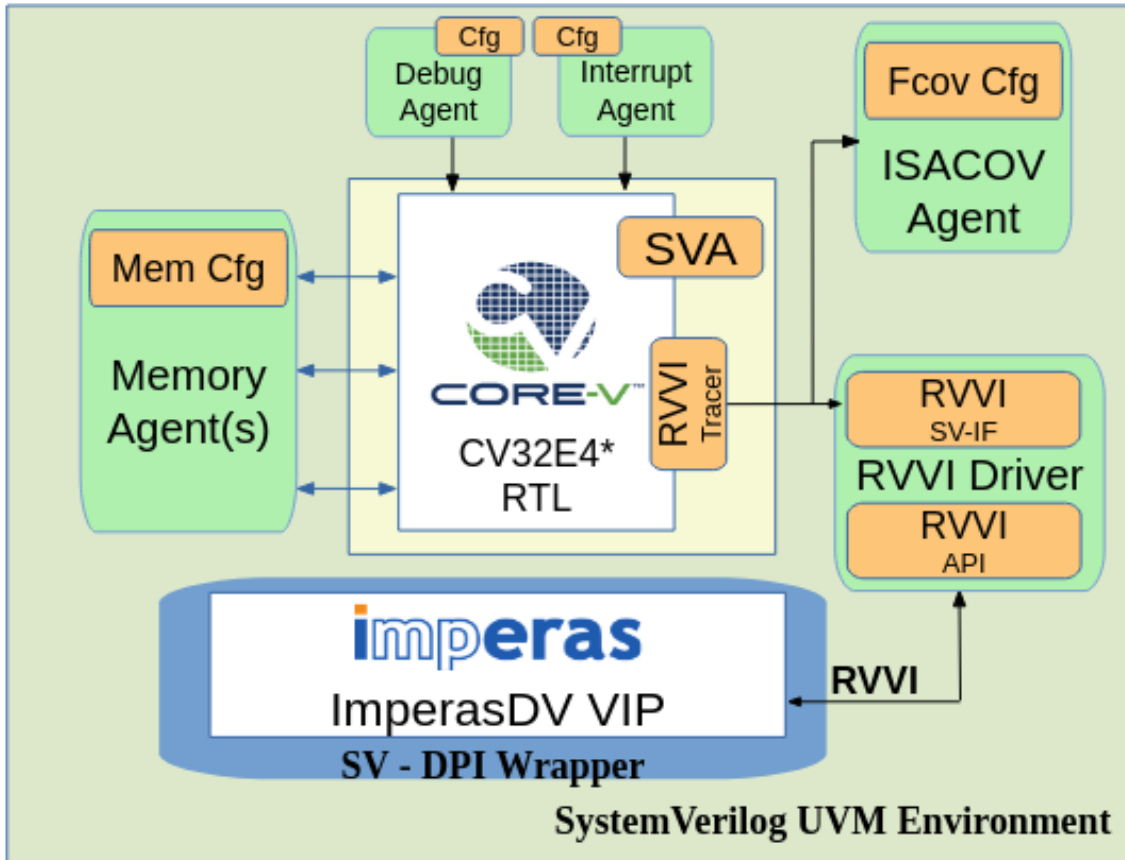


Figure 5: The third generation CORE-V-VERIF environment

Figure 5, above, illustrates the changes implemented to create the third and current generation UVM environment:

- The CORE-V tracer has been replaced with an RVVI-compliant tracer
- The Imperas Reference model has been replaced with ImperasDV verification IP (VIP) that incorporates the reference model
- The step and compare logic has been eliminated

The benefits of adopting an RVVI-compliant tracer have been discussed in previous sections. At the time of writing, the migration from CORE-V's use of RVFI with extensions to RVVI is an ongoing activity.

The most impactful change has been the replacement of the reference model and the step-and-compare logic with a piece of verification IP. The ImperasDV VIP encapsulates a reference model of the target processor, and implements

the functions specified by RVVI-API. It performs the internal state comparisons between the reference model and the DUT using information from the RVVI-TRACE interface, and keeps track of those results in an internal scoreboard. This eliminates the need for the error-prone and complex step-and-compare logic being hand coded in the testbench. Since asynchronous events are now communicated to the VIP using RVVI-TRACE the processor's response to these can now be independently validated.

The following section contains an explanation and example of how an architectural reference model can be utilized to provide validation of a processor's handling of asynchronous inputs such as interrupts and halt requests.

### Handling Asynchronous Events with ImperasDV

As previously discussed, one of the most challenging tasks in processor verification is maintaining a consistent view of program execution between an architectural and micro-architectural representation. Using the predictions and validations from an architectural model in order to verify an RTL implementation is highly desirable, but it is a challenge to provide useful data and useful predictive behavior.

Let's consider a very simple example: how to determine the correct point during program execution to apply an external asynchronous event such as an interrupt. When applying an interrupt to an architectural representation it can be taken immediately (if enabled) upon receipt, causing the processor to take the exception and begin execution at whatever is defined as the interrupt handling address.

In a micro-architectural implementation, it is not so simple. For example, the interrupt input logic may contain oversampling to ensure that the interrupt logic is observed to be active for N clock cycles. Once the logic has decided that an interrupt is active, it then has to be merged into the instruction pipeline at an appropriate time. It may be decided that it is wasteful to discard a complex instruction which executes for 32 cycles if it has already been executing for 28 of those. It's better to take the interrupt latency penalty of 4 cycles, rather than discard and lose 28 cycles. This is one of many micro-architectural performance decisions that must be considered in order to ensure the best throughput versus responsiveness.

Making sure that the architectural model matches these (and many other similar) scenarios can be very difficult. Consider a simple problem of the staircase waveform on interrupts, shown in figure 6 below [11]. This is where the interrupt lines ramp up over time on each clock edge.




Figure 6: The staircase problem

The reference model will receive these net changes as a set, but it is impossible for an architectural model to know which interrupt was the highest priority at the exact time that the DUT would have sampled and propagated that particular value. This leads to a set of legal scenarios, any of which may be executed by the DUT. It is the task of the architectural reference to determine what this set of legal scenarios can be, and to ensure that the path the DUT has chosen belongs to that set. In the scenario above there are 9 legal outcomes: taking any one of the interrupts (0-7) and entering its interrupt service routine, or not taking an interrupt at all and continuing with normal program execution.

This automatic analysis of legal scenarios recently revealed a bug in the OpenHW Group's CV32E40X processor core. It involved the following sequence of events:

A set of randomly-generated external interrupt signals have been propagating into the local interrupts of the processor core. These interrupts are masked by two levels of logic, firstly there is the MIE (Machine interrupt Enable) CSR, and the global interrupt enable field of the MSTATUS register (MSTATUS.MIE). An interrupt is detected by evaluating the following expression:

```
IRQ = MSTATUS.MIE && ((MIE & MIP) != 0x0);
```

**Figure 7: Expression for detecting a valid interrupt**

The expression states that for an IRQ to be pending and enabled we must have the equivalent positional bits True in both MIE and MIP, and the global interrupt enable MSTATUS.MIE must also be True.

The upper 16 MIP bits (local Interrupt 0-15) are a direct representation of the interrupt pins on the core, bearing in mind that there is clocked logic to sample these pins.

When the mret instruction is executed, the expression in Figure 7 evaluates to true. However, since the interrupt pins have been toggling during the interval since the previous instruction was retired it is unknown which interrupt should currently be active. The analysis engine must identify the legal possibilities based on the signal transitions it receives via RVVI-TRACE and the current state of the processor, and ensure that the DUT has executed one of these.

In this instance, the DUT actually did not service any interrupt, it executed the ebreak instruction and entered debug mode instead. Since this did not match any of the legal scenarios an error was flagged. This bug is now captured in the OpenHW Group's GitHub issue tracker: <https://github.com/openhwgroup/cv32e40x/issues/665>

## FUTURE WORK

The CORE-V-VERIF verification environment continues to evolve both in response to new requirements from the expanding roadmap of new cores and to on-going learnings from current and previous verification efforts. An upcoming enhancement to the environment will involve automated generation of functional coverage of retired instructions. This is enabled by a common Tracer interface (RVVI) that supports all information required for instruction coverage and by the standard definition of the RISC-V ISA.

## SUMMARY

The OpenHW Group's Verification task group has been a pioneer in the field of RISC-V processor verification. Through the CORE-V-VERIF environment we have employed different approaches and evaluated the merits and shortcomings of each. With each generation the CORE-V-VERIF environment has improved to become more robust, more reusable, more efficient, and ultimately better at finding RTL bugs. The current generation of CORE-V-VERIF uses RISC-V processor verification IP enabled by the RVVI to realize a comprehensive verification methodology that encompasses asynchronous peripheral events that occur randomly during program execution. This is the state-of-the-art methodology at present, however the verification task group members are highly motivated to continue to innovate and advance the practice of RISC-V processor verification.

## ACKNOWLEDGEMENTS

The authors would like to recognize the participation and contribution to this work of several of the OpenHW Group collaborators: Simon Davidmann and Aidan Dodds of Imperas Software, Greg Tumbush of EM Microelectronic, and Steve Richmond, formerly of Silicon Labs.

## REFERENCES

- [1] <https://openhwgroup.org>
- [2] *RVVI - RISC-V Verification Interface* <https://github.com/riscv-verification/RVVI>
- [3] <https://github.com/openhwgroup/core-v-cores>
- [4] <https://github.com/openhwgroup/core-v-verif>
- [5] [https://docs.openhwgroup.org/projects/core-v-verif/en/latest/quick\\_start.html](https://docs.openhwgroup.org/projects/core-v-verif/en/latest/quick_start.html)
- [6] <https://en.wikipedia.org/wiki/OVPsim>
- [7] <https://www.ovpworld.org/>
- [8] *Jump Start your RISC-V project with OpenHW*. DVCon US, March 1-4, 2021 (Virtual).
- [9] <https://github.com/SymbioticEDA/riscv-formal/blob/master/docs/rvfi.md>
- [10] <https://docs.openhwgroup.org/projects/cv32e40x-user-manual/en/latest/rvfi.html>
- [11] Figure created using <https://wavedrom.com>