

Solving Memory Subsystem Configurations Challenge with SV-Rand and Auto-Config Flow

One Class to Resolve All Parts

Kaushal Vala, VIP R&D, Cadence Design Systems, Ahmedabad, India (kaushalm@cadence.com)

Krunal Kapadiya, VIP R&D, Cadence Design Systems, Ahmedabad, India (krunalk@cadence.com)

Joseph Bauer, VIP R&D, Cadence Design Systems, San Jose, USA (jbauer@cadence.com)

Shyam Sharma, VIP R&D, Cadence Design Systems, San Jose, USA (ssharma@cadence.com)

Dharini SubashChandran, VIP R&D, Cadence Design Systems, Bangalore, India (dhairini@cadence.com)

Ritesh Desai, VIP R&D, Cadence Design Systems, Ahmedabad, India (riteshd@cadence.com)

Pooja Patel, VIP R&D, Cadence Design Systems, Ahmedabad, India (patelp@cadence.com)

Vatsal Patel, VIP R&D, Cadence Design Systems, Ahmedabad, India (vatsalp@cadence.com)

Abstract— Cadence Memory Models use configuration files to uniquely describe the characterization attributes of every legal and real memory device described in protocol standards and vendor data sheet. The Memory Model performs timing checks and protocol responses according to the combination of attribute settings in each device file. As the number of these memory device configurations increases per memory type, a significant challenge arises for the EDA provider and for the design verification engineers using them. Memory part offerings even within one protocol have increased into the tens of thousands while the number of characterization attributes describing a memory and the dependencies between them increases with each protocol generation. Moreover, the set of memory part offerings and attributes evolves over many version updates of the protocol standard and vendor data sheets during user's product development. The EDA provider is challenged to provide and maintain accurate updates for all the evolving real part configurations. In conjunction, the user's challenge is to validate that their SoC's (System on Chip) memory sub-system is compatible across the applicable memory part variations their SoC can externally connect to. To ensure compatibility as evolving updates are taken at regular intervals, the user has had to repeat their process steps to manage their repository snapshot of configurations while incorporating them into their test environment's part selection and coverage tracking.

A new widely adopted solution is available that represents all parts as one class file while simplifying the selection and coverage scoping to the right level of memory differentiating attributes aligned to the specification, vendor, and application. This "SystemVerilog constraint random configuration (SV-Rand)" flexible solution represents all valid parts in a single SystemVerilog class of constraints which resolves evenly across user's application scope of required configurations, while simultaneously providing compatibility coverage for closure over that same scoped set of parts. This constraints class, representing the relational intersection of all configuration settings relative to these memory differentiating attributes, is auto generated from a human friendly form for ease maintaining specification alignment. Additionally, Auto-Config layer in memory subsystem is an automated and scalable solution to reduce verification efforts, achieve faster time to market with no silicon escape.

Keywords— verification; memory subsystem; systemverilog constraint; automation; randomization; ddr5dimm; dfi; ddr-phy; auto-config; registers; timings;

I. RELATED WORK

A. SV-Rand Verification Flow:

While the "SystemVerilog constraint random configuration (SV-Rand)" has been applied to various protocols, we can use as example the DDR5 DIMM memory protocol to illustrate the challenge faced in the prior art. The variety of memory parts offered for a given protocol is determined by certain part differentiating configurations, and this combinational set can

be huge. For example, considering the part differentiating configurations below, DDR5 DIMM has 23 thousand parts of JEDEC and other vendors.

15 Operating Speeds x 5 Densities x 5 Stack-heights x 3 Data-widths x 2 DIMM-widths x 4 DIMM Types x 2 Rank

Memory differentiating attributes determines the settings on the many, often hundreds, of functional and physical characteristics that accurately represents a legal real each part in digital simulations. Directly maintaining hundreds of attribute key-value sets across this high number of configuration files is time-consuming and error prone, but necessary for qualifying the memory sub-systems in System-on-Chips (SoCs) for functional compatibility to these parts requires compliance testing across these many configuration combinations. The real and relevant space of part configurations is determined by these memory differentiating attributes since not all combinations of valid attribute values are legal or applicable to an application's target space, and ignoring legal applicable configurations is not good practice.

A challenge for users is the effort, time, resources to:

- Maintain the iterative download of EDA's provided files of "memory configurations SOMA files (Specification Of Memory Architecture)."
- Integrate them into their environment and validation flow.
- Target applicable configurations scoped within the set of all valid configurations.

This effort would include:

- Convert all "SOMA" files to native format such as SystemVerilog (SV) device classes.
- Compile all the SV device classes.
- Form an SV object queue of the many "SOMA" (for iteration and/or query).
- Form valid queries to meet some selection criteria (valid within available parameters and values and inter-dependencies).
- Query search the queue for the applicable matches which is still iterative.
- Evenly select one amongst the matched queue items.
- Provide and enable coverage tracking to refine targeted scope.

For an emerging memory protocol like DDR5, specification and ballot updates come frequently from the standards body which include timing and feature changes affecting these configurations. These repeated updates again require tedious efforts as just mentioned.

B. Auto-Config Layer:

Memory systems contain two major components, DDR memory controller (MC) and DDR PHY to access DDR memory. To control the data to and from the DRAM (Dynamic random-access memory) device and between the MC and PHY we have the DFI (DDR PHY Interface) [1]. This DFI standard encompasses signals, timing, and programmable parameters which are applicable to all DRAM protocols including DDRs, LPDDRs, Graphic DDRs and High Bandwidth Memory (HBM). The controller targets the timings with respect to clock, while the PHY uses physical timings, which causes interoperability challenges. The Auto-Configuration layer is crucial in simplifying the memory subsystem verification by automating DRAM-PHY configuration and reducing error-prone scenarios related configurations.

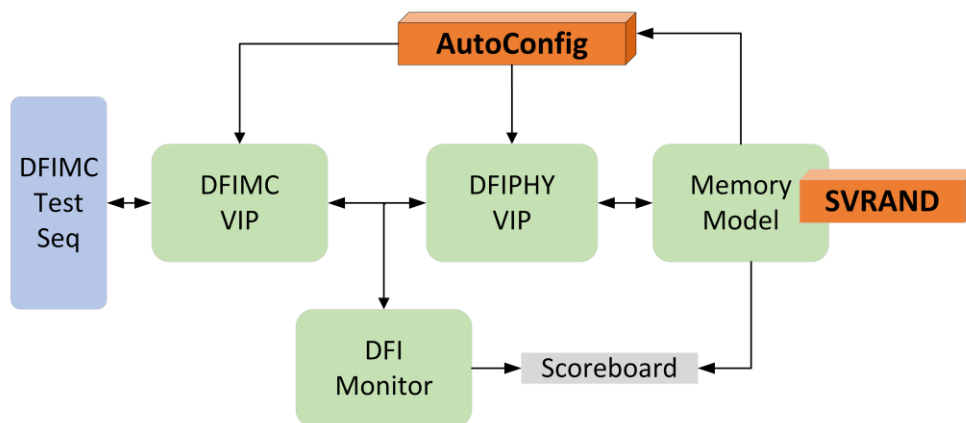


Figure 1. Memory Subsystem Verification Environment

New aged DRAM devices are challenged by high level complexities ^[2]. Multiple vendors support many timing parameters configuration registers.

- The need to configure DFI VIP timings to match delays of the DRAM device is demanding since DRAM timings are defined in term as of absolute values (denoted ns/ ps /fs) while DFI timings are in terms of clock (clk) based on selected frequency ratio.
- DFI protocol encompasses versions for all types of DRAMs. When verifying designs having various DRAMs, it would be cumbersome to have individual DFI solution for every type of DRAM. Having one DFI verification component catering to all the DRAMs would be the ideal solution since verification capabilities like configurability and randomization are similar for memory devices.
- There is a crucial need for configurable and randomizable verification solutions to achieve targeted functional coverage and to ensure no silicon escape. The new system design should be made available in the correct market window, ensuring the technology and architecture are highly accurate and complete.

II. APPLICATION

A. SV-Rand Verification Flow:

The SV-Rand solution automates the generation of these many part-configurations. Instead of maintaining the many thousands of configuration files, one metadata file is now maintained. The metadata file contains the same parameter set as listed in the many “SOMAS,” but instead of a single value associated to each parameter, the set of all possible values are listed, AND for each value any dependency the noted memory differentiating attributes. With every new revision of the protocol specification, it becomes much easier to maintain the one metadata file and with less opportunity for flaws. The full set of parts (“SOMA”) and part coverage can be generated from this metadata by unfolding these dependency connections, i.e., every value combination on the small set of memory differentiating attributes provides the ‘key’ to context align the intersect of a value across the many, often hundreds, of functional and physical memory characterization attributes.

This metadata representation of parameters values and interdependences to the small set of memory differentiating attributes is also auto converted into an equivalent SystemVerilog representative set of valid value constraints and legal implication constraints we call SV-Rand, the “SystemVerilog constraint random memory configurator class” (with examples shared shortly).

The ‘key’ with “SystemVerilog constraint random configuration” and “SOMA” automation is deriving real and legal parameter value combinations from dependencies to a small set of memory differentiating attributes. We refer to these selection attributes as ‘High-Level Parameters,’ (HLPs) which for DDR5 include [Speed, density, width, etc.] i.e., the list mentioned earlier. Users can scope their criteria for part selection on these ‘High-Level Parameters’ and the constraints will resolve dependencies evenly to only ‘real’ legal parameter value combinations.

B. Auto Configure layer:

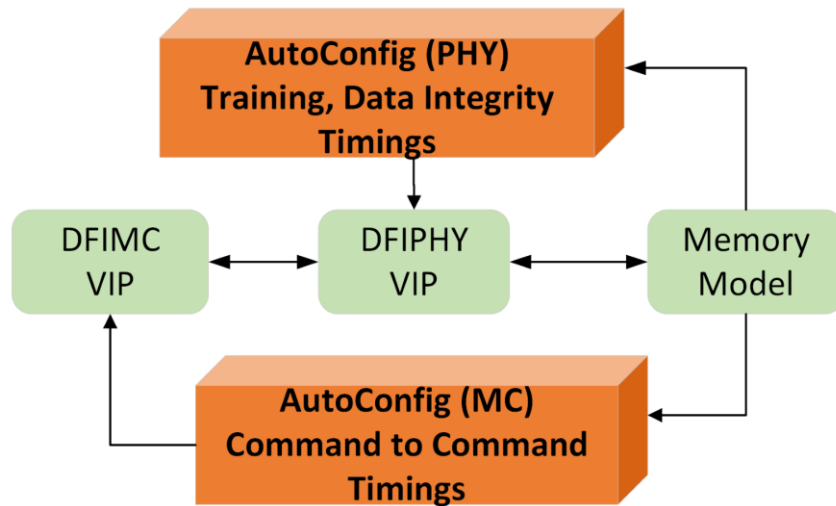


Figure 2. Auto Configuration of DFIMC-DFIPHY VIP

We have introduced an innovative Auto-Configured Layer that automates timing parameter updates for DFIMC/DFIPHY VIPs based on DRAM configuration. This approach streamlines the verification process, ensuring comprehensive coverage of configurations and scenarios.

How it Works:

- DFIMC VIP (Memory Controller): Maintains minimum Command to Command Delay for any memory device, as defined by JEDEC standards. Users connect a DRAM instance to the DFIMC VIP and configure the DRAM clock period. The VIP internally calculates and converts all timings into DFI clock terms, generating traffic accordingly.
- DFIPHY VIP (PHY): Maintains delays during training and Setup/Hold timing phases for the memory interface, adhering to JEDEC standards.

Key Features and Benefits:

- Automated Timing Configuration: Eliminates the manual, error-prone task of synchronizing DDR MC/PHY timing parameters with specific vendor parts.
- Comprehensive Support: Covers all DFI standards up to version 5.1 and memory devices used across various industries.
- Reduced Verification Risks: Ensures the accuracy of configured timings and registers, minimizing failures during verification stages. Users have reported a 70% reduction in effort during PHY/MC configuration, while eliminating the potential errors caused by manual adjustments.

III. RESULTS

A. SV-Rand Verification Flow:

The existing “SOMA” flow has a tedious pre-simulation task to collect all the parts available either through batch or manual configuration download from eMemory website. These downloaded parts must all be converted into SystemVerilog device class files, compiled, queued, and queried during simulation to select an applicable configuration.

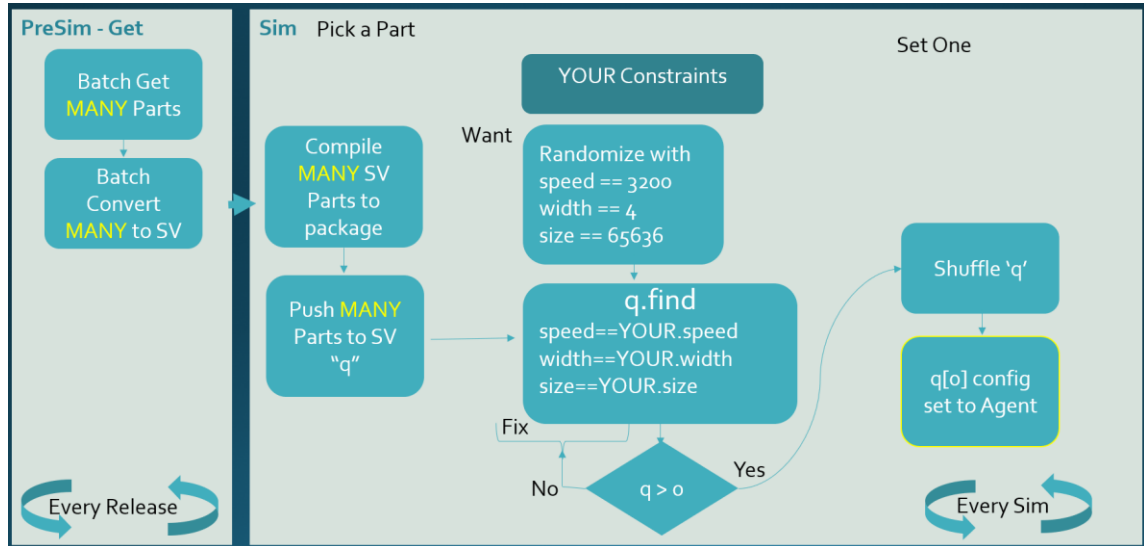


Figure 3. Comparison of existing “SOMA” flow versus new “SV-Rand” solution

For the new “SV-Rand” solution, only one SystemVerilog constraint class file is required. Subsets or specific parts within the space of all real and legal configurations can be targeted in a ‘randomize’ method call having in-line constraints on these ‘High-Level Parameters.’

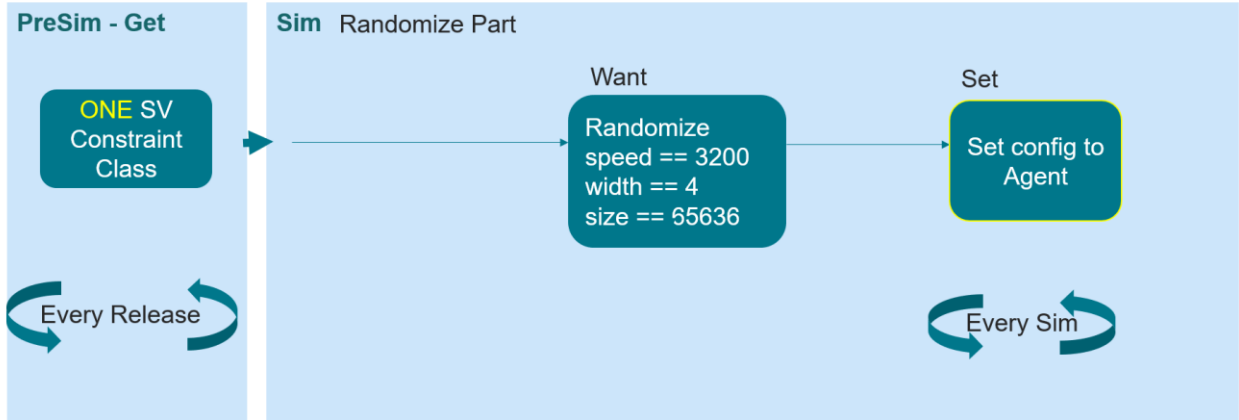


Figure 4. Randomization of SV-Rand to derive a part.

Since the HLP combinational value sets map 1:1 to legal real parts, the relational metadata also unfolds to provide comprehensive memory part coverage in “SV-Rand” for compatibility closure, e.g., DDR5 SDRAM:

Speed	6.67%	1 / 15 (6.67%)	[jedec_ddr5_64g_x8_6400b_part]	100%
Size	10%	2 / 20 (10%)	[jedec_ddr5_64g_x8_6400bn_part]	100%
DataWidth	33.33%	1 / 3 (33.33%)	[jedec_ddr5_64g_x8_6400c_part]	100%
RowWidth	33.33%	1 / 3 (33.33%)	[jedec_ddr5_8g_x8_6400an_part]	100%
ColWidth	50%	1 / 2 (50%)	[jedec_ddr5_8g_x8_6400b_part]	100%
NumBankGroups	50%	1 / 2 (50%)	[jedec_ddr5_8g_x8_6400bn_part]	100%
NumBanks	50%	1 / 2 (50%)	[jedec_ddr5_8g_x8_6400c_part]	100%
StackHeight3ds	20%	1 / 5 (20%)	[jedec_ddr5_16g_x8_6800an_part]	100%
PartNumber	0.02%	1 / 4320 (0.02%)	[jedec_ddr5_16g_x8_6800b_part]	100%

Figure 5. DDR5 SDRAM part configuration coverage

i. Query Part Select Example:

The traditional SOMA approach requires maintaining and compiling many part classes and then creating an object queue of all these valid parts to verify all possible valid configurations, in the case of DDR5-DIMM more than thousands of valid configurations needs to be verified. And still a unique HLP value set per part must be correctly added into each part class as selection criteria for user to query on

Input: DIMM-settings is randomized with parts which only pushed to queue.

```

dimmmSettings settings;
assert( settings.randomize() with {
    speed == spd ;
    numRanks == 2;
    dramDataWidth == 4;
    size == 524288;
    subClass == UDIMM;
    compsIn3dsStack == 1;
}

);
return settings.get_matching_part(settings);

```

Output: DIMM part is selected as it only existed in queue when queue size is more than zero.

```
Query part selected = jedec_ddr5udimm_64gbyte_2r_x4_3200bn, udimm, 2r, 3200, x64, x4, size=524288
```

ii. SV-Rand Solution Example :

Above traditional SOMA approach of putting all valid configurations into a queue for part selection means maintaining, compiling, and queuing many configurations, e.g., DDR5 DIMM has tens of thousands. Whereas the new solution is one class object for a user to randomize and scope applicable configurations as needed to High-Level Parameters.

For example:

1. To select a specific part configuration – In-line constrain all High-Level Parameters

Input: DIMM selection randomized on High Level Parameters

```
JEDEC_RandSettings = new();
assert( JEDEC_RandSettings.randomize() with {
    kind == lrdimm;
    density == 128;
    prank == 2;
    width == 4;
    lrank == 1; // non-3ds
    rate == an3200;
    eccType == withEcc;
} );
$cast( randSettings, JEDEC_RandSettings);
```

Output: Resolves DIMM part

```
jedec ddr5lrdimm 128gbyte 2r x4 3200c
```

2. To wildcard or widen the selection criteria such that ‘randomize’ resolves – Constrain some High-Level Parameters or none

Input: DIMM selection randomized on High Level Parameters

```
JEDEC_RandSettings = new();
assert( JEDEC_RandSettings.randomize() with {
    kind == lrdimm;
    density == 128;
    prank == 2;
    width == 4;
    lrank == 1; // non-3ds
    rate inside {an3200, b3200, bn3200, c3200} ;
    eccType == withEcc;
} );
$cast( randSettings, JEDEC_RandSettings);
```

Output: Resolves DIMM part within a range of devices

```
jedec ddr5lrdimm 128gbyte 2r x4 3200an
jedec ddr5lrdimm 128gbyte 2r x4 3200b
jedec ddr5lrdimm 128gbyte 2r x4 3200bn
jedec ddr5lrdimm 128gbyte 2r x4 3200c
```

B. Auto Configure layer:

The simulations successfully demonstrated the ability of the DFIMC VIP to maintain minimum command-to-command delays and the DFIPHY VIP to accurately implement training, setup, and hold timings, all according to JEDEC specifications.

As depicted in the below waveform, the Activate to Write, Write to Read and back-to-back Reads generated with minimum command delays. These results validate the functionality of both VIPs, demonstrating their capability to accurately manage memory interface timing and ensure compliance with JEDEC standards. This validation is crucial for ensuring the reliability and performance of memory systems based on these VIPs.

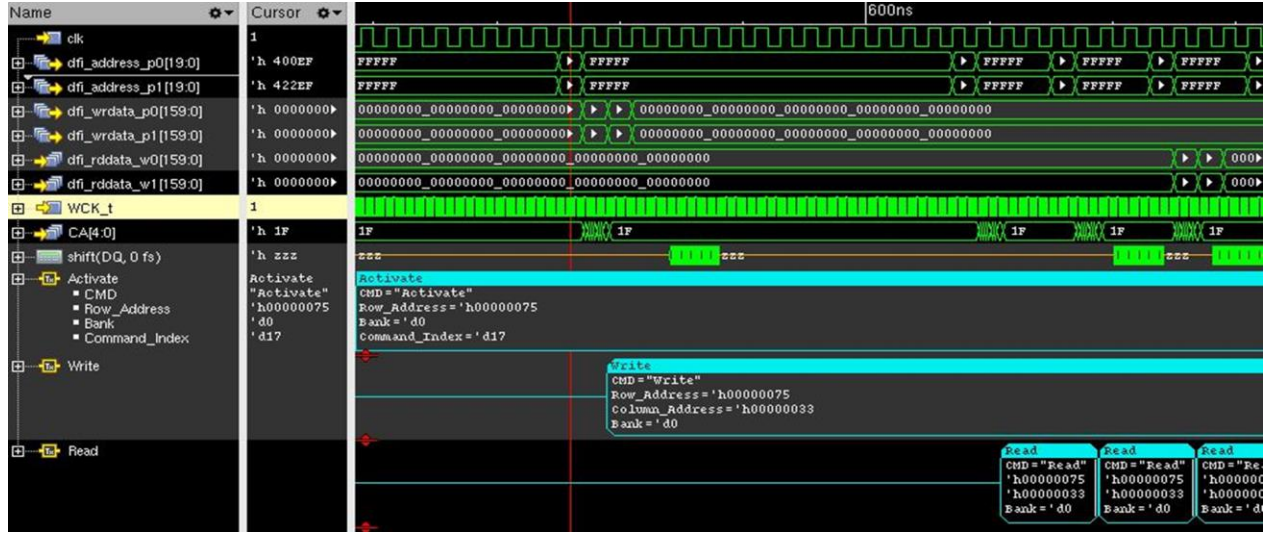


Figure 6. Auto Configuration Layer Results

IV. CONCLUSION

The traditional approach of “SOMA” flow demands significant effort, time, compute resources.

The new “SV-Rand” flow provides:

- Effective, optimized, native, user-friendly way to verify thousands of valid configurations.
- Widely adopted in industry.
- Overall configuration tasks cut by 30%.
- Simplified maintenance for EDA provider and SoC developer.
- Intuitive, comprehensive, fast part selection
- Part compatibility coverage closure measured.

Auto-Config layer provides:

- Considering manual configuration of DFIMC/PHY, it will not only be error-prone task but also debugging this is extremely challenging. There is more time spent on correctness of the parameters than the actual verification.
- Auto-Configuration technique reduces efforts up to 70% with 100% accuracy.
- It is important to test memory sub-systems in an exhaustive environment. This solution also generates all the traffic from the DFIMC with minimum timing allowed between them.

REFERENCES

1. DDR PHY Interface, Version 4.0, p. 15.
2. C. Kim, H. -W. Lee and J. Song, "Memory Interfaces: Past, Present, and Future," in IEEE Solid-State Circuits Magazine, vol. 8, no. 2, pp. 23-34, Spring 2016, doi: 10.1109/MSSC.2016.2546659.