"C" you on the faster side: Accelerating SV DPI based co-simulation

Hari Vinod Balisetty, Parag Goel, Amit Sharma DVCON-2014





Introduction: C Integration Support

- DPI SystemVerilog Standard
 - An interface between SystemVerilog and a foreign programming language: C or C++
- Simple interface to C models
 - Allows SystemVerilog to call a C function just like any other native SystemVerilog function/task
 - Variables passed directly to/from C/C++
 - NO need to write PLI-like applications/wrappers
- Why DPI
 - o Easy of use
 - Allows SystemVerilog to call a C function just like any other native SystemVerilog function/task
 - Direct interface provides better performance
- Support both functions and tasks

DPI: Two-way Communication

Import "DPI"

 SystemVerilog calling C/C++ functions

<pre>import "DPI-C" context task c_test(input int addr);</pre>			
program automatic top; initial c_test (1000); initial c_test (2000); endprogram	<pre>#include <svdpi.h> void c_test(int addr) { }</svdpi.h></pre>		

• Export "DPI"

- C calls SystemVerilog functions
- C calls SystemVerilog (blocking) tasks



DPI: Pure vs. Context Declarations



Co-simulation Using DPI : Transactor Based Verification



Result: The speed improvement over cycle-based can be orders of magnitude faster reaching tens of MHz



Step 1: An 'import' task/function which carries the relevant data to the hardware side.

Step II: The Intermediate C layer maps HW/SW calls with the correct data types and

Step III: forwards the data and invokes the hardware time consuming method through an export task.

Achieving co-simulation performance

Testbench and DUT are synchronized on a cycle basis, and communication overhead occurs at each and every cycle



Cycle-Based Test Bench Communication Overhead Emulated DUT

With Transaction-Based Emulation, synchronization is done only when required

Challenge: Reduce the inefficiencies in *Communication*

- Reduced Frequency of DPI calls
- Enabling efficient communication
 - Appropriate usage of language constructs
 - Recommendations for more efficient data transfers

- Improved communication schemes
 - o Offline processing
 - Leveraging Multi-threading
 - Enabling concurrency in the 'C' domain
 - Efficient File I/O

Reduced Communication Overhead Appropriate Usage Of Language Constructs

- Declare imported functions as 'pure' whenever possible, to allow for more optimizations.
- Preferably use data types that map to native C-data
- Avoid usage of scalar data types (bit/reg/logic) mapped to scalar "unsigned char", and the packed counterparts mapped to canonical form.
- The SystemVerilog-specific types, including packed types (arrays, structures, unions), which have no natural correspondence in C.

SystemVerilog	C (input)	C (out/inout)	SystemVerilo	og C (input)	C (out/inout)
byte	char	char*	bit	svBit	svBit*
shortint	short int	short int*	logic, reg	svLogic	svLogic*
int	int	int*	bit[N:0] reg[N:0] logic[N:0]	const svBitVecVal*	svBitVecVal*
longint	long int	long int*			
shortreal	float	float*		const svLogicVecVal*	svLogicVecVal*
real	double	double*	Open array[1		
string	const char*	char**	(import only)	const svOpenArrayHandle	svOpenArrayHandle
string[n]	const char**	char**	chandle	const void*	void*

Reduced Communication Overhead Reduced Number of Argument in DPI calls

- Each argument should be declared as narrow as possible
 - for example, never use the 'int' data type if the argument can only be 0/
 1.
- Combine them and prune argument lists to reduce the amount of data transferred from the simulator



Reduced Communication Overhead Reduced Frequency of DPI calls

- If DPI calls are very frequent,
 - Each DPI call results in a transfer of at least a n-bit packet header. Additional n-bit words are added to the packet to convey argument values.
 - When passing a limited number of constants, try to create multiple versions of DPI functions to eliminate the need for constants.



Reduced Communication Overhead Reduced Frequency of DPI calls

• Reduction of redundant calls, for example,

 if both f() and g() are functions that take a one bit wide argument each (SV 'bit' data type), then "f(a); f(b); g(c); g(d);" causes a transfer of 4*2(n-bit) words. One way to minimize the total amount of words transferred is to combine calls. In the above example, by creating a new wrapper function ffgg(bit a, bit b, bit c, bit d), we can reduce the total transfer size from 4*2(n-bit) words to 2(n-bit).



Reduced Communication Overhead Tips for efficient data transfers

To enhance data transfers to the host PC (System Verilog),

- Avoid unused bits by using the narrowest data type possible
- Avoid logic vector (4-state) arguments and unpacked structures
- Use arguments whose width is less than or equal to 32 bits
- Concatenate large numbers of short vectors in Verilog and unpack them on the C++ side. This will ensure more data transfer with less communication overhead.
- Bit-vectors aligned with 32-bit or multiple of 32-bits boundaries yields greater performance. So does the use of 64 bit wide integers as they are natively mapped as longint on C-side.

On the C-front:

- If the memory/array is 16-bit/8-bit wide then use appropriate data type for building C-side arrays. For example, use byte to map the 8-bit wide vectors, shortint for 16-bit wide vectors etc.
 - Using wider elements than needed is a bad idea for two reasons:
 - RAM usage goes up, and
 - CPU cache may become clogged with junk.

Reduced Communication Overhead Communication schemes

Offline processing

- Parameters that are passed via the DPI call to the 'C' side are offloaded for offline process,
 - by assigning these to some global data structures defined in the imported domain and the main function resumes the execution of the functionality further.

Leveraging Multi-threading



Reduced Communication Overhead Communication schemes

Efficient File Operations

 <u>Original Case:</u> In this example we have 56 variable sized memories that needs to be loaded using the HEX dump files



Results & Summary

Improvements observed in simulation performance and memory overhead for the following setups:

- A UVM frontend leveraging SV DPI to communicate with a Synthesizable PCIe transactor
- DPI based scoreboard and reference for a video decoder block

Changes Done	Percentage Improvement
Improvement using less number of DPI calls : 100 secs for 35 DPI calls got reduced to 30 secs for 20	3.33X
Using efficient file IO operations : usage of standard file streaming leads to a reduction from 90 secs to 15 secs	6X
Loading binary files against hex files	5X
Introducing threading : in file loading operation – time improved from 50 secs to 10 secs	5X
Usage of correct and compact and native C data types	1.5X
Offline processing on the C-side	2X