

Writing TLM2.0-compliant timed SystemC simulation models for SoCLib

Authors : Alain Greiner, François Pêcheux, Aline Vieira de Mello

1. [A\) Introduction](#)
2. [B\) Single VCI initiator and single VCI target](#)
3. [C\) TLM-T VCI Transaction](#)
4. [D\) VCI initiator Modeling](#)
 1. [D.1\) Member variables](#)
 2. [D.2\) Sending a VCI command packet](#)
 3. [D.3\) Receiving a VCI response packet](#)
 4. [D.4\) Initiator Constructor](#)
 5. [D.5\) Lookahead parameter](#)
 6. [D.6\) VCI initiator example](#)
5. [E\) VCI target modeling](#)
 1. [E.1\) Member variables & methods](#)
 2. [E.2\) Receiving a VCI command packet](#)
 3. [E.3\) Sending a VCI response packet](#)
 4. [E.4\) Target Constructor](#)
 5. [E.5\) VCI target example](#)
6. [F\) VCI Interconnect modelling](#)
 1. [F.1\) Generic network modeling](#)
 2. [F.2\) Arbitration Policy](#)

A) Introduction

This document is still under development.

It describes the modeling rules for writing TLM-T SystemC simulation models for SoCLib that are compliant with the new TLM2.0 OSCI standard. These rules enforce the PDES (Parallel Discrete Event Simulation) principles. In the TLM-T approach, we don't use anymore the SystemC global time, as each PDES process involved in the simulation has its own local time. PDES processes (implemented as SC_THREADS) synchronize through messages piggybacked with time information. Models complying to these rules can be used with the "standard" OSCI simulation engine (SystemC 2.x) and the TLM2.0 protocol, but can also be used also with others simulation engines, especially distributed, parallelized simulation engines.

The examples presented below use the VCI/OCF communication protocol selected by the SoCLib project, but the TLM-T approach described here is very flexible, and is not limited to the VCI/OCF communication protocol.

The interested user should also look at the [general SoCLib rules](#).

B) Single VCI initiator and single VCI target

Figure 1 presents a minimal system containing one single VCI initiator, **my_initiator**, and one single VCI target, **my_target**. The **my_initiator** module behavior is modeled by the SC_THREAD **execLoop()**, that contains an infinite loop. The call-back function **vci_rsp_received()** is executed when a VCI response packet is received by the initiator module.



Unlike the initiator, the target module has a purely reactive behaviour and is therefore modeled as a simple call-back function. In other words, there is no need to use a `SC_THREAD` for this simple target component: the target behaviour is entirely described by the call-back function `vci_cmd_received()`, that is executed when a VCI command packet is received by the target module.

The VCI communication channel is a point-to-point bi-directionnal channel, encapsulating two separated uni-directionnal channels: one to transmit the VCI command packet, one to transmit the VCI response packet.

C) TLM-T VCI Transaction

The TLM2.0 standard allows the user to redefine both the payload and the phases of a transaction. Two classes are defined in `soclib_vci_types` : a `tlmt_vci_transaction` and a `tlmt_phase`".

A payload of type `soclib_vci_types::tlm_payload_type` corresponds to a `tlmt_vci_transaction`. It contains three groups of information:

- TLM2.0 related fields
- TLM-T related fields
- VCI related fields

```
class tlmt_vci_transaction
{
    ...
private:
    // TLM2.0 related fields and common structure

    sc_dt::uint64      m_address;           // address
    unsigned char*     m_data;              // buf
    unsigned int        m_length;           // nword
    tlmt_response_status m_response_status; // rerror
    bool               m_dmi;              // nothing
    unsigned char*     m_byte_enable;       // be
    unsigned int        m_byte_enable_length;
    unsigned int        m_streaming_width;   //

    // TLM-T related fields

    bool*               m_activity_ptr;
    sc_core::sc_time*   m_local_time_ptr;

    // VCI related fields

    tlmt_command         m_command;         // cmd
    unsigned int          m_src_id;          // srcid
    unsigned int          m_trd_id;          // trdid
    unsigned int          m_pkt_id;          // pktid
```

The TLM2.0 compliant accessors allow to set the TLM2.0 related fields, such as the transaction address, the byte enable array pointer and its associated size in bytes, and the data array pointer and its associated size in bytes. The byte enable array allows to build versatile packets thanks to a powerful but slow data masking scheme. Further experiments are currently done to evaluate the performance degradation incurred by the byte formatting. It is therefore possible that the types of the `m_data` and `m_byte_enable` of the `tlmt_vci_transaction` will be changed to `uint32*` in a near future.

Dedicated VCI accessors are used to define the VCI transaction type, that can either be `set_read()` (for read command), `set_write()` (for write command), `set_locked_read()` (for atomic locked read), and `set_store_cond()` (for atomic store conditional). The `set_src_id()`, `set_trd_id()` and `set_pkt_id()` functions respectively set the VCI

source, thread and packet identifiers.

D) VCI initiator Modeling

D.1) Member variables

In the proposed example, the initiator module is modeled by the **my_initiator** class. This class inherits from the standard SystemC **sc_core::sc_module** class, that acts as the root class for all TLM-T modules.

The initiator local time is contained in a member variable named **m_localTime**, of type **sc_core::sc_time**. The local time can be accessed with the following accessors: **addLocalTime()**, **setLocalTime()** and **getLocalTime()**.

```
sc_core::sc_time m_localTime;           // the initiator local time
...
void addLocalTime(sc_core::sc_time t);   // add an increment to the local time
void setLocalTime(sc_core::sc_time& t);  // set the local time
sc_core::sc_time getLocalTime(void);     // get the local time
```

The boolean member variable **m_activity** indicates if the initiator is currently active. It is used by the arbitration threads contained in the **vci_vgmn** interconnect, as described in section F. The corresponding access functions are **setActivity()** and **getActivity()**.

```
bool m_activity;
...
void setActivity(bool t);                // set the activity status (true if the comp
bool getActivity(void);                  // get the activity state
```

The **execLoop()** method, describing the initiator behaviour must be declared as a member function.

Finally, the class **my_initiator** must contain a member variable **p_vci_init**, of type **tlmt_simple_initiator_socket**. This member variable represents the VCI initiator port.

D.2) Sending a VCI command packet

To send a VCI command packet, the **execLoop()** method must use the **nb_transport_fw()** method, that is a member function of the **p_vci_init** port. The prototype of this method is the following:

```
tlm::tlm_sync_enum nb_transport_fw
( soclib_vci_types::tlm_payload_type &payload,           // VCI payload pointer
  soclib_vci_types::tlm_phase_type    &phase,           // transaction phase (TLMT_CMD)
  sc_core::sc_time                     &time);          // local time
```

The first argument is a pointer to the payload, the second represents the phase (TLMT_CMD in this case), and the third argument contains the initiator local time.

The **nb_transport_fw()** function is non-blocking. To implement a blocking transaction (such as a cache line read, where the processor is stalled during the VCI transaction), the model designer must use the SystemC **sc_core::wait(x)** primitive (**x** being of type **sc_core::sc_event**): the **execLoop()** thread is then suspended, and will be reactivated when the response packet is actually received.

D.3) Receiving a VCI response packet

To receive a VCI response packet, a call-back function must be defined as a member function of the class

my_initiator. This call-back function (named **vci_rsp_received()** in the example), is executed each time a VCI response packet is received on the **p_vci_init** port. The function name is not constrained, but the arguments must respect the following prototype:

```
tlm::tlm_sync_enum vci_rsp_received
( soclib_vci_types::tlm_payload_type &payload,      // payload
  soclib_vci_types::tlm_phase_type   &phase,        // transaction phase (TLMT_RSP)
  sc_core::sc_time                   &time);        // response time
```

The return value (type `tlm::tlm_sync_enum`) is not used in this tlmt implementation, and must be systematically set to `tlm::TLM_COMPLETED`.

In the general case, the actions executed by the call-back function depend on the response transaction type (**m_command** field), as well as the **pktid** and **trdid** fields. For sake of simplicity, the call-back function proposed in the example below does not make any distinction between transaction types.

D.4) Initiator Constructor

The constructor of the class **my_initiator** must initialize all the member variables, including the **p_vci_init** port. The **vci_rsp_received()** call-back function being executed in the context of the thread sending the response packet, a link between the **p_vci_init** port and this call-back function must be established.

The **my_initiator** constructor for the **p_vci_init** object must be called with the following arguments:

```
p_vci_init.register_nb_transport_bw(this, &my_initiator::vci_rsp_received);
```

D.5) Lookahead parameter

The SystemC simulation engine behaves as a cooperative, non-preemptive multi-tasks system. Any thread in the system must stop execution after at some point, in order to allow the other threads to execute. With the proposed approach, a TLM-T initiator will never stop if it does not execute blocking communication (such as a processor that has all code and data in the L1 caches).

To solve this issue, it is necessary to define -for each initiator module- a **lookahead** parameter. This parameter defines the maximum number of cycles that can be executed by the thread before it is automatically stopped. The **lookahead** parameter allows the system designer to bound the de-synchronization time interval between threads.

A small value for this parameter results in a better timing accuracy for the simulation, but implies a larger number of context switches, and a slower simulation speed.

D.6) VCI initiator example

```
//////////////////////////////// my_initiator.h //////////////////////////////////
```

```
//////////////////////////////// my_initiator.cpp //////////////////////////////////
```

E) VCI target modeling

In the proposed example, the **my_target** component handles all VCI commands in the same way, and there is no error management.

E.1) Member variables & methods

The class **my_target** inherits from the class **sc_core::sc_module**. The class **my_target** contains a member variable **p_vci_target** of type **tlmt_simple_target_socket**. This object has 3 template parameters, that are identical to those used for declaring initiator ports (see above). It contains a call-back function as described below.

E.2) Receiving a VCI command packet

To receive a VCI command packet, a call-back function must be defined as a member function of the class **my_target**. This call-back function (named **vci_cmd_received()** in the example), is executed each time a VCI command packet is received on the **p_vci_target** port. The function name is not constrained, but the arguments must respect the following prototype:

```
tlmt::tlm_sync_enum vci_cmd_received
( soclib_vci_types::tlm_payload_type &payload,      // VCI payload pointer
  soclib_vci_types::tlm_phase_type   &phase,        // transaction phase
  sc_core::sc_time                   &time);        // time
```

E.3) Sending a VCI response packet

To send a VCI response packet the call-back function **vci_cmd_received()** use the **nb_transport_bw()** method, that is a member function of the class **tlmt_simple_target_socket**, and has the same arguments as the **nb_transport_fw()** function. Respecting the general TLM2.0 policy, the payload argument refers to the same **tlmt_vci_transaction** object for both the **nb_transport_fw()** and **nb_transport_bw()** functions, and the associated call-back functions. The **set_response_status** field must be documented for all transaction types, but only two values are used in this TLM-T implementation:

- TLMT_OK_RESPONSE
- TLMT_ERROR_RESPONSE

For a reactive target, the response packet time is computed as the command packet time plus the target intrinsic latency.

```
payload.set_response_status(soclib::tlmt::TLMT_OK_RESPONSE);
phase = soclib::tlmt::TLMT_RSP;
time = time + (nwords * UNIT_TIME);
p_vci_target->nb_transport_bw(payload, phase, time);
```

E.4) Target Constructor

The constructor of the class **my_target** must initialize all the member variables, including the **p_vci_target** port. The **vci_cmd_received()** call-back function being executed in the context of the thread sending the command packet, a link between the **p_vci_target** port and the call-back function must be established. The **my_target** constructor must be called with the following arguments:

```
p_vci_target.register_nb_transport_fw(this, &my_target::vci_cmd_received);
```

E.5) VCI target example

```
//////// my_target.h //////////////////////////////////////////
```

```
//////// my_target.cpp //////////////////////////////////////////
```

F) VCI Interconnect modelling

The VCI interconnect used for the TLM-T simulation is a generic interconnection network, named **VciVgmn**. The two main parameters are the number of initiators, and the number of targets. In TLM-T simulation, we don't want to reproduce the cycle-accurate behavior of a particular interconnect. We only want to simulate the contention in the network, when several VCI initiators try to reach the same VCI target. Therefore, the network is actually modeled as a complete cross-bar : In a physical network such as the multi-stage network described in Figure 2.a, conflicts can appear at any intermediate switch. In the **VciVgmn** network described in Figure 2.b, conflicts can only happen at the output ports. It is possible to specify a specific latency for each input/output couple. As in most physical interconnects, the general arbitration policy is round-robin.



F.1) Generic network modeling

There is actually two fully independent networks for VCI command packets and VCI response packets. There is a routing function for each input port, and an arbitration function for each output port, but the two networks are not symmetrical :

- For the command network, the arbitration policy is distributed: there is one arbitration thread for each output port (i.e. one arbitration thread for each VCI target). Each arbitration thread is modeled by a `SC_THREAD`, and contains a local clock.
- For the response network, there are no conflicts, and there is no need for arbitration. Therefore, there is no thread (and no local time) and the response network is implemented by simple function calls.

This is illustrated in Figure 3 for a network with 2 initiators and three targets :



F.2) Arbitration Policy

As described above, there is one **cmd_arbitration** thread associated to each VCI target. This thread is in charge of selecting one timed request between all possible requesters, and to forward it to the target. According to the PDES principles, the arbitration thread must select the request with the smallest timestamp. The arbitration process must take into account the actual state of the VCI initiators: For example a DMA coprocessor that has not yet been activated will not send request and should not participate in the arbitration process. As a general rule, each VCI initiator must define an **active** boolean flag, defining if it should participate to the arbitration. This **active** flag is always set to true for general purpose processors. Any arbitration thread receiving a timed request is resumed. It must obtain an up to date timing & activity information for all its input channels before making any decision. To do that, the `LocalTime` and `ActivityStatus` of all VCI initiators are considered as global variables, that can be accessed (for read only) by all arbitration threads. The arbitration policy is the following : The arbitration thread scans all its input channels, and selects the smallest time between the active initiators. If there is a request, this request is forwarded to the target, and the arbitration thread local time is updated. If not, the thread is descheduled and will be resumed when it receives a new request.

For efficiency reasons, in this implementation, each arbitration thread constructs - during elaboration of the simulation - two local array of pointers (indexed by the input channel index) to access the `LocalTime` and `ActivityStatus` variables of the corresponding VCI initiators. To get this information, each arbitration thread uses the `nb_transport_bw()` *function on all its VCI target ports, with a a dedicated phase called* `soclib::tlmt::TLMT_INFO`. *The payload argument refers to the same tlmt_vci_transaction object as the two other*

phases (TLMT_CMD and TLMT_RSP).

```
for (size_t i=0;i<m_nbinit;i++) {  
    phase    = soclib::tlmt::TLMT_INFO;  
    m_RspArbCmdRout[i]->p_vci->nb_transport_bw(payload, phase, rspTime);  
    m_array[i].activity = payload.get_activity_ptr();  
    m_array[i].time = payload.get_local_time_ptr();  
}
```

As the net-list of the simulated platform must be explicitly defined before constructing those LocalTime and ActivityStatus arrays, the vgm hardware component provides an utility function **fill_time_activity_arrays()** that must be called in the SystemC top-cell, before starting the simulation.