

You Are In a Maze of Twisty Little Sequences, All Alike – or Layering Sequences for Stimulus Abstraction

Rich Edelman
Mentor Graphics
rich_edelman@mentor.com

Adam Rose
Mentor Graphics
adam_rose@mentor.com

Andreas Meyer
Mentor Graphics
andy_meyer@mentor.com

Raghu Ardeishar
Mentor Graphics
raghu_ardeishar@mentor.com

Jason Polychronopoulos
Mentor Graphics
jason_polychronopoulos@mentor.com

ABSTRACT

This paper will demonstrate building layered stimulus using OVM sequences and sequencers. Virtual sequences and virtual sequencers will be demonstrated by building a small collection of examples that can be used in layered stimulus verification environments. The main contribution of this paper is a new layering component that performs the standard layering task while minimizing user programming without requiring exotic connectivity, extended components or the use of the factory.

Categories and Subject Descriptors

[Hardware Verification]: Functional Simulation and Verification – *class based SystemVerilog, OVM library, sequences, sequencers, drivers.*

General Terms

Verification, Simulation, Layered Sequences.

Keywords

Sequences, sequencers, layering, register based verification, MVC, VIP, Verification IP.

1. INTRODUCTION

Complex verification environments and modern communication protocols often have many parallel operations active at a time, and are difficult to model with a linear series of programming statements.

Sequences can be used to parallelize behavior, but bring along certain complexities that are difficult to manage for new users. This paper will discuss ways to hide the gory details of layering sequences, and the “hardness” about sequences, while enabling the productivity and infinite flexibility afforded by them.

Many kinds of systems fall into the area where sequences and layering are appropriate, including a system where the DUT examines multiple layers of a protocol or any system where higher level “instructions” need to be broken down into their lower level constituents.

Using sequences in these environments can make life easier, since sequences and sequencers are natively parallel and have arbitration and other “communicating process” hooks already built in.

A user of such a verification environment typically uses an OVM verification component (OVC) which has a library of sequences and a sequencer. This OVC is usually combined with other OVC types to build up a system. Each OVC is operating independently, and can be coordinated with a higher “layer”. This is the simplest aspect of layering discussed in this paper.

Layering stimulus is a common practice to achieve reuse and to raise abstraction levels. Protocols can be layered, functional decomposition can be expressed through layering – any higher to lower level abstraction is a possible layering target.

Layering allows a high level programming abstraction to be used for stimulus writing. Then the higher level programming abstraction is broken down into a lower level programming abstraction. The lower level programming abstraction is what comes with the OVC building block – it is the verification library that is built-in to the OVC. It can be used stand-alone, but most verification teams will prefer to use a higher level abstraction. This is the motivation for a layered stimulus approach – breaking the desired higher level abstraction into the available lower level abstraction.

For example, consider a register layer running on a low level bus transport layer. The high level register write() is implemented in terms of lower level bus transactions specific to a particular bus protocol, such as OCP bus transactions. This is a more complex aspect of layering. For a register layering situation, the higher level uses READS and WRITES. The lower level is specific to a bus protocol. For example a lower level AHB implementation knows how to wiggle the pins on an AHB bus, while a lower level OCP implementation knows how to wiggle the pins on an OCP bus. Each of these lower levels is very different, but the higher level reads and writes can run on either one, with the proper translation in place.

Layering will be discussed in this paper as applicable to packet or protocol layering, OVC layered stimulus and register stimulus generation over an OVC layer. This layering will be achieved with “virtual connections” or with layered abstractions or both.

Layered verification and layered software is not a new concept. A simple function call is a layer. The lower level function implements some basic behavior. The upper level function calls the lower level function to perform activities. For example, a lower level function could implement the multiply command using shifts and adds. The upper level function then could implement a higher level function like an (x, y) coordinate system. The (x, y) function could then be used at an ever higher level in an image processing function. Using layers in software design is now taken for granted to reduce complexity and promote reuse.

Function calls in software are layers. In the OVM [1] function calls can also be layers. Additionally, OVM sequences [2] can also be layered. Sequences are really just function calls [3] on steroids. A sequence is a special kind of function call – it is a “functor”. The implementation of the function is generally in the “body()” function of the sequence, but there are also other useful hooks like the pre_body(), pre_do(), post_do() and post_body() to name a few. Using the flexible sequence API allows for creating powerful and

interesting tests and layers of tests which reduce complexity and promote reuse.

2. OVERVIEW AND BACKGROUND

Sequences work together with sequencers and drivers to manage tests on an interface. Details of this system is beyond the scope of this paper, but a brief summary of certain functionality and API will be discussed here [4] [5], as related to the layering concept – Figure 1.

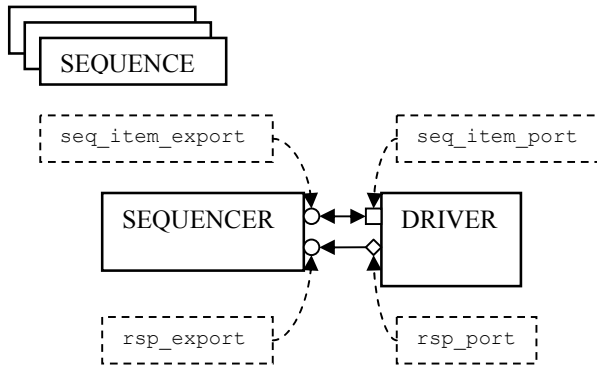


Figure 1 - Sequencer / Driver interface

3. LIONS AND TIGERS AND BEARS! OH, MY!

The collection of APIs associated with a sequence class is used to implement a function call, and to provide other interfaces for synchronization and arbitration of other sequences. A basic sequence exists to generate a "transaction". The transaction may be randomized, may be the "next" transaction in a sequence or may be related to some state of the system.

Another role of the sequence is to manage other sequences instead of generating a transaction. In the OVM, such a "virtual sequence" is really just a collection of SystemVerilog code that starts other sequences. It's a sequence coordinator. The sequence is really just a function call which does "something" that is important for stimulus generation or monitoring.

A driver in the OVM is SystemVerilog code which manages an interface; perhaps a bus or some other communication to another block. Usually a driver is managing a bus and knows how to wiggle the pins of that bus, causing bus traffic. It may also know how to monitor the bus and create transactions from bus activity.

A driver normally "gets" a transaction from elsewhere and then causes that transaction to be realized as pin wiggles. Bus traffic whether AHB or AXI or OCP, is being driven by the driver in response to the transaction.

A driver needs a transaction and a sequence generates a transaction.

We could just hook up a sequence to a driver. But most interesting test scenarios aren't just one sequence. There may be many prioritized sequences which compete for the opportunity to send a transaction. For example, two sequences which are generating bus traffic transactions are running in parallel to simulate two different transfers interleaved on the same bus.

A sequencer is used to arbitrate these competing prioritized sequences. A sequencer manages "who goes next" – among other

things. The sequencer manages the relationship between sequence and driver, and is usually a vanilla first-in-first-out arbitrator. It can be programmed for many different behaviors, all beyond the scope of this paper [4].

A monitor can be created separately from the driver which interacts with the bus. It interacts with the same interface as the driver, but is passive. It samples pin wiggles, or other low level communication (just as the driver drives that low level communication). The monitor recognizes a series of pin wiggles as a transaction (for example a READ), and then creates a transaction and publishes it.

3.1 Sequencer API

The sequencer has complex behavior, but can be interacted with in simple ways. For this paper we're only concerned with the driver connection and interface, and the way a sequence provides a transaction to a driver. We won't consider any of the more complex sequencer behavior.

The sequence creates an item and starts it. Starting an item causes the sequencer to interact with the driver and perform arbitration. A sequencer has a REQ/RSP export, and expects to be connected to a corresponding REQ/RSP port. Usually that corresponding port is on the driver. The sequencer expects to be connected to a "driver interface".

Sequence:

```
create_item()
start_item()
  sqr.wait_for_grant()
  seq.pre_do();
finish_item()
  seq.mid_do();
  sqr.send_request();
  sqr.wait_for_item_done();
  seq.post_do();
```

Driver:

```
seq_item_port.get()
sqr.get()
  m_req_fifo.peek();
  item_done();
```

3.2 Driver API

The driver has a simple behavior. It exists to just be told what to do, and then do it. The 'what to do' is a transaction. The transaction arrives from a sequencer which has arbitrated possibly competing sequences. Once a driver receives a transaction it may cause activity on the interface it is managing. Each driver implementation is protocol specific, and can vary widely, but the basics are shared among all drivers. The driver process is a loop which tests the sequencer to see if a transaction is available. For example, a driver loop can loop forever, doing a get() to retrieve a transaction from the sequencer. Aside from get(), there are other ways a driver can interact with a sequencer, but each results in a transaction being produced and acted upon by the driver.

When the driver has executed the transaction it supplies a response back to the sequencer. Those responses can be immediate or take time, depending on the protocol.

Simple driver design – get and transaction and execute it.

```

task run();
  ovcl_sequence_item t;
  ovcl_sequence_item rsp;
  forever begin
    seq_item_port.get(t);
    ovm_report_info("driver", t.convert2string());

    // Send to BUS / Wiggle pins
    ...
    // Construct response
    rsp = new();
    rsp.n = t.n;
    rsp.set_id_info(t);
    rsp_port.write(rsp);
  end
endtask

```

4. VIRTUAL SEQUENCES AND SEQUENCERS

Virtual sequences are just sequences whose job is to start other sequences. Virtual sequences do not produce a transaction, but rather cause other sequences to run, which can cause other sequences to run, until finally there are transactions (or some other work) created. The job of a sequence is to produce a transaction, or cause some work. A virtual sequence is a convenient description of sequences that don't directly produce transactions – but instead cause other sequences to produce transactions.

In Figure 2 a virtual sequence 'seqX' starts three sub-sequences in parallel (seqA, seqB and seqC). These three parallel sequences control the interface 1, 2 and 3, testing some interaction of parallel traffic.

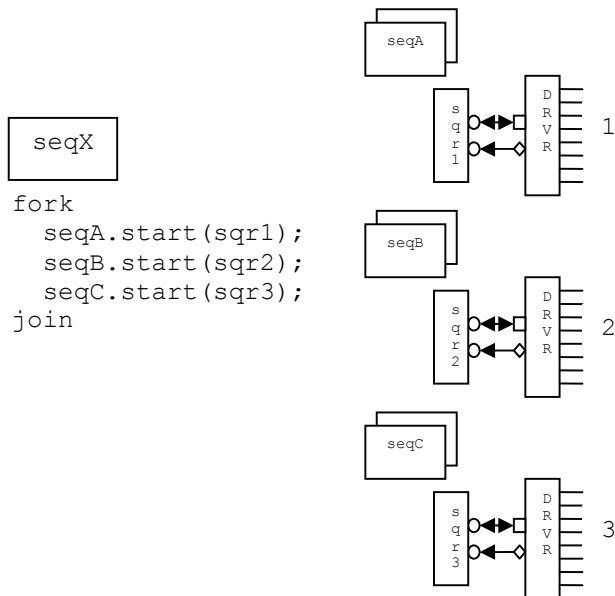


Figure 2 - Virtual sequence - seqX

In the same way that virtual sequences are sequences that don't produce transactions, virtual sequencers are sequencers that are not connected to a real driver. Instead they interact with other "sub-sequencers". Figure 2 could describe a virtual sequencer – the sequencer that seqX is running on.

5. LAYERED SEQUENCES

A sequence can be written to encapsulate some behavior and produce a transaction. It is a function call. Just as function calls can call other function calls, sequences can call other sequences – as seen in the virtual sequence example above.

Aside from a sequence calling another sequence, another kind of layering is important. This kind of layering is the layering that occurs when a sequence generates a transaction, and that transaction causes a different sequence to generate sub-transactions.

For example, a high level transaction might be generated to move 1024 bytes of data from A to B. This transaction would get "interpreted" or "translated" into a lower level transaction that was supported on the particular interface. The 1024 byte transfer from A to B might get mapped into eight 128 byte transfers. In Figure 3 an upper level seqX might generate a 1024 byte transfer, which gets received at the lower level and translated into the eight 128 byte transfers.

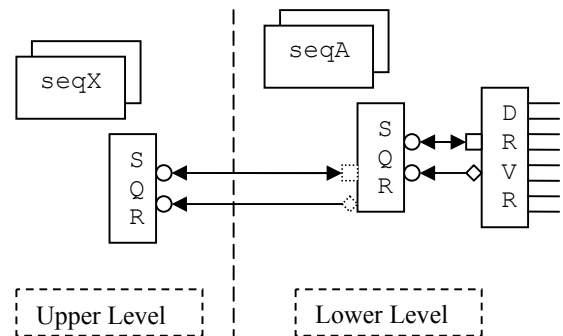


Figure 3 - Layered sequences

Layered sequences are used to reuse lower level sequences. Layered sequences are built supporting higher level constructs like packet protocols and register stimulus over an existing OVC.

5.1 Packets

A packet described here is a collection of bytes, possibly formatted to describe a network protocol or other structured data. A packet can also just be a collection of bytes. Building layered packets consists of structuring or un-structuring the bytes to represent a different abstraction level.




Figure 4 - One packet to one packet with detail added

Figure 4 shows a higher level collection of bytes which are converted to a lower level by adding a header and a checksum.

Figure 5 shows a higher level collection of bytes which is broken into 4 equally sized sub-blocks. Each sub-block has a header and checksum added.

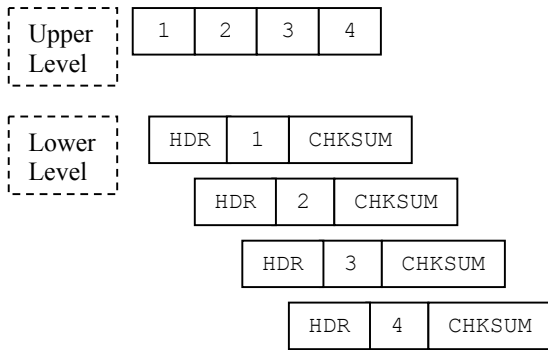


Figure 5 - One packet split into 4 sub-packets

Figure 6 shows a higher level abstraction which gets its headers and checksums striped, and then each payload block – 1, 2, 3, 4 – is concatenated into one large transaction.

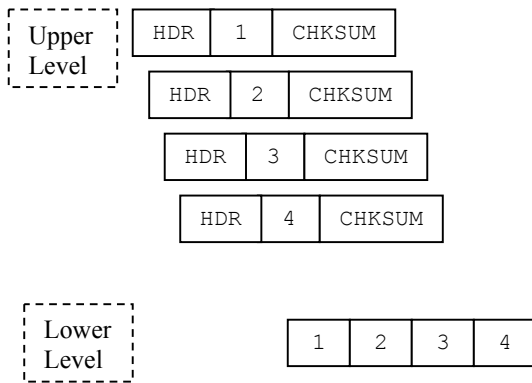


Figure 6 - Many packets consolidated to one

5.2 Instructions

Instructions can be implemented using layers – an instruction is a symbolic command at the higher level, for example MOV32 ADDR1, ADDR2, which is converted into a bus transaction that moves 32 bits from ADDR2 to ADDR1. [6]




Figure 7 - Machine instructions to microcode

This example of layering demonstrates higher level instructions like ‘ADD x’ and ‘STORE x’ being broken down into lower level instructions like ‘IR[add] -> MAR’, ‘A + MDR -> A’.

Using this high level layer, a software program could be written as a test program. When such a test program executes, it results in “commands” running – like ‘Store 20’ and ‘Add 20’. Those commands execute at the higher level, and get translated into lower level commands by the layering system and the actual hardware machine underneath.

5.3 Registers

Registers are an example of where layered sequences can be applied easily. A register transaction is defined as a READ or a WRITE with a register name specified and a data field provided for write, or a data returned for a read.




Figure 8 - Register layering

Registers layering is one of the easiest to achieve, since the layering translation is easy. In order to do the translation from higher to lower level, the register name is looked up in an address map. Given the register name, the address map returns the register address. In the example above, “regA” gets translated to the address ‘1002’.

6. USING THE LAYERING COMPONENT

As seen in Figure 9 to layer sequences requires and implies certain recommendations, such as connection by TLM ports, phased build(), connect(), etc. These details are quite tedious and error prone. The layering component allows for easier integration and use.

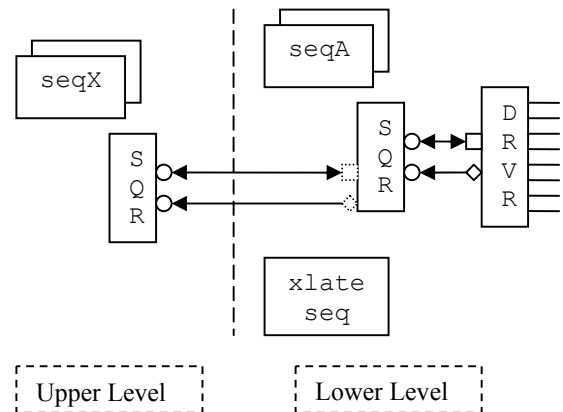


Figure 9 - Sequence translation

The layering component is a component that provides for hooks into the OVM infrastructure to provide construction, connection, and run phasing. It allows for creation of the required ports, the creation of the translation sequence itself, and the connection of the translation interface (driver interface) to the upper level sequencer.

6.1 Details

The layering component implements the connect(), build() and run() phases to achieve the layering. See the Appendix 1 for a sample implementation.

In build(), the layering component builds the required upper sequencer automatically. In the OVM, components must be built before the run phase, most typically in the build phase.

In `connect()`, the layering component builds the lower level ports required to connect to the upper layering sequencer. The ports are constructed in the connect phase since we can guarantee here that the build phase has passed, and the lower sequencer has been created. These lower level or downstream ports are used as a “way-point” for connecting the sequencer and driver. The lower level ports – the so-called “driver interface” are constructed with a parent pointer that is not the traditional “this” pointer. Instead, these ports are constructed with a parent pointer of the downstream sequencer. Essentially, the ports are built into the downstream sequencer, although they do not exist syntactically in any class definition. They are built dynamically. Once constructed, these ports are connected to the upper sequencer. Now the upper sequencer thinks it is connected to a driver style interface – which is really the two ports created here which are “hosted” or “parented” in the downstream sequencer.

In `run()`, the layering component performs its final chore. Once this chore is complete the layering component has no other functionality, and does nothing else. This final functionality is to start the upper and lower level sequences required for this particular layering. For example, in register layering, the upper level sequence started is a regular register sequence, and the lower level sequence is a special translation sequence which knows how to convert between a register transaction and a downstream transaction. See Appendix 2 for a sample register transaction translator sequence.

The translator sequence is running on the lower level sequencer. It “knows” how to create lower level transactions. In the example the lower level sequence has helper functions, `read()` and `write()` which work with the sequence API to cause lower level transactions to be created. The job of the translator sequence is to call these lower level helper functions with the correct arguments. Before the helper functions can be called the translator sequence must “get” a transaction – just like a driver would. This is the driver interface that has been built. The upper level sequencer thinks it is sending a transaction to an upper level driver. Really the transaction is being sent to a correctly typed port which is operated on like a driver would operate on it. The translator sequence uses the driver API – calling `port.get()` and `rsp_port.write()`, etc. As far as the upper level sequencer is concerned, the translator sequence is a driver.

6.2 Avoiding the Details

Building the connections, constructing the components and starting the sequences is a tedious task. These details can be easily avoided by using the layering component. The layering component is programmed using a simple 4-tuple. One of the hidden benefits of this approach is that no extended sequencer is required. Layering can be added to any existing sequencer/driver combination by instantiating an appropriate layering component, and registering the sequences to be run. Adding the layering component required less than 10 additional lines of code changed in an existing testbench.

A layering component needs to be constructed:

```
register_layer reg_layer = new("reg_layer", null);
```

Once constructed, the specific layering must be registered:

```
reg_layer.add_register_db(
    "simple_el.sequencer", // An existing sequencer.
    "translation_seq",   // The built-in
                        // translation sequence.

    "layered_reg_sqr",   // A sequencer type
                        // to create.
    "layered_reg_seq"   // The register sequence
                        // to start.
);
```

The behavior of the layering component is guided by a data structure containing the names of the sequencers and sequences. An alternate implementation can be used which uses types instead of strings for better error checking.

The data structure contains the following fields, ‘lower_sequencer’ – the name of the lower level sequencer, ‘translation_seq’ – the name of the sequence which will run on the lower level sequencer, ‘upper_sequencer’ – the name of the upper level sequencer that will be instantiated and started automatically, and ‘upper_sequence’ – the name of the sequence that will be started on the upper level sequencer (upper_sequencer).

```
typedef struct {
    string lower_sequencer;
    string translation_seq;
    string upper_sequencer;
    string upper_sequence;
    ...
} layer_db_t;
```

7. CONCLUSION

The layering component is a relatively unobtrusive way to manage the complexities of sequences, sequencers and drivers when the goal of the layering is to add a new set of stimulus or monitoring to an already existing infrastructure. For example, an existing bus level infrastructure might benefit from register based coverage using a register layering component with the OVM Register Package. Another example might be register based tests layered on top of the existing bus tests.

The layering component is not limited to register testing and monitoring. It can be used in any situation where an upper level transaction must be converted back and forth to a lower level transaction.

8. REFERENCES

- [1] OVM World Download, www.ovmworld.org
- [2] OVM User Guide.
- [3] Edelman, Rich, “Sequences in SystemVerilog”, DVCON 2008
- [4] OVM Reference Guide.
- [5] Meyer, Andreas, “Overview of Sequence Based Stimulus Generation in OVM 2.1”. [unpublished]
- [6] Crews, Thad “LMMS: An 8-bit Microcode Simulation of the Little Man Computer”, <http://www.citidel.org/bitstream/10117/117/6/LMMS.pdf>

APPENDIX 1 – REGISTER LAYERING LIBRARY EXAMPLE

```

package register_layer_pkg;
import ovm_pkg::*;
`include "ovm_macros.svh"

import ovm_register_pkg::*;

class layered_register_sequencer
  extends ovm_register_sequencer#(
    ovm_register_transaction,
    ovm_register_transaction);

  `ovm_sequencer_utils(
    layered_register_sequencer)

  function new(string name,
    ovm_component p);
  super.new(name, p);
  `ovm_update_sequence_lib_and_item(
    ovm_register_transaction)
  count = 0; // No automatic start
endfunction
endclass : layered_register_sequencer

class layered_register_sequence
  extends register_sequence_all_registers
  #(ovm_register_transaction,
    ovm_register_transaction);

  `ovm_sequence_utils(
    layered_register_sequence,
    layered_register_sequencer)

  function new(string name =
    "layered_register_sequence");
  super.new(name);
endfunction
endclass : layered_register_sequence

typedef struct {
  string lower_sqr;
  string translation_seq;
  // Runs on the lower_sqr

  string upper_sqr;
  string upper_seq;
  // Runs on the upper_sqr

  ovm_register_sequencer
  #(ovm_register_transaction,
    ovm_register_transaction)
  register_sqr;
  ovm_sequencer_base lower_sqr;
} layer_db_t;

class register_layer
  extends ovm_component;
  // Database of the layering
  layer_db_t layer_db[string];

  function new(string name,
    ovm_component p);
  super.new(name, p);
endfunction

  local function ovm_sequencer_base
  m_build_sequencer(
    string sequencer_name);

    ovm_sequencer_base sqr;
    string sequencer_instance_name;
    static int count = 0;

    // Create an instance name.
    $sformat(sequencer_instance_name,
      "%s%0d", sequencer_name, count++);

    $cast(sqr,
      factory.create_component_by_name(
        sequencer_name, get_full_name(),
        sequencer_instance_name, this));
    return sqr;
  endfunction

  local function ovm_sequencer_base
  m_find_sequencer(
    string sequencer_name);

    ovm_sequencer_base sqr;
    $cast(sqr, ovm_top.find(
      sequencer_name));
    return sqr;
  endfunction

  function void build();
  super.build();
  // Build all the register sequencers.
  foreach (layer_db[s])
    $cast(layer_db[s].register_sqr,
      m_build_sequencer(
        layer_db[s].upper_sqr));
  endfunction

  local function void m_connect(
    ovm_register_sequencer
    #(ovm_register_transaction,
      ovm_register_transaction)
    register_sqr,
    ovm_sequencer_base lower_sqr);

    ovm_seq_item_pull_port
    #(ovm_register_transaction,
      ovm_register_transaction)
    lower_sqr_register_seq_item_port;

    ovm_analysis_port
    #(ovm_register_transaction)
    lower_sqr_register_rsp_port;

    // Create some appropriate 'upstream'
    // ports, and slip them into
    // lower_sqr. Bus_sqr is none the
    // wiser. Notice: the parent pointer
    // is lower_sqr.
    lower_sqr_register_seq_item_port
    = new("register_seq_item_port",
      lower_sqr);
    lower_sqr_register_rsp_port
    = new("register_rsp_port",
      lower_sqr);

    // Pay attention here! We're calling
    // connect on the ports we just
    // slipped into the 'lower_sqr'. And
    // we're connecting to the upstream
    // register sequencer that we
    // magically created behind the
    // scenes in the build() phase above.
    lower_sqr_register_seq_item_port.connect(
      register_sqr.seq_item_export);
    lower_sqr_register_rsp_port.connect(
      register_sqr.rsp_export);
  endfunction

function void connect();
super.connect();
// Find the bus sequencers as named.
// We can find them in connect(),
// since we are guaranteed that they
// exist.
foreach (layer_db[s])
  $cast(layer_db[s].lower_sqr,
    m_find_sequencer(
      layer_db[s].lower_sqr));

  // Register all the ports, and do
  // the connects
  foreach (layer_db[s])
    m_connect(
      layer_db[s].register_sqr,
      layer_db[s].lower_sqr);
endfunction

// Called by the user to "setup" a
// layering.
function void add_layer_db(
  string lower_sqr,
  translation_seq,
  register_sequencer,
  upper_seq
);

  layer_db[lower_sqr].lower_sqr
  = lower_sqr;
  layer_db[lower_sqr].translation_seq
  = translation_seq;
  layer_db[lower_sqr].upper_sqr
  = register_sequencer;
  layer_db[lower_sqr].upper_seq
  = upper_seq;
endfunction

  local task start_sequence(
    string sequence_name,
    ovm_sequencer_base sqr;
    ovm_sequence_base seq;

    // Get a sequence from the factory.
    $cast(seq,
      factory.create_object_by_name(
        sequence_name, get_full_name(),
        sequence_name));

    fork
      ovm_sequencer_base m_sqr = sqr;
      ovm_sequence_base m_seq = seq;
      m_seq.start(m_sqr);
    join_none
  endtask

  task run();
  // Start all the bus
  // translation sequences.
  // (the downstream side).
  foreach (layer_db[s])
    start_sequence(
      layer_db[s].translation_seq,
      layer_db[s].lower_sqr);

  // Start all the register sequences.
  // (the upstream side).
  foreach (layer_db[s])
    start_sequence(
      layer_db[s].upper_seq,
      layer_db[s].register_sqr);
  endtask
endclass : register_layer
endpackage

```

APPENDIX 2 – REGISTER LAYERING TRANSLATOR SEQUENCE EXAMPLE

```
virtual class translation_sequence_base
  extends simple_sequence_base;
  `ovm_sequence_utils(
    translation_sequence_base,
    simple_sequencer)

  ovm_register_map rm;

  pure virtual task write(
    bit[31:0]addr, bit[31:0]data);
  pure virtual task read (
    bit[31:0]addr, output bit[31:0]data);

  function new(string name);
    super.new(name);
  endfunction

  ovm_seq_item_pull_port
  #(ovm_register_transaction,
    ovm_register_transaction)
  register_seq_item_port;

  ovm_analysis_port
  #(ovm_register_transaction)
  register_rsp_port;

task body();
  ovm_object o;
  if (m_sequencer.get_config_object(
    "register_map", o))
    $cast(rm, o);
  else
    ovm_report_fatal(
      "translator_sequence",
      "Cannot find 'register_map'");

  register_layer::lookup_ports(
    m_sequencer,
    register_seq_item_port,
    register_rsp_port);

  forever begin
    BV rd_data, addr;
    bit valid_address;
    ovm_register_transaction
    register_req, register_rsp;

    // Get a request from upstream.
    // This is a register transaction.
    register_seq_item_port.get(
      register_req);

    ovm_report_info("translator",
      $sprintf("reg transaction '%s'",
        register_req.convert2string()));

    addr =
      rm.lookup_register_address_by_name(
        register_req.name,
        valid_address);

    // Create the register response
    register_rsp = new();

    if (!valid_address) begin
      ovm_report_error(
        "translator_sequence",
        $sprintf(
          "Register '%s' is not mapped",
          register_req.name));
    end
    else begin
      if (register_req.op ==
        ovm_register_pkg::WRITE)
        write(addr, register_req.data);
      else
        read(addr, rd_data);

      register_rsp.copy_req(
        register_req);
      register_rsp.set_id_info(
        register_req);

      if (register_rsp.op ==
        ovm_register_pkg::READ)
        register_rsp.data = rd_data;
    end

    // Send a response back upstream.
    register_rsp_port.write(
      register_rsp);
  end
endtask
endclass
```