

A general method for SystemC modeling of RISC processors

Authors : Alain Greiner, François Pécheux, Nicolas Pouillon

1. A) General principles
2. B) Generic ISS API
3. C) ISS internal organisation
4. D) Generic cache controller
5. E) CABA modeling
6. F) TLM-T modeling

The goal of the method presented here is to simplify the SystemC modeling of a specific class of embedded processors : The method is well suited to 32 bits RISC processors, with one single instruction issue per cycle, and blocking instruction and data caches.

A) General principles

The method relies on three basic principles :

- The processor core is modeled as a generic ISS (Instruction Set Simulator).
- This ISS is wrapped in appropriate wrappers for several types of simulation models : CABA, TLM-T and PV.
- All processors types use the same generic cache controller.

On one hand, the same ISS is encapsulated in different wrappers to generate several simulation models, corresponding to several abstraction levels: CABA (Cycle-Accurate Bit-Accurate), TLM-T (Transaction Level Models with Time), and PV (Programmer View, untimed). On the other hand, it is possible to use the same wrapper for different types of processor architectures. As illustrated below, all simulation models can be obtained as the cartesian product of the ISS set, by the wrappers set.

	CABA Wrapper	TLM-T Wrapper	PV Wrapper
ISS MIPS3000	CABA Model MIPS	TLM-T Model MIPS	PV Model MIPS
ISS PPC405	CABA Model PPC	TLM-T Model PPC	PV Model PPC
ISS OpenRISC	CABA Model OpenRISC	TLM-T Model OpenRISC	PV Model OpenRISC

The method has been demonstrated for the MIPS3000 and PPC 405 processors, and can be simply extended to the OpenRISC, Sparc, Nios, and MicroBLAZE processors.

This modeling approach supposes that all ISS implement the same generic API (Application Specific Interface), as this API must be independant from both the procesor architecture, and the wrapper type.

The proposed method makes the assumption that the processors use the **VcIXcache** cache controller available in the SoCLib library to interface the VCI interconnect. Such modular approach allows to share the modeling effort of the L1 cache controller. The fonctionnal validation and debug of this component has been a tedious task, and such reuse is probably a good policy. Nevertheless, a clean procedural interface has been defined between the processor core, and the cache controller, and the cache behaviour can be easily modified if required.

Finally this generic approach has been exploited to develop the gdbServer module that is mandatory to help the debug of the multi-tasks application software running on the MP-SoC architectures modeled with SoCLib. This tool can be used for all simulation models compliant with the method described below.

B) Generic ISS API

As explained in the introduction, the modeling method relies on a generic ISS API, usable by any 32 bits RISC processor, and by the three wrappers CABA, TLM-T & PV. The Instruction Set Simulator corresponding to a given processor handles a set of registers defining the processor internal state. The API described below defines a procedural interface to allows the various wrappers to access those registers. The main access function is the **step()** function, that executes one ISS step : For an untimed model (PV wrapper) one step corresponds to one instruction. For a timed model (CABA wrapper or TLM-T wrapper), one step corresponds to one cycle.

- **inline void reset()**

This function reset all registers defining the processor internal state.

- **inline bool isBusy()**

This function is only used by timed wrappers (CABA & TLM-T). In RISC processors, most instructions have a visible latency of one cycle. But some instructions (such as multiplication or division) can have a visible latency larger than one cycle. This function is called by the CABA and TLM-T wrappers before executing one step : If the processor is busy, the wrapper calls the **nullStep()** function. If the processor is available, the wrapper may call the **step()** function to execute one instruction.

- **inline void step()**

This function executes one instruction. All processor internal registers can be modified.

- **inline void nullStep()**

This function performs one internal step of a long instruction.

- **inline void getInstructionRequest (bool & req , enum InsAccessType & type, uint32_t & address)**

This function is used by the wrapper to obtain from the ISS the instruction request parameters. The **req** parameter is true when there is a valid request. The **address** parameter is the instruction address. The **type** parameter can have the values defined below:

```
enum InsAccessType {
    RC ,    // Read Instruction Cached
    RU ,    // Read Instruction Uncached
}
```

- **inline void getDataRequest (bool &req , enum DataAccessType & type, uint32_t & address, uint32_t & wdata)**

This function is used by the wrapper to obtain from the ISS the data request parameters. The **req** parameter is true when there is a valid request. The **address** parameter is the data address, and the **wdata** parameter is the data value to be written. The **type** parameter is defined below :

```
enum DataAccessType {
    RW ,    // Read Word Cached
    RH ,    // Read Half Cached
    RB ,    // Read Byte Cached
    RZ ,    // Cache Line Invalidate
    WW ,    // Write Word
    WH ,    // Write Half
    WB ,    // Write Byte
}
```

```

    SC ,    // Store Conditional Word
    LL ,    // Load Linked Word
}

```

- **inline void setInstruction (bool error, uint32_t ins)**

This function is used by the wrapper to transmit to the ISS, the instruction to be executed (**ins** parameter). In case of exception (bus error), the **error** parameter is set.

- **inline void setDataResponse (bool error, uint32_t rdata)**

This function is used by the wrapper to transmit to the ISS, the response to the data request. In case of a read request, the **rdata** parameter contains the read value. In case of exception (bus error), the **error** parameter is set. In any case, this function must reset the ISS data request.

- **inline void setWriteBerr ()**

This function is used by the wrapper to signal asynchronous bus errors, in case of a write acces, that is non blocking for the processor.

- **inline void setIrq (uint32_t irq)**

This function is used by the wrapper to signal the current value of the interrupt lines. For each processor, the number of interrupt lines must be defined by the ISS variable **n_irq**.

C) ISS internal organisation

As an example, we present the general structure of the MIPS R3000 ISS (chronogram of figure 1). The instruction fetch, instruction decode, and instruction execution are done in one cycle. A specific register **r_npc** is introduced to model the delayed branch mechanism : the instruction following a branch instruction is always executed. The load instructions are executed in two cycles, as those instructions require two cache access (one for the instruction, one for the data). The ISS can issue two simultaneous request for the instruction cache, and the data cache, but those requests are done for different instructions.



The **r_pc** et **r_npc** registers contain respectively the current instruction address, and the next instruction address. The wrapper can obtain the PC content using the **getInstructionRequest()** function, fetch the instruction in the cache (or in memory in case of MISS), and propagate the requested intruction to the ISS using the **setInstruction()** function. The wrapper starts the instruction execution using the **step()** function. The general registers **r_gp**, as well as the **r_mem** registers defining the possible data access, are modified. If, at the end of cycle (i) the **r_mem** registers contain a valid data access, this access will be performed during the next cycle, in parallel with the execution of instruction executed at cycle (i+1).

From an implementation point of view, a specific ISS is implemented by a class **processorIss**. This class inherits the class **genericIss**, that defines the prototypes of the access function presented in section B, (defined as virtual functions).

D) Generic cache controler

The hardware component **VciXcache** is a generic cache controler, that can be used by various processor cores. It contains two separated instruction and data caches, but has a single VCI port to acces the VCI interconnect. The cache line width, and the cache size are defined as independant parameters for the data cache and the instruction

cache. On the processor side, the cache controller can receive two requests at each cycle : one instruction request (read only), and one data request (read or write). Those requests, and the corresponding responses are transmitted through a normalised interface described below. Both instruction and data caches are blocking : the processor is supposed to be frozen in case of MISS (uncached read acces are handled as MISS). Both caches are direct mapping, and the write policy for the data cache is WRITE-THROUGH. The cache controller contains a write buffer supporting up to 8 fposted write requests. In case of successive write requests to contiguous addresses, the cache controller will build a single VCI burst. Therefore, the procesor can be blocked in case of MISS on a read request, but is generally not blocked in case of write request. Finally, in order to garanty a strong ordering memory consistency, the **VciXcache** controler sequencialize the memory accesses, strictly respecting the access ordering defined by the processor on the **VciXcache** interface. As the VCI interconnect does not garanty the in order delivery property, the cache controller waits the VCI response packet corresponding to transaction (n) before sending the VCI command packet corresponding to transaction (n+1).

To communicate with the processor, the CABA model of the **VciXcache** component contains two ports defined below :

```
class IcacheCachePort {
    sc_in<bool> req; // valid request
    sc_in<sc_dt::sc_uint<2> > type; // instruction access type
    sc_in<sc_dt::sc_uint<32> > adr; // instruction address
    sc_out<bool> frz; // frozen processor
    sc_out<sc_dt::sc_uint<32> > ins; // instruction
    sc_out<bool> berr; // bus error
}

class DcacheCachePort {
    sc_in<bool> req; // valid request
    sc_in<sc_dt::sc_uint<4> > type; // data access type
    sc_in<sc_dt::sc_uint<32> > adr; // data address
    sc_in<sc_dt::sc_uint<32> > wdata; // data to be written
    sc_out<bool> frz; // frozen processor
    sc_out<sc_dt::sc_uint<32> > rdata; // read data
    sc_out<bool> berr; // bus error
}
```

E) CABA modeling

The CABA modeling for a complete CPU (processor + cache) is presented in figure 2. The processor ISS is wrapped in the generic CABA wrapper, implemented by the class **IssWrapper**. The class **IssWrapper** contains the member variable **m_iss** representing the processor ISS. The type of the **m_iss** variable - defining the type of the wrapped processor - is specified by the template parameter **iss_t**. The class **IssWrapper** inherit the class **caba::ModuleBase**, that is the basis for all CABA modules.



To communicate with the **VciXcache**, the **IssWrapper** class contains two member variables **p_icache**, of type **IcacheProcessorPort** and **p_dcache**, of type **DcacheProcessorPort**. It contains also the member variable **p_irq**, that is a pointer to an array of ports of type **sc_in<bool>**. This array represents the interrupt ports. The number N of interrupt ports depends on the wrapped processor, an is defined by the **n_irq** member variable of the **iss_t** class.

The SystemC code for the generic CABA wrapper is presented below :

```
template<typename iss_t>
class IssWrapper : Caba::BaseModule
{
public :
```

```

////////// ports //////////
sc_in<bool> *p_irq ;
IcacheProcessorPort p_icache ;
DcacheProcessorPort p_dcache ;
sc_in<bool> p_resetsn ;
sc_in<bool> p_clk ;

////////// constructor //////////
IssWrapper(sc_module_name insname,
           int ident ) :
    BaseModule(insname),
    m_iss(ident)
{
    p_icache("icache") ;
    p_dcache("dcache") ;
    p_resetsn("resetsn") ;
    p_clk("clk") ;
    for (uint32_t i = 0 ; i < iss_t::n_irq ; i++) {
        new(&p_irq[i]) sc_in<bool> ("irq", i) ;
    }
    m_ins_asked = false ;
    m_data_asked = false ;
    SC_METHOD(transition);
    dont_initialize();
    sensitive << p_clk.pos();
    SC_METHOD(genMoore);
    dont_initialize();
    sensitive << p_clk.neg();
}

private :

//////////
iss_t m_iss ;

////////////////////////////////////
void transition()
{
    bool ifrz = p_icache.frz.read() ;
    bool ireq = p_icache.req.read() ;
    bool iberr = p_icache.berr.read() ;

    bool dfrz = p_dcache.frz.read() ;
    bool dberr = p_dcache.berr.read() ;
    bool dreq = p_dcache.req.read() ;

    if ( ! p_resetsn.read() ) {
        m_iss.reset();
        return;
    }

    if ( ireq ) {
        m_iss.setInstruction( iberr, p_icache.ins.read())
    }

    if ( dberr && ( !dreq || dfrz ) ) {
        m_iss.setWriteBerr() ;
    } else if ( dreq ) {
        m_iss.setDataResponse( dberr, p_dcache.rdata.read() ) ;
    }

    if ( m_iss.isBusy() || ifrz || dfrz ) { // Processor frozen or busy
        m_iss.nullStep();
    } else { // Execute one instruction:
        m_iss.step();
    }
}

```

```

        // report interrupts
        uint32_t irqword = 0;
        for ( size_t i=0; i<(size_t)iss_t::n_irq; i++ ) { if (p_irq[i].read()) irqword |= (1<<i); }
        m_iss.setIrq(irqword);

    } // end transition()

//////////
void genMoore()
{
    bool ins_req ;
    enum InsAccessType ins_type ;
    uint32_t ins_address ;
    bool data_req ;
    enum DataAccessType data_type ;
    uint32_t data_address ;
    uint32_t data_wdata ;

    m_iss.getDataRequest( data_req, data_type, data_address, data_wdata ) ;
    m_iss.getInstructionRequest( ins_req, ins_type, ins_address ) ;

    m_ins_asked = ins_req ;
    m_data_asked = data_req ;

    p_icache.req = ins_req ;
    p_icache.type = ins_type ;
    p_icache.adr = ins_address;
    p_dcache_req = data_req ;
    p_dcache_type = data_type ;
    p_dcache.adr = data_address;
    p_dcache.wdata = data_wdata;
} // end genMoore

```

F) TLM-T modeling

The TLM-T modeling for a complete CPU (processor + cache) is presented in figure 3. To increase the simulation speed, the TLM-T wrapper is the cache controller itself, and it is implemented as the class **VciXcache**. This class contains the SC_THREAD **execLoop()** implementing the PDES process, and the **m_time** member variable implementing the associated local clock. The class **VciXcache** inherit the class **tlmt::ModuleBase**, that is the basis for all TLM-T modules. This class contains the member variable **m_iss** representing the processor ISS. The type of the **m_iss** variable is defined by the template parameter **iss_t**.



The class **VciXcache** contain a member variable **p_vci**, of type **VciInitPort**, to send VCI command packets, and receive VCI response packets. This class contains also the member variable **p_irq**, that is a pointer to an array of ports of type **SynchroInPort**. This array represents the interrupt ports. The number N of interrupt ports depends on the wrapped processor, and is defined by the **n_irq** member variable of the **iss_t** class.

The **execLoop()** function contains an infinite loop. One iteration in this loop corresponds to one cycle for the local clock, (or more, as the thread is suspended in case of MISS).

The cache behavior is specifically described by the **cacheAccess()** method, that is a member function of the class **VciXcache**, and is called by **execLoop()** at each cycle. This function has the following prototype :

```

void cacheAccess( icache_request_t  *ireq,
                  dcache_request_t  *dreq,
                  xcache_response_t  *rsp)

```

The **icache_request_t**, **dcache_request_t**, and **xcache_response_t** classes represent the instruction and data requests, and the cache response respectively :

```
class icache_request_t {
bool    valid ;
enum InsAccessType  type ;
uint32_t  address ;
}
class dcache_request_t {
bool    valid ;
enum DataAccessType  type ;
uint32_t  address ;
uint32_t  wdata ;
}
class xcache_response_t {
bool  iber ;
uint32_t  instruction ;
bool  dber ;
uint32_t  rdata ;
}
```

The **cacheAccess()** function determines the actions to be done :

- In case of data or instruction MISS, the **cacheAccess()** function sends the proper VCI command packet on the **p_vci** port, and the **exedcLoop()** thread is suspended.
- In case of data write, the the **cacheAccess()** function sends the proper VCI command packet on the **p_vci** port, but the **exedcLoop()** thread is not suspended.

At each iteration in the execution loop, the **cacheAccess()** method updates the local clock (variable **m_time**) :

- The local time is simply incremented by one cycle, if the cache controller is able to answer immediately to the processor requests.
- The local time is updated using the date contained in the VCI response packet in case of MISS.

The SystemC TLM-T model for the **VciXcache** module is presented below :

```
template<typename iss_t, typename vci_param>
class VciXcache<iss_t> : tlmt::BaseModule {

public :
//////// ports //////////
VciInitiatorPort<vci_param>  p_vci ;
SynchroInPort  * p_irq ;

//////// constructor /////
VciXcache (sc_module_name  name,
           uint32_t  initiatorIndex,
           uint32_t  processorIdent,
           uint32_t  lookahead,
           uint32_t  dcache_nlines,
           uint32_t  dcache_nwords,
           uint32_t  icache_nlines,
           uint32_t  icache_nwords)
    BaseModule(name),
    m_iss(processorIdent),
    m_time(0)
    {
        p_vci (« vci », this, &VciXcache::rspReceived, &m_time) ,
        for (uint32_t i = 0 ; i < iss_t::n_irq ; i++) {
            new(&p_irq[i])SynchroInPort ("irq", i, this, &VciXcache::irqReceived) ;
        }
        m_initiator_index = initiatorIndex ;
    }
```

```

        m_counter = 0 ;
        m_lookahead = lookahead ;
        m_icache_nlines = icache_nlines ;
        mi_icache_nwords = icache_nwords ;
        m_dcache_nlines = dcache_nlines ;
        m_dcache_nwords = dcache_nwords ;
        SC_THREAD(execLoop) ;
    } // end constructor

private :
////////// member variables
tlmt_time m_time ;
sc_event m_rsp_received ;
iss_t m_iss ;
uint32_t m_rsp_error ;
uint32_t m_dcache_nlines ;
uint32_t m_dcache_nwords ;
uint32_t m_icache_nlines ;
uint32_t m_icache_nwords ;
uint32_t m_initiator_index ;
uint32_t m_lookahead ;
uint32_t m_counter ;
bool m_irqpending[iss_t::n_irq] ;
uint32_t m_irqtime[iss_t::n_irq] ;
vci_cmd_t m_cmd ;
////////// thread
void execLoop()
{
    icache_request_t icache_req ; // The Icache request
    dcache_request_t dcache_req ; // The Dcache request
    xcachel_response_t xcachel_rsp ; // The Xcache response
    uint32_t irqword ;
    while(1) {
        // execute one cycle
        if (m_iss.isBusy() {
            m_iss.nullStep() ;
        } else {
            ////////// cache access
            m_iss.getInstructionRequest(icache_req.valid,
                                       icache_req.type,
                                       icache_req.address) ;
            m_iss.getDataRequest(dcache_req.valid,
                                dcache_req.type,
                                dcache_req.address,
                                dcache_req.wdata)
            xcachelAccess(&icache_req, &dcache_req, &xcachel_rsp) ;
            if ( icache_req.valid ) m_iss.setInstruction(xcachel_rsp.iber, xcachel_rsp.instruction)
            if ( dcache_req.valid ) m_iss.setDataResponse(xcachel_rsp.dber, xcachel_rsp.rdata) ;
            ////////// handling interrupts
            irqword = 0 ;
            for ( size_t i = 0 ; i < iss_t::n_irq ; i++) {
                if( m_irqpending[i] && m_irqtime[i] <= get_time()) irqword |= (1<<i);
            }
            m_iss.setIrq(irqword) ;
            ////////// handling asynchronous bus error
            if ( m_write_error ) {
                setWriteBerr() ;
                m_write_error = false ;
            }
            m_iss.step() ;
        } // end cycle
        // lookahead management
        m_counter++ ;
        if (m_counter >= m_lookahead) {
            m_counter = 0 ;
            wait(SC_ZERO_TIME) ;
        }
    }
}

```

```

        } // end if lookahead
    } // end while(1)
} // end execLoop()

////////////////////////////////////
void cacheAccess(icache_request_t ireq,
                 dcache_request_t dreq,
                 xcache_response_t rsp)
{
    ...
} // end cacheAccess()

////////////////////////////////////
void rspReceived(vci_rsp_t rsp,
                uint32_t time)
{
    if ( rsp.cmd == VCI_CMD_write ) { // asynchronous bus error signaling
        m_write_error = ( rsp.error != 0 ) ;
    }
    if ( rsp.cmd == VCI_CMD_READ ) { // time update & cache activation
        m_time.updateTime( time + rsp.length ) ;
        notify( m_rsp_received ) ;
    }
} // end rspReceived()

////////////////////////////////////
void irqReceived(bool val,
                uint32_t time,
                size_t index)
{
    m_irqpending[index] = val ;
    m_irqtime[p_irq[index]] = time ;
} // end irqReceived()
} // end class VciXcache

```