

SHIM: A Language for Hardware/Software Integration

Stephen A. Edwards

*Department of Computer Science, Columbia University
1214 Amsterdam Avenue, New York, New York, 10027*

Abstract

Virtually every system designed today is an amalgam of hardware and software. Unfortunately, software and circuits that communicate across the hardware/software boundary are tedious and error-prone to create. This suggests a more automatic way to synthesize them.

This paper presents the SHIM language, which combines imperative C-like semantics for software and RTL-like semantics for hardware to allow a unified description of hardware/software systems. Hardware processes and software functions communicate through shared variables, hardware for which is automatically synthesized by the SHIM compiler, which generates C and synthesizable VHDL.

I demonstrate the effectiveness of the language by re-implementing an I²C bus controller. The SHIM source is half the size of an equivalent manual implementation, slightly faster, and has a smaller memory footprint. Partial and complete hardware implementations in SHIM are also presented, showing that SHIM is succinct and effective.

Key words: Hardware-software codesign, language design, interface synthesis

1 Introduction

As integrated circuit technology advances relentlessly, the size and complexity of a typical design continues to spiral upward. As always, managing complexity is the designer's greatest challenge. The design must be right the first time and be completed faster than before. While validation methods such as simulation and formal verification work well to discover mistakes, moving to higher levels of abstraction, such as from the gate level to the register transfer

¹ Edwards is supported by an NSF CAREER award, a grant from Intel corporation, an award from the SRC, and by New York State's NYSTAR program.

² Email: sedwards@cs.columbia.edu

level, is more effective because it helps designers to avoid mistakes in the first place. In this paper, I describe a language that raises the abstraction level for hardware/software systems.

My intention with SHIM (Software/Hardware Integration Medium), the language I propose here, is to provide seamless communication between hardware and software modules. It arose in part from observing beginning design students tackle combined hardware/software systems: they understood C well, and could learn VHDL, but had difficulty making the two worlds communicate.

Rather than propose a completely new semantics for SHIM, I chose to integrate the two well-known, well-established semantics typically used to design hardware and software systems: C-like imperative semantics (single-threaded) for the software and register-transfer level (concurrent) semantics for the hardware. By itself, each is deterministic, but their combination turns out not to be because the relative speeds of the hardware and software are not specified. I adopted this model because it closely reflects current industrial practice, in which validation must often be performed with a cycle-accurate processor simulator running in concert with a cycle-accurate hardware simulation.

While it would be an interesting exercise to consider other, perhaps deterministic, semantics, the objective of this work was to explore the consequences of adopting commonly-used techniques.

The syntax of SHIM mimics C, which I chose both for its familiarity (C, C++, Java, and C# programmers all know it well) and its succinctness. Although from a technical standpoint, language syntax is nearly irrelevant, it can be a significant barrier to user adoption of a language.

The SHIM compiler takes a single specification—a collection of software functions and hardware processes—and generates both C and synthesizable VHDL source. No attempt is made to do automatic partitioning between the two domains: it is the user’s responsibility to mark everything as hardware, software, or shared. This is partially to simplify the compiler, but the more serious issue is the very different semantics of the two domains.

Hardware and software in SHIM must be coded differently; although the syntax is similar, is usually not possible to turn a software function into hardware by writing `hw` in front of it. This is perhaps unfortunate, but I believe inescapable. When I began this project, I wanted to make switching functionality between the two domains this easy, but the semantics of RTL and software are just too different: how many clock cycles should a piece of software code take, and how is concurrency implemented in software? Hardware-like software and software-like hardware languages (e.g., Esterel [2] and Handel-C [4]) have been proposed, but all deviate significantly from software and RTL semantics and none are widely accepted. See my review of C-like hardware languages for more details [9,10].

Aside from minor improvements in RTL syntax (SHIM’s RTL can be half as verbose as the equivalent VHDL), SHIM’s major contribution is the automatic synthesis of communication between the software and hardware components

```

module timer {
    shared uint:32 counter;

    hw void count() {
        counter = counter + 1;
    }

    out void reset_timer() {
        counter = 0;
    }

    out uint get_time() {
        return counter;
    }
}

```

(a)

```

#ifdef _TIMER_H
#define _TIMER_H
extern void reset_timer(void);
extern unsigned int get_time(void);
#endif /* _TIMER_H */

```

(b)

```

signal counter : UNSIGNED(31 downto 0);

count : process(Clk)
begin
    if Clk'event and Clk = '1' then
        counter <= counter + 1;
        if cs1 = '1' and RNW = '0' then
            if offset = 0 then
                counter <= DBus;
            end if; end if; end if;
        end process count;

read_shared_variables : process(Clk)
begin
    if Clk'event and Clk = '1' then
        if cs1 = '1' and RNW = '1' then
            DBus_out <= read_data;
        else DBus_out <= "0"; end if;
        if offset = 0 then
            read_data <= counter;
        end if; end if;
    end process read_shared_variables;

```

```

#include "timer.h"
#include "xio.h"
#define IO_BASE 0xfeff0200
#define counter (IO_BASE + 0x0)

void reset_timer() {
    XIo_Out32(counter, 0);
}

unsigned int get_time() {
    return XIo_In32(counter);
}

```

(c)

(d)

Fig. 1. (a) A simple SHIM program: a hardware timer. (b) The C header file generated by the SHIM compiler. (c) Generated C source. (d) A fragment of generated VHDL. These two processes implement the SHIM *count* process and the ability to read the value of *count* from software. In the first process, *count* is incremented or read from the bus if chip select and write are true. The second process places the value of *count* on the bus if chip select and read are true.

based on a shared memory model. While algorithmically simple, such synthesis simplifies the designer's task, avoids errors, and improves code portability.

Figure 1a shows a simple SHIM program that represents a hardware timer. It manages the shared variable `counter`, whose value is stored in the hardware

but can also be read and written by the two software functions `reset_timer` and `get_time`. From this description, the compiler produces a C header file (Figure 1b) that describes the external software interface for the module, a C source file containing code implementing `reset_timer` and `get_time` (Figure 1c), and VHDL for a peripheral that attaches to a processor bus and implements the `count` process and circuitry that allows the shared variable to be read and written from the C program (Figure 1d).

Much of the generated code in Figure 1 is specific to the target for which SHIM currently generates code—a Xilinx FPGA containing a Microblaze soft processor core communicating with peripherals over an on-chip peripheral bus (OPB)—but illustrates the code the SHIM compiler generates. The C code in Figure 1c uses two Xilinx-supplied macros—`XIo_Out32` and `XIo_In32`—that write to and read from arbitrary memory locations to communicate with the synthesized hardware module. The `#defines` above them contain the addresses to which the SHIM compiler has assigned the timer peripheral.

The VHDL code in Figure 1d is only a fragment of the complete implementation. Not shown is the chip-select logic that generates `cs1` whenever the processor is accessing one of this peripheral’s memory addresses; the `offset` variable contains the low-order address bits that indicate which shared variable is being written (there may be many in general; this example only has one). The `RNW` signal indicates whether the processor is reading from or writing to a peripheral, `DBus` contains data coming from the processor, and `DBus_out` sends data to the processor.

This example illustrates how SHIM simplifies coding hardware/software systems: the SHIM source is only fifteen lines, but the SHIM compiler generates twenty lines of C and nearly 100 lines of VHDL from it, code a designer would have otherwise had to write manually. The fractional improvement is high for such a small example because most of the code is related to the bus interface, but large examples remain smaller and easier to code.

As its name suggests, SHIM is designed for creating “glue” that connects hardware and software. It specifically does not try to be a general-purpose software language or a full hardware description language. Instead, it provides facilities for connecting subsystems written in hardware and software and external interfaces: externally-visible functions and variables in software, ports in hardware. Figure 4 illustrates how these facilities can be used.

SHIM can also be thought of as a language for simultaneously writing peripherals and their device drivers. As such, it is applicable to systems where hardware can be customized, such as ASICs with processor cores; field-programmable gate arrays, especially those with processors such as the Virtex II Pro; and more programmable SoCs that consist of multiple processor cores, hard peripherals, and a substantial amount of programmable logic. This latter class of chip, a platform with a mix of programmable hardware and software, seems a likely architecture of the future since the cost of designing complete chips is skyrocketing; it seems more likely that standard programmable

platforms will grow more popular. When the hardware is provided and immutable, the NDL language [8] would be better-suited.

SHIM currently generates code (hardware and software) for the Xilinx Microblaze soft processor core driving IBM’s CoreConnect On-Chip Peripheral Bus (OPB). I chose this configuration because we have target boards from XESS and a development environment from Xilinx, but there is nothing OPB-specific about the SHIM language. The synthesis code is about 1/6th of the compiler and could be easily retargeted to a different processor and bus.

2 Related Work

Unlike other hardware/software codesign systems, SHIM uses imperative C semantics for software and RTL for hardware. Polis [1] describes its systems with communicating extended finite-state machines. COSMOS (Jerraya et al. [13]), COSYMA (Ernst et al. [11]), and CoWare (Bolsens et al. [3]) all use concurrently-running processes communicating through remote procedure calls (RPC). FIFO-based communication among concurrently-running processes is another choice (see, e.g., Gupta and De Micheli [12]).

Although such higher-level semantics enable automatic hardware/software partitioning (a focus of earlier work), they raise efficiency issues. While RPC is natural in software, it seems overly sequential for hardware. Furthermore, some techniques employ high-level hardware synthesis, which industry has largely rejected due to efficiency concerns. COSYMA describes a hardware/software system using a C-like imperative language, partitions it, and passes certain processes to high-level synthesis.

Most other work proposes synthesizing communication mechanisms far more complicated than shared memory. Jerraya et al. [5] propose synthesizing wrappers. CoWare layers protocols. It is unclear whether these approaches are general enough and produce efficient hardware.

SHIM synthesizes bus-based communication, but is not wedded to it. Chou et al. [6,7] target microcontrollers communicating through a limited number of I/O pins.

Other languages, such as Mérillon et al.’s Devil [15] (generates C macros for communicating with peripherals), Thibault et al.’s GAL [16] (generates graphics adapter drivers), and Conway et al.’s NDL [8] (synthesizes Unix device drivers), focus only on synthesizing software and assume the hardware is given.

Lavagno and Sentovich’s ECL [14], which combines Esterel-like synchrony [2] with imperative C, partially inspired SHIM. Like SHIM, their compiler uses simple rules to split a program into the two domains, but they interleave Esterel and C at the statement (instead of function) level and they do not target hardware/software systems.

3 The SHIM Language

SHIM was designed to be simple. A SHIM program is a module containing global variables and functions. Each variable is either unmarked, marked `hw`, or marked `shared` indicating it is to be visible in software only, hardware only, or both. The translation of hardware- and software-only variables is straightforward; the state of a shared variable is held in hardware and the compiler synthesizes circuitry that allows it to be read and written from software. Each variable can also be marked `in` or `out`, indicating its value comes from outside the module or is visible outside (variables are only module-visible by default).

Functions may either be unmarked (indicating software) or marked `hw`. The translation of a software function is nearly one-to-one. A hardware function becomes a concurrently-running synchronous hardware process clocked by the bus to which the synthesized peripheral is attached. Currently, all hardware functions take no arguments, do not return a value, and may not be explicitly called (they are implicitly called once per clock cycle).

SHIM supports bit vectors, arrays of bit vectors, and string literals. Bit vectors can be manipulated as integers and are either signed (`int`) or unsigned (`uint`). The width of each bit vector may be specified or may be omitted (it defaults to a standard value). Array dimensions must be given.

The body of a SHIM function contains the familiar set of C statements, `if`, `switch`, `for`, `while`, and `return`; expressions; and local variable declarations. Expressions follow the usual C syntax with a few extensions: applying the array index operator to a bit vector returns a bit, e.g., `a[1]` is the second least-significant bit of vector `a`. Looping statements are not allowed in hardware processes.

Software functions have the usual sequential semantics; there is a single program counter; bit vectors are call-by-value, and arrays are call-by-reference.

Hardware processes have RTL semantics. Every hardware process is effectively invoked once per clock cycle. Each shared or hardware-only variable may be written by a single process only, a syntactic constraint common in RTL, but may be read by multiple processes. Outputs are latched, meaning that if a process writes a variable, the new value can be seen by other processes only in the next cycle, but the same process can see the new value in the same cycle.

When software-only variables have an initial expression, the value is assigned once when the variable enters scope (e.g., when the program starts for global variables, and when control enters the enclosing block for local variables) and stored. Software-only variables marked `const` must be given an initial value that can be determined at compile-time and may not be written.

Hardware-only variables may also be assigned to an expression, but their semantics differ: hardware variables assigned to an expression always take the value of that expression (i.e., are effectively continuous assignments) and may not be assigned in other processes. Such variables generate combinational

logic and conceptually execute after all the processes have executed for the cycle and updated their outputs. That is, they can be used to communicate among processes, but only across clock-cycle boundaries.

4 Compiling SHIM

The SHIM compiler comprises some 3000 lines of OCAML code. The largest single module (about 500 lines) performs the interface synthesis. Much of it comes from explicit construction of SHIM code for the bus controller.

The compiler's structure is typical: an automatically-generated scanner and parser generates an abstract syntax tree, which is sent through a static semantic analyzer that resolves names and types and dismantles certain constructs to produce an AST-like intermediate representation.

The interface generator takes the interface-agnostic IR from the dismantler, enumerates the shared variables (a simple walk through the symbol table), and assigns each an I/O address. It transforms the body of each software function by changing reads and writes of shared variables to the appropriate I/O function call.

Interface synthesis for hardware is more complicated. At the end of each process, the compiler adds code that reads from the bus each shared variable written in that process. This generates at most one read of each shared variable since each such variable may be written by at most one hardware process (software functions have no such restriction). Furthermore, placing the read at the end of the process means a write from the software domain takes precedence over a write from hardware. I chose these semantics because hardware writes every cycle are a common idiom (e.g., the `count` variable in Figure 1). Users who dislike software taking precedence over hardware can use separate shared variables for each communication direction.

Two additional hardware processes are synthesized: one reads all the unwritten shared variables from the bus (e.g., those shared variables that are only written from software, never from hardware). The other copies each shared variable to the bus when requested (i.e., when software generates a read cycle to an I/O address).

Finally, the interface synthesis procedure adds interface ports for all the bus signals and processes that decode the address bus and speak the bus protocol. After this, the IR contains all the necessary interface code.

Once the interface synthesis procedure runs, C source, C header, or VHDL source is generated. Syntax-directed translators transform the IR into simple ASTs for C or VHDL, which are then pretty-printed. Using another intermediate representation guarantees syntactically-correct output.

Example	SHIM	C	VHDL
I ² C Software (by hand)		259	133
I ² C Software	171	175	136
I ² C Bit-level	283	173	337
I ² C Byte-level Receive	299	163	358
I ² C Byte-level Send/Receive	323	116	344
Timer	15	20	98

The first line is a reference implementation of an I²C bus controller written without the aid of SHIM. The next four lines list SHIM implementations with the same functionality with an increasing fraction of the system in hardware. The final line lists statistics for the simple example in Figure 1. For all but the first line, the SHIM column represents the number of lines written by the designer; the C and VHDL columns are lines of code generated by the SHIM compiler.

Table 1
Lines of code for various examples.

5 An Example: An I²C Bus Controller

To test SHIM, I re-implemented a simple hardware/software interface with it. In an embedded systems class, we use XSB-300E boards from XESS Corporation that contain, among other things, a Xilinx Spartan IIE FPGA and a Philips SAA7114H video decoder. An I²C bus connecting the two can write to the 7114's many configuration registers. The I²C bus is a low-speed two-wire clocked bus designed for exactly such an application, and is very lenient about timing. The bus has no maximum delays between transitions so it can be controlled completely from software.

One of my students (Marcio Buss) implemented a simple all-software I²C bus controller to program the 7114. Its only task is to send a series of commands to write over eighty configuration registers. Marcio spent a few days writing its 259 lines of C and 133 lines of VHDL. The VHDL interfaces with the On-chip Peripheral Bus and provides two shared variables that control the two I²C bus pins. The C program handles the bus protocol.

Table 1 shows how the handwritten implementation of the I²C bus controller compares to designs written in SHIM. The software version implements the protocol in software; the hardware does little more than control the pair of tri-state drivers called for by the bus interface. The other versions of the I²C controller do more in hardware: the first implements the bit-level protocol in hardware (each byte is sequenced by software), the second moves the complete receive functionality to hardware, and the third performs both send and receive in hardware (bytes are still sequenced in software).


```

shared out bool SCL;    // I2C clock
shared out bool SDA;    // I2C data out
shared out bool SDA_oe; // Output enable for data
shared bool SDA_data;   // I2C data in

void send(uint:8 byte) {
    SDA_oe = 0; delay();
    for (int i = 7 ; i >= 0 ; i = i - 1) {
        SDA = (byte & 0x80) >> 7; delay();
        SCL = 1; byte = byte << 1; delay();
        SCL = 0; delay();
    }
    SDA_oe = 1; delay();
    SCL = 1; delay();
    bool acknowledge_received = SDA_data;
    if (!acknowledge_received)
        xio.print("Acknowledge not received\r\n");
    delay();
    SCL = 0; delay();
}

```

Fig. 2. A fragment of the SHIM code in the all-software version of the I²C protocol that sends a single byte to a slave and looks for an acknowledgement. Each output bit is placed on the data wire and the clock is toggled, then the data wire is set to read, the clock is toggled, and the acknowledge bit from the slave is read.

The first two lines of this table are most telling: Marcio wrote $259 + 133 = 392$ lines of C and VHDL to implement this controller, while I was able to achieve the same functionality in 171 lines of SHIM, a reduction of over 55%.

Figure 2 shows a fragment of the all-software SHIM implementation of the I²C controller. Controls for the two I²C wires, SDA and SCL, are shared one-bit variables, and the *send* function simply toggles them appropriately.

I then implemented more of this controller in hardware using SHIM, the sizes of which I report in the next three lines of the table. By design, changing the implementation of a function from software to hardware is not as simple as marking it as *hw* (SHIM hardware uses RTL semantics, which are very different from C's). The main difference is that sequencing must be coded as a state machine using, say, a *switch* statement (see Figure 3). This is the main source of the additional 152 lines of SHIM.

Figure 3 is a hardware fragment similar to the software of Figure 2. It toggles the clock and data lines to send a single byte of data. Most of the code is now devoted to sequencing since traditional RTL, and by design SHIM, requires all state machines' next-state functions to be explicit. Additional complexity comes from synchronizing with software, done here with a four-phase handshake. The controller raises *ready* in the IDLE state to indicate

```

stared uint:8 sreg;           // Send/receive shift register
shared uint:5 state;         // Controller state
shared bool ready;          // true => controller idling
shared uint:3 command;       // Command for the controller
shared const uint:3 IDLE    = 0; // Commands
shared const uint:3 SEND    = 2;

hw void controller() {
  const uint:5 IDLE = 0;  const uint:5 SEND1 = 5;
  const uint:5 SEND2 = 6;  const uint:5 IDLE0 = 24;
  uint:3 bit_counter;

  if (reset) state = IDLE;
  if (i2c_clock) {
    ready = 0;
    switch (state) {
    case IDLE:
      ready = 1;
      switch (command) {
      case START:  state = START1;  break;
      case SEND:   state = SEND1;   break;
      case RECEIVE: state = RECV1;   break;
      case STOP:   state = STOP1;   break;
      default:     state = IDLE;    break;
      }
      break;
    case SEND1: SDA_oe = 0; bit_counter = 0;  state = SEND2; break;
    case SEND2: SDA = sreg[7];                state = SEND3; break;
    case SEND3: SCL = 1; sreg = sreg << 1; bit_counter = bit_counter - 1;
                                                         state = SEND4; break;
    case SEND4: SCL = 0; if (bit_counter == 0) state = SEND5;
                                                         else state = SEND2; break;
    // Receive Acknowledge
    case SEND5: SDA_oe = 1;                    state = SEND6; break;
    case SEND6: SCL = 1;                      state = SEND7; break;
    case SEND7: acknowledge_received = SDA_in; state = SEND8; break;
    case SEND8: SCL = 0;                      state = IDLE0; break;
    case IDLE0: if (command == IDLE_BIT)      state = IDLE;
                                                         else state = IDLE0; break;
      }
    }
  }
}

void send(uint:8 byte) {
  sreg = byte;
  command = SEND_BIT; while (ready) ;
  command = IDLE_BIT; while (!ready) ;
}

```

Fig. 3. A fragment of the all-hardware implementation of the I²C bus controller in SHIM. This shows part of the main hardware state machine responsible for toggling the I²C clock and data lines to send a byte (cf. Figure 2) and the software function that invokes it with a handshake.

it is ready to accept a command. The software then writes a command such as `SEND` into the `command` variable, which steps the state machine through a state sequence ending with `IDLE0`. In this state, the controller waits until it receives `IDLE` from the software, which sends it back to the `IDLE` state. The `send` function at the bottom of Figure 3 performs the software half of this handshake.

The SHIM version of a controller can be faster than its handwritten equivalent. In the handwritten version of the I²C controller, which was not written for efficiency, Marcio chose to pack the four control bits for the SDA and SCL lines into a single I/O location. While this simplifies address decoding hardware, it requires a read-modify-write operation to change a single bit from software. The SHIM-generated code is more efficient since it can modify each control bit individually.

The size of the SHIM-generated C code can also be superior to handwritten code. The object file for the all-software version of the I²C controller generated by SHIM is only 2106 bytes. By contrast, the handwritten C is over twice the size (4370 bytes). The difference is due to additional function calls and read-modify-write operations.

Figure 4 illustrates how SHIM integrates existing software components through C’s usual function-call mechanism and hardware components through ports (wires). The fragment from the I²C bus controller in Figure 4a shows how ports are defined in the hardware model (variables marked `hw out` and `hw in` are output and input ports on the synthesized VHDL model respectively). For the moment, the actual integration takes place in a manually-written VHDL wrapper (Figure 4b) that simply instantiates both the I²C controller core and a couple of I/O pads (the `IOBUF_F_12s`) and connects them accordingly. It would be technically easy to automatically generate such a wrapper for connecting together SHIM-generated hardware and foreign IP blocks, although this has not yet been done.

6 Conclusions

This paper presented the SHIM language and its compiler. SHIM fuses two widely-accepted computational models—single-threaded imperative software and register-transfer-level hardware—to allow hardware/software systems to be written in a unified language. Declaring a SHIM variable *shared* allows it to be read and written from both hardware and software; the SHIM compiler synthesizes this interface.

I demonstrated the effectiveness of SHIM on an example—an I²C bus controller—and showed that it can reduce by half the number of lines necessary to describe such a system. The savings comes mostly from the automatic synthesis of a bus interface.

SHIM insists that the designer mark each process or function as being software or hardware. One reviewer wondered whether a modeling style could

```

opb_i2ccontroller_0_i : i2chw_core
GENERIC MAP (
    c_baseaddr => X"feff0200",
    c_highaddr => X"feff02ff",
    c_opb_awidth => 32,
    c_opb_dwidth => 32)
PORT MAP (
    OPB_Clk => OPB_Clk,
    OPB_ABus => OPB_ABus,
    OPB_BE => OPB_BE,
    OPB_DBus => OPB_DBus,
    OPB_RNW => OPB_RNW,
    OPB_Rst => OPB_Rst,
    OPB_select => OPB_select,
    OPB_seqaddr => OPB_seqaddr,
    Sln_DBus => Sln_DBus,
    Sln_errAck => Sln_errAck,
    Sln_retry => Sln_retry,
    Sln_toutSup => Sln_toutSup,
    Sln_xferAck => Sln_xferAck,
    SDA => SDA,
    SDA_oe => SDA_oe,
    SDA_in => SDA_in,
    SCL => SCL,
    SCL_oe => SCL_oe
);

sda_pad : IOBUF_F_12 port map (
    I => SDA,
    IO => VID_I2C_SDA,
    O => SDA_in,
    T => SDA_oe
);

scl_pad : IOBUF_F_12 port map (
    I => SCL,
    IO => VID_I2C_SCL,
    O => open,
    T => SCL_oe
);

```

(a)

(b)

Fig. 4. Interfacing SHIM with existing components (IP blocks). (a) The interface to the tri-state drivers (pins) used by the I²C bus controller and a function that calls Xilinx library functions. (b) The (handwritten) VHDL code that connects the bus controller to both the bus and these pins.

allow code to be easily migrated from software to hardware. The answer is a qualified yes, but the style is quite restrictive. The hardware aspects of the language are largely a subset of the software aspects, so writing software in a hardware-friendly way (e.g., prohibit loops and function calls) would make it easily migrated to software, but this is fairly restrictive.

Along the same lines, another reviewer asked whether the SHIM compiler could perform architectural exploration instead of insisting the designer prescribe the hardware/software boundary. While such exploration is certainly useful, SHIM was designed primarily to make it easier to get complex designs functionally correct. Functional correctness is less of an issue during architectural exploration; higher-level languages would be more appropriate.

I believe SHIM is successful in its aim to fuse two computational models but it raises the question of whether these two models are the best choices. A particularly glaring issue is that SHIM models are not easy to simulate. This is due to the models themselves: the two domains run asynchronously and while the hardware is timed, the software effectively is not, meaning that the behavior of the system may be nondeterministic or at least very difficult to predict without careful modeling of software timing, such as by using an instruction-set simulator. The timer example in Figure 1 is perhaps the simplest example illustrating this problem: the hardware effectively counts the number of clock cycles between calls of `reset_timer` and `get_time`, which is a complicated function of the processor and C compiler used to implement the system. Of course, such a performance timer is often desired for analyzing software, but more frequently the behavior of hardware/software systems is meant to be timing independent.

My feeling is that the shared-variable model of hardware/software communication, while the de facto standard, is too flexible and more complicated, but timing-independent protocols are almost always implemented on top of it. For example, a peripheral usually provides status registers and interrupts that inform the software when it is ready for the next command, and expects that software will obey its protocol. In implementing the I²C bus controller described in Section 5, I implemented a four-phase handshake protocol to ensure that the software waited for the hardware to complete its task before starting the next. Although this works and is robust in practice, it can be error-prone and is certainly more verbose than it needs to be.

Part of the problem is that the hardware/software boundary represents true parallelism and concurrency. Peripherals run truly independently from their processors (in most cases, that is the point of a peripheral), and so are truly concurrent systems posing all the classical problems such as races and deadlocks. However, since hardware/software systems are neither completely hardware nor software, using classical software techniques such as semaphores and critical regions or the standard hardware technique of a global clock seems inappropriate, although not impossible.

I leave this question for future work: what is an appropriate general, timing-independent model of computation for hardware/software systems? Single-threaded imperative software and RTL hardware with shared variables, such as in SHIM, is reasonably general, but is too low-level, easily nondeterministic, and error-prone. There must be something higher-level that avoids tedious manual protocol implementation and ensures correctness.

References

- [1] Balarin, F., P. Giusto, A. Jurecska, C. Passerone, E. Sentovich, B. Tabbara, M. Chiodo, H. Hsieh, L. Lavagno, A. Sangiovanni-Vincentelli and K. Suzuki, “Hardware-Software Co-Design of Embedded Systems: The POLIS Approach,” Kluwer, Boston, Massachusetts, 1997.
- [2] Berry, G. and G. Gonthier, *The Esterel synchronous programming language: Design, semantics, implementation*, Science of Computer Programming **19** (1992), pp. 87–152.
- [3] Bolsens, I., H. J. De Man, B. Lin, K. Van Rompaey, S. Vercauteren and D. Verkest, *Hardware/software co-design of digital telecommunication systems*, Proceedings of the IEEE **85** (1997), pp. 391–418.
- [4] Celoxica, “Handel-C Language Reference Manual,” (2003), rM-1003-4.0.
URL <http://www.celoxica.com>
- [5] Cesário, W., A. Baghdadi, L. Gauthier, D. Lyonard, G. Nicolescu, Y. Paviot, S. Yoo, A. A. Jerraya and M. Diaz-Nava, *Component-based design approach for multicore SoCs*, in: *Proceedings of the 39th Design Automation Conference*, New Orleans, Louisiana, 2002, pp. 789–794.
- [6] Chou, P., R. B. Ortega and G. Borriello, *Synthesis of the hardware/software interface in microcontroller-based systems*, in: *Proceedings of the IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, San Jose, California, 1992, pp. 488–495.
- [7] Chou, P., R. B. Ortega and G. Borriello, *Interface co-synthesis techniques for embedded systems*, in: *Proceedings of the IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, San Jose, California, 1995, pp. 280–287.
- [8] Conway, C. L. and S. A. Edwards, *NDL: a domain-specific language for device drivers*, in: *Proceedings of Languages, Compilers, and Tools for Embedded Systems (LCTES)*, Washington, DC, 2004, pp. 30–36.
- [9] Edwards, S. A., *The challenges of hardware synthesis from C-like languages*, in: *Proceedings of the International Workshop on Logic Synthesis (IWLS)*, Temecula, California, 2004, pp. 509–516.
- [10] Edwards, S. A., *The challenges of hardware synthesis from C-like languages*, in: *Proceedings of Design, Automation, and Test in Europe (DATE)*, Munich, Germany, 2005, pp. 66–67.
- [11] Ernst, R., J. Henkel, T. Benner, W. Ye, U. Holtmann, D. Herrmann and M. Trawny, *The COSYMA environment for hardware/software cosynthesis of small embedded systems*, Microprocessors and Microsystems **20** (1996), pp. 159–166.
- [12] Gupta, R. K. and G. De Micheli, *Hardware-software cosynthesis for digital systems*, IEEE Design & Test of Computers **10** (1993), pp. 29–41.

- [13] Ismail, T. B., M. Abid and A. Jerraya, *COSMOS: A codesign approach for communicating systems*, in: *Proceedings of the 3rd International Workshop on Hardware/software Co-Design*, Grenoble, France, 1994, pp. 17–24.
- [14] Lavagno, L. and E. Sentovich, *ECL: A specification environment for system-level design*, in: *Proceedings of the 36th Design Automation Conference*, New Orleans, Louisiana, 1999, pp. 511–516.
- [15] Mérillon, F., L. Réveillère, C. Consel, R. Marlet and G. Muller, *Devil: An IDL for hardware programming*, in: *Proceedings of the 4th Symposium on Operating Systems Design and Implementation (OSDI)*, San Diego, California, 2000, pp. 17–30.
- [16] Thibault, S. A., R. Marlet and C. Consel, *Domain-specific languages: from design to implementation application to video device drivers generation*, *IEEE Transactions on Software Engineering* **25** (1999), pp. 363–377.