BusMOP: a Runtime Monitoring Framework for PCI Peripherals

Rodolfo Pellizzoni, Patrick Meredith, Marco Caccamo, Grigore Roşu
Department of Computer Science, University of Illinois at Urbana-Champaign
{*rpelliz2, pmeredit, mcaccamo, grosu*}*@cs.uiuc.edu*

Abstract

COTS peripherals are heavily used in the embedded market, but their unpredictability is a threat for high-criticality real-time systems: it is hard or impossible to formally verify COTS components. Instead, we propose to monitor the runtime behavior of COTS peripherals against their assumed specifications. If violations are detected, then an appropriate recovery measure can be taken. Our monitoring solution is decentralized: a monitoring device is plugged in on a peripheral bus and monitors the peripheral behavior by examining read and write transactions on the bus. Provably correct (w.r.t. given specifications) hardware monitors are synthesized from high level specifications, and executed on FPGAs, resulting in zero runtime overhead on the system CPU. The proposed technique, called BusMOP, has been implemented as an instance of a generic runtime verification framework, called MOP, which until now has only been used for software monitoring. We experimented with our technique using a COTS data acquisition board.

I. INTRODUCTION

The real-time embedded system industry is progressively moving towards the use of Commercial-Off-The-Shelf (COTS) components in an attempt to reduce costs and time-to-market, even for highly critical systems like those deployed by the avionic industry. While specialized hardware and software solutions are sometimes available for such markets, their average performance and ease of integration is lagging behind the development of COTS components. For example, a commercial plane like the Boeing 777 uses the SAFEbus backplane [10], which, while specially designed to meet the hard real-time constraints of an avionic system, is only capable of transferring data up to 60 Mbps. On the other side, a modern COTS peripheral bus such as PCI Express 2.0 [16] can reach transfer speeds of 16 Gbyte/s, over three orders of magnitude greater than SAFEbus.

Unfortunately, when trying to use COTS for building high-integrity, real-time embedded systems, current engineering practices face significant challenges. While one can capture relevant assumptions about COTS as formal specifications, they are hard or impossible to formally verify: this is both because manufacturers are unwilling to disclose details of their implementation, for fear of losing competitive edge, and because the increase in performance is often matched by a similar increase in design complexity (out-of-order execution and branch prediction are examples of this trend in CPU design). Modern COTS peripherals running in master mode are particularly challenging. A master peripheral can directly communicate with all other elements in the system, including main memory and other peripherals, thus reducing the load on the CPU. On the other side, providing fault-containment becomes extremely hard: a misbehaving, low-criticality master peripheral could very well disrupt the entire system.

Based on the above discussion, our proposal for the safe integration of COTS peripherals in critical embedded systems is to use *runtime monitoring*: the peripheral requirement specifications are checked at runtime against its current observable behavior. If any violation is detected, then a suitable recovery action can be taken to restore the system to a safe state. The validity of the runtime monitoring approach has been proved in the field of software engineering by a large number of developed tools and techniques (see Section VII). However, applying runtime monitoring to our scenario poses some new challenges. First of all, the behavior of a COTS peripheral is controlled both by the hardware of the peripheral itself and by its software driver, hence we must check the correctness of their interactions. Second, master peripherals can directly interact with the rest of the system without requiring any action by the CPU. Based on these two considerations, our monitoring solution must be able to detect and check all communication between the peripheral and the rest of the system. Finally, runtime monitoring typically comes with an unforgivable price: runtime overhead. We can split such overhead in two components: 1) overhead due to the observation and generation of relevant events 2) overhead due to running a monitor at each event to check if any property of the specification is violated. Both types of overhead tend to be unpredictable and thus unsuitable for real-time computation.

To combat these problems, we propose a distributed monitoring technique based on the development of a *monitoring device*. The idea is to introduce an additional hardware component into the system that can check all peripheral communication and perform recovery actions, when necessary. Assuming "sniffing" data transfers does not add delay to the system, our solution prevents the first type of overhead. The second type of overhead is removed by running all monitors directly on the device, adding no runtime overhead to the CPU. Additionally, the system can run completely undisturbed as long as no recovery action is needed.

The speed of modern COTS communication architectures rules out the possibility of a software implementation for the device; instead, all logic is implemented on a reconfigurable FPGA. Finally, to show that a monitored system is safe, we need to prove that the monitoring logic monitors, indeed, the right properties. In our system, this is ensured by automatically synthesizing the monitoring logic from formal requirements specification, so that it is "correct by construction". In particular, we leverage on the Monitor Oriented Programming (MOP)[5] framework (see Section II), which is highly extensible and supports multiple formalisms, by creating a new MOP instance: BusMOP.

Illustrative Example. An example of BusMOP can be seen below. This example is a property used in the case study of Section VI and related to the behavior of Counter 2, a counter on the PCI703A board we used in our experiments.

```
logic = ERE

declarations : {
  signal cntrlCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal cntrlOld : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event countDisable : memory write address = base1 + X"220"
                      dbyte value in "---------------0"
event cntrlMod : memory write address in base1 + X"220"
                 {
                   cntrlOld <= cntrlCurrent;
                   cntrlCurrent <= value(15 downto 0);
                 }
event countEnable : memory write address = base1 + X"220"
                     dbyte value in "---------------1"

pattern: ((countEnable countDisable) + cntrlMod + countDisable)*

violation handler : {
  mem_reg <= '1';
  address_reg <= base1 + X"220";
  -- roll back to the previous cntr_cntrl2 value
  value_reg(15 downto 0) <= cntrlOld;
  cntrlCurrent <= cntrlOld;
  enable_reg <= "0011";
}
```

This property, called SafeCounterModify, requires that any modification to cntr_cntrl2, the control register for Counter 2, happens only while the counter is not in use. This modification is captured by the cntrlMod event, because cntr_cntrl2 is at address X"220". The counter can be enabled/disabled by modifying bit 0 of cntr_cntrl2 (captured by the countEnable/countDisable events; "-" is the VHDL 'don't care").

logic = ERE tells BusMOP that the property will be expressed using an extended regular expression pattern. The declarations section declares two monitor-local registers, cntrlCurrent and cntrlOld, and initializes them to 0. These registers will hold the current and previous values of the cntr_cntrl2 register. This allows us to repair the register when/if the property is *violated* by writing the old value to the register on the peripheral itself (the value_reg assignment), and forcing the current value the monitor stores to be the previous value, as can be seen in the violation handler section of the specification. The pattern itself, in the pattern section, matches any trace that consists of a cntr_cntrl2 modification, a disable of the counter, or an enable followed by a disable. The pattern is followed by *, allowing it to match repeatedly. The only way to violate this pattern, then, is to see a modification after an enable that is *not* followed by a disable first.

Using extended regular expressions is not the only possible way to express the property. In particular, we can also express the same SafeCounterModify property using past time linear temporal logic (PTLTL).

```
logic = PTLTL

declarations : {
  signal cntrlCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal cntrlOld : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event countDisable : memory write address = base1 + X"220"
                      dbyte value in "---------------0"
event cntrlMod : memory write address in base1 + X"220"
                 {
                   cntrlOld <= cntrlCurrent;
                   cntrlCurrent <= value(15 downto 0);
```

```
                  }
event countEnable : memory write address = base1 + X"220"
                    dbyte value in "---------------1"

formula: (cntrlMod) and (*)((not countDisable) S countEnable)

validation handler : {
  mem_reg <= '1';
  address_reg <= base1 + X"220";
  -- roll back to the previous cntr_cntrl2 value
  value_reg(15 downto 0) <= cntrlOld;
  cntrlCurrent <= cntrlOld;
  enable_reg <= "0011";
}
```

The main difference compared to the ERE example is that a failure to meet the specification is represented here by a *validation* of the PTLTL formula. In particular, the formula itself states if cntr_cntrl2 has been modified and the counter has not been disabled since the last time it was enabled, then we must recover.

The implementation of the events, declarations, and the actions available to handlers is explained in Section V-B. The formula/pattern implementation, and the use of handlers is explained in Section V-C. An online trial is available on our website [4], and can be used to generate monitoring code for any property in either ERE or PTLTL form.

Key contributions. We provide three main contributions. First, in Section IV we describe the design of a monitoring device for the PCI/PCI-X bus (a brief overview of PCI is presented in Section III). The monitoring device can be plugged in on a PCI bus segment, and monitor all peripherals attached to the segment. Whenever peripheral activity fails to conform to the specification, the device can perform a *corrective* action: either bring the peripheral back to a safe state if the error is recoverable, or otherwise disconnect it from the system. While certain implementation decisions are necessarily specific to our choice of PCI, we believe that the general design principles and lessons learned can be applied to most other communication architectures. Second, in Section V we provide a new instantiation of the MOP framework, called BusMOP, that is able to generate hardware modules; the generated monitoring logic is then integrated with the rest of the monitoring device design and synthesized on the FPGA. Third, in Section VI we show the feasibility of the overall approach by applying our technique together with the developed monitoring device to check a COTS data acquisition board. Our experimental results reveal that the monitoring device is able to detect, and recover from, errors caused by faults in the driver that we discovered after manually inspecting it. We conclude by discussing related work in Section VII, and providing final remarks and future work in Section VIII.

II. The MOP Framework

Monitoring-Oriented Programming (MOP) ([5] and citations there) is a formal framework for system development and analysis, in which the developer specifies desired properties using *definable* specification formalisms, along with code to execute when properties are violated or validated; it is important to note that a failure to confirm to the specification can be expressed as either the validation or violation of a property, see Section VI for examples. Monitoring code is then automatically generated from the specified properties and integrated together with the user-provided code into the original system. MOP is a highly extensible and configurable runtime verification framework; currently there are two MOP instances: JavaMOP and BusMOP (the instance described in this paper).

Property specifications consist of event definitions, which are instance dependent (e.g., pointcuts in JavaMOP and bus transactions or interrupts in BusMOP), and logical formulae or patterns, which are not. The user is allowed to extend the MOP framework with his/her own logics via *logic plugins* which encapsulate the monitor synthesis algorithms. This extensibility of MOP is supported by a layered architecture which separates monitor generation and monitor integration. By standardizing the protocols between layers, modules can be added and reused easily and independently. By providing language specific shells, logic plugins can be reused between several different MOP instances. A graphical representation of the architecture can be seen in Figure 1.

The formula or pattern designates which “traces” (observed series of events) are valid or invalid. Because extended regular expression (ERE) and past-time linear temporal logics (PTLTL) are the two plugins used in this paper, we will describe which traces are valid or invalid for ERE patterns and PTLTL formulae. For EREs, valid traces are those which are strings in the language represented by the ERE, with events treated as the letters in the alphabet of the language. Neutral traces (which trigger no handlers) are prefixes of strings in the language, while violations are invalid strings. For PTLTL formulae, valid traces are any traces for which the formula evaluates to true, invalid traces are those for which the formula evaluates to false; there are no neutral traces. For more information on regular languages and temporal logic see [14] and [7], respectively.

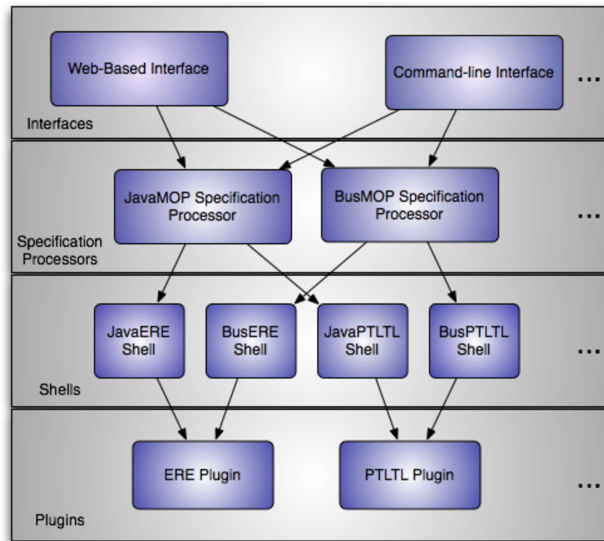

Fig. 1. MOP Architecture.

III. PCI Bus Overview

The Peripheral Component Interconnect (PCI) is the current standard family of communication architectures for motherboard/peripheral interconnection in the personal computer market; it is also widely popular in the embedded domain [16]. The standard can be divided in two parts: a *logical* specification, which details how the CPU configures and accesses peripherals through the system controller, and a *physical* specification, which details how peripherals are connected to and communicate with the motherboard. While the logical specification has remained largely unaltered since the introduction of the original PCI 1.0 standard in 1992, several different physical specifications have emerged since then.

One of the main features of the logical layer is plug-and-play (automatic configuration) functionality. On start-up, the OS executes a PCI base driver which reads information from special configuration registers implemented by each PCI-compliant peripheral and uses them to configure the system. Of peculiar importance is a set of up to 6 Base Access Registers (BARs). Each BAR represents a request by the peripheral for a block of addresses in either the I/O or memory space; the PCI base driver is responsible for accepting such requests, allocating address blocks and communicating back the chosen addresses to the peripheral, by writing in the BARs. To communicate with the peripheral, the CPU can, then, issue write and read commands, called *transactions*, to either I/O or memory space; each peripheral is required to implement *bus slave* logic, which decodes and responds to transactions targeting all address spaces allocated to the peripheral. Typically, address spaces are used to implement either registers, which control and determine the logical status of the peripheral, or data buffers. Peripherals *can* also implement *bus master* logic: they can autonomously initiate read and write transactions to either main memory or the address space of another peripheral. Master mode is typically used by high-performance peripherals to perform a DMA transfer, i.e., transfer data from the peripheral to a buffer in main memory. The peripheral's driver can then read the data directly from memory, which is much faster than issuing a read transaction on the bus. Finally, each peripheral is provided with an interrupt line that can be used to send interrupts to the CPU.

There are two main flavors of physical architecture: PCI/PCI-X is parallel, while PCI-E is serial but runs at much higher frequency (2.5Ghz against up to 133Mhz for PCI-X). We have focused on PCI/PCI-X[1], which implements a shared bus architecture. The logical PCI tree is physically divided into bus segments, and most bus wires are shared among all peripherals connected to a single segment. Each transaction seen on the bus consists of an address phase, which provides the initial address in either memory or I/O space, followed by one or multiples data phases, each of which carries up to 32 or 64 bits of data for PCI/PCI-X, respectively (individual bytes can be masked using *byte enables*). Since each bus segment is shared, arbitration is required to determine which master peripheral is allowed to transmit at any one time. Arbitration uses two active-low, point-to-point wires between the peripheral and the bus segment arbiter, REQ# and GNT#. A standard request-grant handshake is used, where the peripheral first lowers REQ# to request access to the bus, and the arbiter grants permission to start a new transaction by lowering GNT#.

Examples PCI read/write burst transactions are shown in Figures 2(a), 2(b) (for refer to [16] for detailed bus specifications). The entity that starts a transaction, either the CPU or a master peripheral, is known as the *initiator*, while the entity that receives the transaction, either a slave peripheral or main memory, is known as the *target*. All signals shown are active low. During

[1]We also plan to extend our design to PCI-E; see Section VIII.

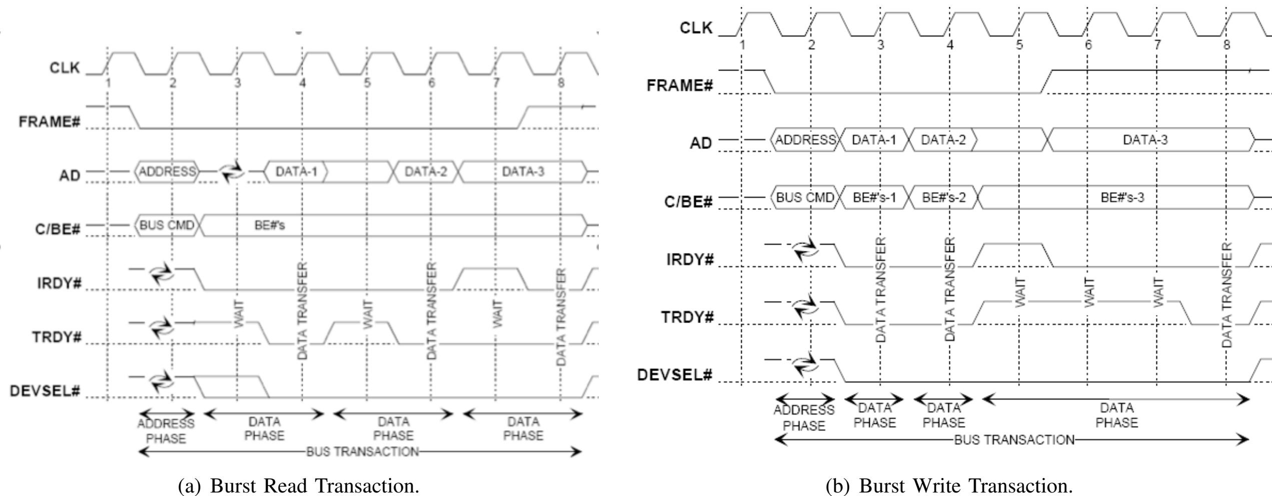

(a) Burst Read Transaction.

(b) Burst Write Transaction.

Fig. 2. PCI Transactions.

the address phase, the AD wires contain the address for the first data phase, while C/BE# determines the type of transaction: memory or I/O, read or write. During each data phase, the value is carried in AD and the address is implicitly incremented by 4/8 for a 32/64 bits bus respectively. C/BE# carries a set of byte enables for the value in AD; this permits to read/write only some of the bytes in AD in each data phase. The beginning and end of the transaction is signaled by the initiator using the FRAME# signal, while the target uses DEVSEL# to signal that it has correctly decoded the address as being part of its address space. Finally, the IRDY# and TRDY# signals can be used by the initiator and target respectively to introduce wait cycles in the transaction: a data transfer happens only when both signals are driven low.

IV. Monitoring Device

We designed a prototype monitoring device based on a Xilinx ML455 board [18] using a mixed VHDL/Verilog register transfer level (RTL) description. The board is outfitted with a Virtex-4 FPGA and is can be plugged into a standard 3.3Volts PCI/PCI-X socket. The FPGA implements both a slave and a master peripheral module, together with the monitoring modules. Events for the system are specified in terms of read/write data transfers on the bus and interrupt requests; the device continuously "sniffs" all ongoing activities on the bus, and it is therefore able to monitor communication for all other peripherals located on the same bus segment. Whenever a failure to meet the specification is detected, the device can execute a recovery action using strategies based on the detected error.

For a vast category of errors that involves incorrect interaction between the peripheral and its software driver, it is often possible to recover from the failure by forcing the peripheral into a consistent state. The monitoring device implements a master module, and can therefore initiate transactions on the bus. For example, consider a common type of error, where the driver fails to validate some input from the user and as a result writes an invalid value to a register in the peripheral. We can recover by rewriting the register with a valid value. However, if the error is caused by a fault in the peripheral hardware, interacting with registers may not be enough to bring the peripheral to a consistent and safe state.

We propose a mechanism that lets the monitoring device disconnect the faulty peripheral from the bus. We developed a simple hardware device, the *peripheral gate* [17], that is able to force the REQ# signal from the peripheral to the bus arbiter to be high; hence, the peripheral never receives the grant and it is prohibited from initiating any further transaction on the bus[2]. The peripheral gate is implemented based on a PCI extender card, i.e., a debug card that is interposed between the peripheral card and the bus and provides easy access to all signals. A clarifying picture for monitoring of a single peripheral is provided in Figure 3(a). The monitoring device can output a *stop* signal, which closes the gate when active high. Finally, sometimes the monitoring device cannot perform a suitable recovery action by itself, but there is a higher level actor, such as the OS or the system user, that can provide better recovery; examples include complex software operations such as restarting the driver or the whole PCI stack, and physically interacting with the peripheral. In this case, the best strategy is to communicate the failure to the chosen actor. The study of OS-level reliability techniques is outside the scope of this paper; instead, for our prototype design we implemented a RS-232 controller that can be used to send information to the user over a serial connection.

The reader should notice that the nature of our implementation is such that if a trace is seen, which does not conform to a specification, as a consequence of a bus transaction, that specific bus transaction can not be prevented from propagating to the rest of the system. For example, if a faulty peripheral performs a write transaction to an area in main memory which is

[2]While technically it is always possible for a faulty peripheral to disrupt the bus by altering the state of the signals, in practice the described approach is effective since access to the bus is mediated by three-state buffers enabled by GNT#.

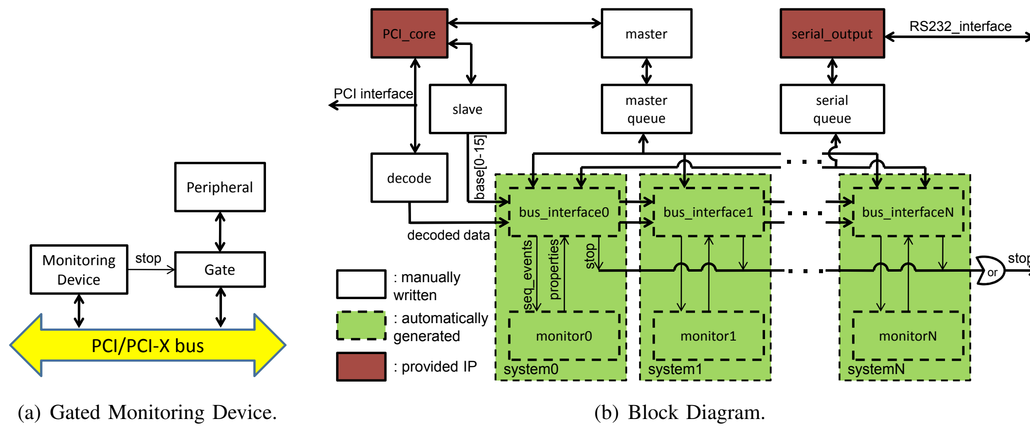

(a) Gated Monitoring Device.

(b) Block Diagram.

Fig. 3. Monitoring Device.

not supposed to modify, we can detect the error, disconnect the peripheral and report the failure to the OS/user. However, the information in the overwritten area will be lost. As part of our future work, we are working to implement an interposed monitoring device: by sitting between the bus and a peripheral, it will be able to buffer all transactions that target that specific peripheral or are initiated by it. If a property is validated/violated, it is then possible to take *preventive* measures (i.e., either discard or modify the transaction before propagating it). While this solution will provide a higher degree of reliability, there is a price to be paid in terms of increased communication delay due to buffering in the monitoring device.

A simplified block diagram for the monitoring device is depicted in Figure 3(b). While we do not show them for simplicity, all modules receive clock and reset signals from the PCI/PCI-X bus (at either 33 or 66Mhz). We distinguish three types of blocks: 1) blocks provided by Xilinx as proprietary intellectual properties (IPs); 2) manually coded RTL modules provided by BusMOP, which are independent of the peripheral specification; 3) automatically generated RTL modules, which are dependent on the specification (see Section V). PCI transaction signals are routed to two different modules: the PCI_core and the decode module.

The PCI_core module is a hard IP that implements all logic required to handle basic PCI functionalities such as plug-and-play. Bus slave and bus master logic is implemented by the slave and master modules, respectively. In particular, slave implements a set of 16 registers, base0 through base15. Since the OS configures the BAR registers at system boot, a peripheral cannot directly determine the location of address blocks used by another peripheral. Hence, the OS must also write the locations of the address blocks allocated to monitored peripherals in the base registers. The decode module is used to simplify event generation. It translates all transactions on the bus (except for those initiated by the monitoring device itself) into a series of I/O or memory reads/writes, one for each data phase, as well as the occurence of an interrupt, and forwards the translated information to the monitoring logic.

The system0, . . ., systemI, . . ., systemN blocks implement the monitoring logic for each of N user specified properties. Each systemI block consists of two automatically generated modules: bus_interfaceI contains all logic that depends on the specific choice of communication interface (PCI bus), while monitorI contains all logic that depends on the formal language used to specify the property. This separation provides good modularity and facilitates code reuse. bus_interfaceI first receives as input the decoded bus signals and generates events, which are sequentialized by the events_sequentializer submodule (see Section V-B), and then passed to monitorI using the seq_events wires. monitorI checks whenever the formula for the I-th property is validated/violated and passes the information back to bus_interfaceI, which can then execute three types of recovery: 1) disconnect a monitored peripheral from the bus using the stop signal; 2) send information to the user using the serial_output module, which implements a RS-232 transmitter; 3) start a write transaction on the bus using the master module. Finally, since it is possible for multiple systemI modules to initiate recovery at the same time, we provide queuing functionalities for serial_output and master in modules master_queue and serial_queue, respectively.

An important note is relative to the time the time elapsed from event detection to executing an handler. In the current implementation, it takes two clock cycles from event detection to informing the monitor (plus one clock cycle for each additional sequentialized event), and another clock cycle to execute the recovery. Since our current monitors can execute in one clock cycle and no property in our example specification has more than two sequentializing events, we can always recover within 5 clock cycles of the event triggering the violation/validation of the property. This time is short enough to execute a recovery action before a faulty peripheral is allowed to start a new transaction, as PCI arbitration overhead prevents a peripheral from transmitting immediately. However, for properties with more sequentializing events, or slower monitors, this could constitute a problem. A possible solution is to clock all modules of systemI at higher speed. In particular, as part of future work we plan to clock systemI at 233Mhz, or four times the speed of PCI. A phase-locked loop on the FPGA can be used to keep the new faster clock in phase with the input PCI clock, hence eliminating the need for asynchronous buffers that

would introduce additional delays.

V. Property Specifications

Properties are specified using a domain specific event syntax, and formulae or patterns written in the logic of a particular plugin. Additional monitor state can also be declared using the declarations section. The violation handler and validation handler sections allow for arbitrary code to be executed on the occurrence of a violation or validation, respectively. An example of how they are used can be seen in the example in Section I. Currently, we have support for the extended regular expression (ERE) and past-time temporal logic (PTLTL) MOP Plugins, and adding most of the others will require a minimal amount of work, as only the monitor component changes from one logical specification formalism to another. This means that properties may be specified, formally, using an ERE pattern or a PTLTL formula.

A property is implemented in two main modules, a bus_interface, which generates logical events from bus traffic and handles monitor recovery, and a monitor implementing a property specification in hardware. A more detailed description will be given below.

A. Events

A formal description of the event syntax (using Backus Naur Form (BNF) [12] extended with $[p]$ and $\{p\}$, denoting zero or one repetitions of p and zero or more repetitions of p, respectively) can be seen below:

$$
\begin{array}{rcl}
\langle \textit{Event} \rangle & ::= & \text{``}event\text{''}\,\langle \textit{ID} \rangle : \langle \textit{Expression} \rangle \\
\langle \textit{Expression} \rangle & ::= & \langle \textit{MemoryOrI/O} \rangle \langle \textit{ReadOrWrite} \rangle \text{``}address\text{''}\,\text{``}=\text{''} \\
 & & \langle \textit{ArithmeticExp} \rangle \text{``}value\text{''}\,[\text{``}not\text{''}]\,\text{``}in\text{''}\,\langle \textit{Range} \rangle \\
 & & [\langle \textit{Action} \rangle] \\
 & | & \langle \textit{MemoryOrI/O} \rangle \langle \textit{ReadOrWrite} \rangle \text{``}address\text{''}\,\text{``}in\text{''} \\
 & & \langle \textit{Range} \rangle [\text{``}\{\text{''}\,\langle \textit{Action} \rangle\,\text{``}\}\text{''}] \\
 & | & \text{``}interrupt\text{''}\,[\langle \textit{Action} \rangle] \\
\langle \textit{MemoryOrI/O} \rangle & ::= & \text{``}memory\text{''} \quad | \quad \text{``}io\text{''} \\
\langle \textit{ReadOrWrite} \rangle & ::= & \text{``}read\text{''} \quad | \quad \text{``}write\text{''} \\
\langle \textit{Action} \rangle & ::= & \text{``''}\,\langle \textit{Arbitrary VHDL code} \rangle\,\text{``''} \\
\langle \textit{Range} \rangle & ::= & \langle \textit{ArithmeticExp} \rangle [\text{``},\text{''}\,\langle \textit{ArithmeticExp} \rangle] \\
\langle \textit{ArithmeticExp} \rangle & ::= & \langle \textit{Number} \rangle \quad | \quad \langle \textit{ID} \rangle \\
 & | & \langle \textit{ArithmeticExp} \rangle\,\text{``}+\text{''}\,\langle \textit{ArithmeticExp} \rangle \\
 & | & \langle \textit{ArithmeticExp} \rangle\,\text{``}-\text{''}\,\langle \textit{ArithmeticExp} \rangle \\
 & | & \langle \textit{ArithmeticExp} \rangle\,\text{``}\&\text{''}\,\langle \textit{ArithmeticExp} \rangle \\
\langle \textit{Number} \rangle & ::= & \langle \textit{VHDL number or bitstring} \rangle \\
\langle \textit{ID} \rangle & ::= & \langle \textit{Capital or lower case letter} \rangle \{\langle \textit{LetterOrDigit} \rangle\} \\
\langle \textit{LetterOrDigit} \rangle & ::= & \langle \textit{Capital or lower case letter} \rangle \quad | \quad \langle \textit{Digit} \rangle
\end{array}
$$

There are three basic types of events in BusMOP: I/O accesses, memory accesses, and interrupts. It is important to distinguish between I/O and memory events because they require different enable functionality and different read/write signals. I/O and memory events must specify at least an address, which may be an arithmetic expression over identifiers, VHDL numbers, addition, subtraction, and concatenation, and whether the event is a read or a write. An I/O or memory event may also specify a value range, which is the value of the address read or written by the bus transaction. Ranges can consist of a single arithmetic expression, or a pair of comma separated arithmetic expression denoting the minimum and maximum values that may trigger the event (thus, ranges are inclusive). Value ranges must also specify a size, byte, dbyte (16 bits), or qbyte (32 bits), so that the correct comparison code and byte enables can be generated (values smaller than a byte require masking the proper bits). Address ranges are used in events that *do not* have specified value ranges. The reason for this is that when a value range is specified, the code generator must generate the proper byte enables based on address alignment, and alignment does not make sense for ranges. Address ranges are useful for some properties, e.g. a property that monitors accesses to a certain buffer in memory.

B. The bus_interface Module

The generated VHDL code of the bus_interface module for the ERE example in Section I is shown below.

```
library IEEE; use IEEE.STD_LOGIC_1164.ALL; use
IEEE.STD_LOGIC_ARITH.ALL; use IEEE.STD_LOGIC_UNSIGNED.ALL;
entity
bus_interface0 is
```

```
   Generic ( NUM_EVENTS : INTEGER );
    Port ( clk : in  STD_LOGIC;
           rst : in  STD_LOGIC;
           value : in  STD_LOGIC_VECTOR (31 downto 0);
           address : in  STD_LOGIC_VECTOR (31 downto 0);
           interrupt : in STD_LOGIC;
           io_read : in  STD_LOGIC;
           io_write : in  STD_LOGIC;
           mem_read : in  STD_LOGIC;
           mem_write : in  STD_LOGIC;
           enable : in  STD_LOGIC_VECTOR (3 downto 0);
           base0 : in  STD_LOGIC_VECTOR (31 downto 0);
           base1 : in  STD_LOGIC_VECTOR (31 downto 0);
           base2 : in  STD_LOGIC_VECTOR (31 downto 0);
           base3 : in  STD_LOGIC_VECTOR (31 downto 0);
           base4 : in  STD_LOGIC_VECTOR (31 downto 0);
           base5 : in  STD_LOGIC_VECTOR (31 downto 0);
           base6 : in  STD_LOGIC_VECTOR (31 downto 0);
           base7 : in  STD_LOGIC_VECTOR (31 downto 0);
           base8 : in  STD_LOGIC_VECTOR (31 downto 0);
           base9 : in  STD_LOGIC_VECTOR (31 downto 0);
           base10 : in  STD_LOGIC_VECTOR (31 downto 0);
           base11 : in  STD_LOGIC_VECTOR (31 downto 0);
           base12 : in  STD_LOGIC_VECTOR (31 downto 0);
           base13 : in  STD_LOGIC_VECTOR (31 downto 0);
           base14 : in  STD_LOGIC_VECTOR (31 downto 0);
           base15 : in  STD_LOGIC_VECTOR (31 downto 0);
           events : out  STD_LOGIC_VECTOR (2 downto 0);
           properties: in STD_LOGIC_VECTOR(1 downto 0);
           io_v : out  STD_LOGIC;
           mem_v : out  STD_LOGIC;
           stop : out  STD_LOGIC;
           address_v : out  STD_LOGIC_VECTOR (31 downto 0);
           value_v : out  STD_LOGIC_VECTOR (31 downto 0);
           enable_v : out  STD_LOGIC_VECTOR (3 downto 0);
           serial_out : out  STD_LOGIC_VECTOR (7 downto 0));
end bus_interface0;

architecture Behavioral of bus_interface0 is

   signal cntrlCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
   signal cntrlOld : STD_LOGIC_VECTOR(15 downto 0) := X"0000";

   signal event0address : STD_LOGIC_VECTOR(31 downto 0);
   signal event1address : STD_LOGIC_VECTOR(31 downto 0);
   signal event2address : STD_LOGIC_VECTOR(31 downto 0);
   signal event_val : STD_LOGIC_VECTOR(NUM_EVENTS-1 downto 0);
   signal io_reg : STD_LOGIC;
   signal mem_reg: STD_LOGIC;
   signal stop_reg: STD_LOGIC;
   signal address_reg: STD_LOGIC_VECTOR (31 downto 0);
   signal value_reg: STD_LOGIC_VECTOR (31 downto 0);
   signal enable_reg: STD_LOGIC_VECTOR (3 downto 0);
   signal serial_reg: STD_LOGIC_VECTOR (7 downto 0);

begin

   events <= event_val;
```

```
event0address <= base1 + X"220";
event_val(0) <= '1' when
  (
  mem_write = '1'
  and event0address(31 downto 2) = address(31 downto 2)
  and ((
       event0address(1 downto 0) = "00"
       and enable(1 downto 0) = "00"
       and (value(15 downto 0) = "---------------0"))
      or (event0address(1 downto 0) = "10"
       and enable(3 downto 2) = "00"
       and (value(31 downto 16) = "---------------0"))
  )
  ) else '0';


event_val(1) <= '1' when
  mem_write = '1' and (address = base1 + X"220") else '0';

event2address <= base1 + X"220";
event_val(2) <= '1' when
  (
  mem_write = '1'
  and event2address(31 downto 2) = address(31 downto 2)
  and ((
       event2address(1 downto 0) = "00"
       and enable(1 downto 0) = "00"
       and (value(15 downto 0) = "---------------1"))
      or (event2address(1 downto 0) = "10"
       and enable(3 downto 2) = "00"
       and (value(31 downto 16) = "---------------1"))
  )
  ) else '0';

io_v <= io_reg;
mem_v <= mem_reg;
stop <= stop_reg;
address_v <= address_reg;
value_v <= value_reg;
enable_v <= enable_reg;
serial_out <= serial_reg;

SIDE_EFFECTS : process( clk, rst )
   begin
      if(rst = '1') then
         io_reg <= '0';
         mem_reg <= '0';
         stop_reg <= '0';
         address_reg <= (others => '0');
         value_reg <= (others => '0');
         enable_reg <= (others => '0');
         serial_reg <= (others => '0');

      elsif(clk'EVENT and clk = '1') then
          if event_val(1) = '1' then

                cntrlOld <= cntrlCurrent;
                cntrlCurrent <= value(15 downto 0);
```

```
            end if;

            if(properties = "10") then
              --------------------------------------

  mem_reg <= '1';
  address_reg <= base1 + X"220";
  -- roll back to the previous cntr_cntrl2 value
  value_reg(15 downto 0) <= cntrlOld;
  cntrlCurrent <= cntrlOld;
  enable_reg <= "0011";

              --------------------------------------
             elsif(properties = "01") then
              --------------------------------------

              --------------------------------------
             else
               io_reg <= '0';
               mem_reg <= '0';
               stop_reg <= '0';
               address_reg <= (others => '0');
               value_reg <= (others => '0');
               enable_reg <= (others => '0');
               serial_reg <= (others => '0');
           end if;
        end if;


      end process; --SIDE_EFFECTS

end Behavioral;
```

The code for the declarations, and handler sections is copied verbatim into the VHDL module. Because of this copying, the code must be written in VHDL. The events are expanded to combinatoric statements implementing the specified logic; the PCI standard has rather complex rules governing byte ordering/enabling, which is reflected in the VHDL statements. The output of the combinatoric statements is assigned to an events wire vector, which is connected to the monitor module through an event_sequentializer submodule. Each index in the bus corresponds to the truth value of a specific event, numbered with the 0'th index as the first event, and the n'th index as the n'th event from top to bottom in the specification. This ordering is important, because it directs the event linearization performed by the event_sequentializer submodule.

The event_sequentializer is necessary because the logical formalisms expect linear, disjoint events. The event_sequentializer takes coincident events and sends them to the monitor in subsequent clock cycles, in ascending index order, using the seq_events wire vector. Therefore, if events(0) and events(3) occur in the same cycle, the monitor will see 0 followed by 3. To see why simultaneous events are possible, consider, again, Figure **??** from Section I. The cntrlMod event is asserted whenever the cntr_cntrl2 register (base1 + X"220") is written. Because both the countEnable and countDisable events require writes to the same address as the cntrlMod event, any time countEnable or countDisable are triggered, a cntrlMod is also triggered. As the property tries to enforce the policy that all modifications happen when the counter is not enabled, we must serialize events such that cntrlMod happens *after* a countDisable and *before* a countEnable. The ordering of events in Figure **??**, is consistent with this, because countDisable is listed before cntrlMod, which is listed before countDisable.

The execution of the violation handler is handled by a SIDE_EFFECTS process: it is only executed if the monitor module denotes that the property has been violated. The situation is similar for a validation handler, save that it is executed only when the formula or pattern is validated. As can be seen in the Figure 3(b), the monitor module reports the validation, violation, or neutral state of the monitored property, via the properties wire vector, to the bus_interface module. The SIDE_EFFECTS process also handles the actions associated with each event, like writing to the cntrlOld and cntrlCurrent register in the example.

Several actions are available in validation and violation handlers. Aside from manipulating any local state of the monitor (such as the write to cntrlCurrent in the example), the bus_interface module makes available several registers which can be used used to execute the recovery actions detailed in Section IV. The registers are summarized in the table below:

	Write Interface
io_reg	write request in I/O space
mem_reg	write request in memory space
address_reg	write address
value_reg	write value
enable_reg	byte enables
serial_reg	ASCII value to serial output
stop_reg	Peripheral gate control

As seen in the example, we perform a memory write to the cntr_cntrl2 register of its previous value. The address_reg is used to denote the address of cntr_cntrl2 (base1 + X"202"), while the value_reg is set to the old value of cntr_cntrl2, the mem_reg is asserted to tell the PCI bus that the write performed is a memory write, and the byte enables are set to "0011" to denote that the lower two bytes must be written.

C. The monitor Module

The monitor module is responsible for monitoring the property given serialized events. It encompasses the logic of the formula, and it is the only portion of our system dependent on the logical formalism used.

Extended Regular Expressions. Extended regular expressions (EREs) are the normal regular expressions [14], extended with negation. The same plugin used for JavaMOP's [5] EREs is used to transform the provided ERE to a minimized deterministic finite automata (DFA) defined in generic code. We convert the generic code to Verilog; the generated code for the example in Section I is shown below.

```
module monitor0(clk, rst, empty, events,properties);
   parameter NUM_EVENTS = 3;
   parameter NUM_PROPERTIES = 2;
   input clk;
   input rst;
   input empty;
   input [NUM_EVENTS-1:0] events;
   output [NUM_PROPERTIES-1:0] properties;
   reg [NUM_PROPERTIES-1:0]  properties_reg;
    reg [1:0] state;
   assign properties = properties_reg;
   always @(posedge clk or posedge rst) begin

      if (rst) begin
         properties_reg <= 0;
         state <= 0;
      end else begin
         if (events[NUM_EVENTS-1:0] != 0) begin
         // properties 0 == nothing, properties 1 == success
         // properties 2 == failure
            case (state)
               0: begin
                     state <= (events[2])?1:(events[1])?0:(events[0])?0:0;
                     properties_reg <= (events[2])?0:(events[1])?1:(events[0])?1:2;
                   end
               1: begin
                     state <= (events[0])?0:0;
                     properties_reg <= (events[0])?1:2;
                   end
               default : begin state <= 0; properties_reg <= 2; end
            endcase
         end
         else //if events[NUM_EVENTS-1:0] == 0
            properties_reg <= 0;

      end
```

```
   end
endmodule
```

The current state of the DFA is kept in the state register. On each clock cycle, the current state of the DFA and the event are consulted to see if the property is violated or validated, and what state to transition to. The output properties wire is set to 1 if the property has been validated during the current clock cycle, 2 if the property has been violated, and 0 otherwise. Note that violations of EREs are tricky, because, if used normally, a DFA, once it reaches a violation state, will report a violation every event (because there is no valid transition out of the violation state). We chose to reset the DFA to the initial state whenever a violation is encountered, to avoid this problem. ERE pattern is as follows:

$$
\begin{array}{lcl}
\langle \mathit{Pattern} \rangle & ::= & \text{“}\mathit{epsilon}\text{”} \quad | \quad \langle \mathit{Event\ Name} \rangle \\
& | & \text{“}\sim\text{”} \langle \mathit{Pattern} \rangle \quad | \quad \langle \mathit{Pattern} \rangle \text{“}*\text{”} \\
& | & \langle \mathit{Pattern} \rangle \text{“}+\text{”} \langle \mathit{Pattern} \rangle \quad | \quad \langle \mathit{Pattern} \rangle \langle \mathit{Pattern} \rangle
\end{array}
$$

“epsilon” is the empty string, “ $\sim$ ” is negation, “ $*$ ” is zero or more repetitions, “+” is logical *or*, and $\langle \mathit{Pattern} \rangle$ $\langle \mathit{Pattern} \rangle$ represents concatenation.

Past-time Linear Temporal Logic. Past-time Temporal Linear Logic (PTLTL) [7] extends normal propositional logic with *temporal* operators. We modified the PTLTL plugin used in JavaMOP to make it more suitable for implementation as a logic circuit. The original, generic code output by the plugin used a number of sequential assignments to an array of truth values. We take this sequential code and, using back substitution, change the sequential code into a series of parallel assignments. The resulting assignments are *entirely* parallel, allowing the operation of the monitor to be contained within a single clock cycle. A more in depth explanation of this transformation is omitted, but will appear in an upcoming technical report on PTLTL. The generated code for the example in Section I is reported below.

```
module monitor0(clk, rst, empty, events, properties);
   parameter NUM_EVENTS = 3;
   parameter NUM_PROPERTIES = 2;
   input clk;
   input rst;
   input empty;
   input [NUM_EVENTS-1:0] events;
   output [NUM_PROPERTIES-1:0] properties;
   reg [NUM_PROPERTIES-1:0]  properties_reg;
   reg [1:0] b;
   assign properties = properties_reg;
   always @(posedge clk or posedge rst) begin

      if (rst) begin
         properties_reg <= 0;
         b <= 0;
      end else if (events[NUM_EVENTS-1:0] != 0) begin
         // properties 0 == nothing, properties 1 == success
         // properties 2 == failure
         if (events[1] && b[1])
             properties_reg <= 1;
         else
             properties_reg <= 2;

         //parallel assignments
         b[0] <= events[2] || ~events[0] && b[0] ;
      b[1] <= events[2] || ~events[0] && b[0];
      end else // if events[NUM_EVENTS-1:0] == 0
          properties_reg <= 0;
   end
endmodule
```

The syntax for PTLTL formulas is as follows:

$$
\begin{array}{lcl}
\langle \textit{Formula} \rangle & ::= & \text{“}\textit{true}\text{”} \quad | \quad \text{“}\textit{false}\text{”} \quad | \quad \langle \textit{Event Name} \rangle \\
& | & \text{“}\textit{not}\text{”} \langle \textit{Formula} \rangle \quad | \quad \langle \textit{Formula} \rangle \text{“}\textit{and}\text{”} \langle \textit{Formula} \rangle \\
& | & \langle \textit{Formula} \rangle \text{“}\textit{or}\text{”} \langle \textit{Formula} \rangle \\
& | & \langle \textit{Formula} \rangle \text{“}\textit{implies}\text{”} \langle \textit{Formula} \rangle \\
& | & \text{“}[*]\text{”} \langle \textit{Formula} \rangle \quad | \quad \text{“}\langle * \rangle\text{”} \langle \textit{Formula} \rangle \\
& | & \text{“}(*)\text{”} \langle \textit{Formula} \rangle \quad | \quad \langle \textit{Formula} \rangle \text{“}\textit{S}\text{”} \langle \textit{Formula} \rangle
\end{array}
$$

“not”, *“and”*, *“or”*, and *“implies”* are the ordinary logic operators. “[∗]”, “⟨∗⟩”, “(∗)”, and *“S”* are temporal operators denoting always in the past, eventually in the past, previously, and since, respectively.

A design decision relating to both logics we have implemented, and all future logics, is that properties cannot be violated or validated before an event arrives. Without this assumption, the example ERE property would be valid at start up. This creates a problem: to correctly trigger recovery actions in the bus_interface module, we require that the properties wire be set to 1/2 (for a validation/violation respectively) for only one clock cycle. The solution we adopted is simply to set properties to zero when no event is detected. An additional problem is that without the assumption, a single event in ERE could cause a violation followed immediately by a validation (since we reset the monitor on violation) in the same clock cycle. This could in turn trigger both a validation and violation handler at the same time, which is something we can not support. JavaMOP has the same functionality, but in JavaMOP it is due to the fact that a monitor does not *exist* before the first event, whereas in BusMOP, the monitor exists as soon as the FPGA is configured.

VI. Case Study: The PCI703A ADC/DAC Board

In this section, we show how our runtime monitoring technique can be applied to a concrete case by providing specification and runtime experiments for a specific COTS peripheral, the PCI703A board [8]. PCI703A is a high performance Analog-to-Digital/Digital-to-Analog Conversion (ADC/DAC) peripheral for the PCI bus. In particular, it can perform high-speed, 14-bits precision ADC at a rate of up to 450,000 conversions/s, and transfer data to main memory in bus master mode. At the same time, the peripheral is simple enough that we were able to carefully check all provided hardware manuals and to manually inspect its Linux driver; specifying formal properties for a peripheral clearly requires a deep understanding of its inner working. In our proposed model, the peripheral's manufacturer is responsible for writing the runtime specification. In this sense, the formal specification can be thought of as a correctness certification provided by the manufacturer, as long as the user employs a monitoring device and recovery actions can be proved to restore the system to a safe state.

To better mimic what we think would be a typical process for a COTS manufacturer, we produced a requirement specification for the PCI703A in two steps. First, we prepared a detailed description of the communication behavior of the peripheral in plain English. Then, we converted this informal description into a formal set of events and formulae for BusMOP. Inspection of the driver revealed two software faults, both of which can cause errors that are detected and recovered by the monitoring device. While in this case we could have prevented errors by simply removing the faults, we argue that drivers for more complex peripherals can be thousands of lines long and neither code inspection nor testing is sufficient to remove all bugs. We further injected additional faults in the driver to test all written formal properties. It would have been nice to also show recovery for hardware faults, but we did not find any in the tested peripheral and injecting faults in the hardware is difficult.

The rest of the section is organized as follows. In Section VI-A we provide an overview of PCI703A and we detailed English specification. In Sections VI-B, VI-C we detail the two driver faults together with a set of formal properties used to detect the generated errors. Finally, in Section VI-D we provide additional formal properties based on the English specification; we do not claim that the list of formal properties is exhaustive, but rather, we focus on providing useful examples of BusMOP usage.

A. Informal Specification

A block diagram for the PCI703A is shown in Figure VI-A. The bus slave logic implements two memory address blocks in BAR0 and BAR1, used for conversion data and control registers, respectively; the corresponding base addresses are written in base0 and base1 in the monitoring device. The ADC Control and DAC Control blocks control the ADC/DAC operations and write/read data into internal FIFOs. The DMA Control block can be programmed to move data between each FIFO and main memory using bus master functionality. Finally, the Counter Timers block implements four counters with different functionalities. In what follows, we provide a detailed description of the communication requirements for the Counter Timers, ADC Control and DMA Control blocks, ignoring the DAC functionality for simplicity. Unless otherwise noted, in the description all registers are 16-bits, memory mapped, relative to either BAR0 or BAR1.

1) Counter Timers subsystem: The Counter Timers module implements four counters, Counter 0 through Counter 3. Counter 0 and 1 are user programmable and can be used either for debugging purposes or to trigger a DA conversion. Counter 3 is also user programmable and produces an external output. Finally, Counter 2 is not meant to be user programmable; it is to be used exclusively to generate the clock for AD conversions. Each counter has two registers: Counter 0 and 1 are user programmable

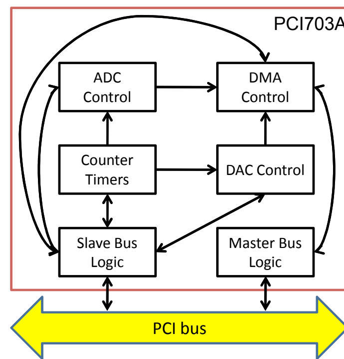

Fig. 4. PCI703A Diagram.

and can be used either for debugging purposes or to trigger a DA conversion. Counter 3 is also user programmable and produces an external output. Finally, Counter 2 is not meant to be user programmable; it is to be used exclusively to generate the clock for AD conversions. Each counter has two registers, cntr_cntrl(0-3) (at hexadecimal address 200-210-220-230 relative to BAR0) and cntr_divr(0-3) (at hexadecimal address 208-218-228-238 relative to BAR0). The initial value is loaded into the counter from cntr_divr(0-3) and the counter counts down, generating an output when it reaches 0. Then the counter loads its value again from cntr_divr(0-3) and so on until the counter is disabled.

cntr_cntrl(0-3): control register for Counter (0-3). A list of relevant bits follows:

- Bit0, CNTR_ENABLE: enables (1) / disables (0) the counter (counter counts down only when enabled, each time the counter transitions from enable to disable it loads the initial value from cntr_divr(0-3)).
- Bit2-1, CNTR_SRC: determines the source clock for the counter. Values 00 or 01 are admissible. Values 10 and 11 are prohibited for Counter 2, admissible for Counters 0,1,3.
- Bit3, output mode: 0 generates a single pulse when the counter reaches 0, 1 toggles each time the counter reaches 0.
- Bit4, interrupt mode: If set to 1 the counter generates an interrupt when it counts to 0.

A counter must be disabled before any other change to cntr_cntrl(0-3) or to cntr_divr(0-3). Setting bit 4 for Counter 2 results in additional interrupts that the driver receives, but discards in the interrupt handler (not a correctness problem but it impacts performance since calling an interrupt handler is costly).

2) DMA Control subsytem: The DMA subsystem is used to move data from the AD Fifo (where ADC data is stored) to a buffer in main memory. It consists of two 32 bits registers: dma_cfg (BAR0 + C4) and dma_addr (BAR0 + C8).

dma_addr: Contains the physical address of the buffer in main memory.

dma_cfg: control register for the DMA controller. A list of relevant bits follows:

- Bit 0, DMA_DONE_BIT: set to 1 by the hardware after a DMA operation ends. Must be reset before a new DMA operation starts.
- Bit 5, DMA_ENABLE_BIT: enables (1) / disables (0) the DMA controller. Must be enabled during a DMA operation.
- Bit 3, DMA_DIR_BIT: must be 0 (indicates a data transfer from HW to memory).
- Bit 4, DMA_REQUEST_BIT: writing 1 to this bit starts a DMA operation.
- Bit 8, INT_STATUS_BIT: set to 1 by the hardware after the DMA controller generates an interrupt. Must be reset by writing 0 to it before a new DMA operation starts.
- Bit 9, INT_ENABLE_BIT: enables/disables the generation of an interrupt when a DMA operation ends.
- Bit 31-16, DMA_BUF_SIZE: size of the transfer in bytes.

The DMA subsystem will perform memory writes into main memory after a DMA operation has been started (DMA_ENABLE_BIT=1, DMA_DIR_BIT=0, write 1 to DMA_REQUEST_BIT). It should only write in the addresses between DMA_ADDR and DMA_ADDR+DMA_BUF_SIZE-1; it can write less than DMA_BUF_SIZE bytes but no more than that. After the operation ends, the driver must read DMA_DONE_BIT set to 1 (if INT_ENABLE_BIT has been set, then the driver must have received an interrupt and also read INT_STATUS_BIT set to 1 - note that INT_STATUS_BIT and DMA_DONE_BIT can be read with a single read transaction). The DMA configuration must not be changed while a DMA operation is in progress.

3) ADC Control subsystem: The ADC subsystem is used to perform analog-to-digital conversion. Since the subsystem is quite complex, we first provide a quick overview of its operational behavior:

- The driver must first provide a channel list. The channel list is a circular list of channels (analog inputs). The ADC module goes through the list one channel at the time and performs an AD conversion on that channel each time it receives a triggering event. After it reaches the end of the list it goes back to the first channel (but several events are triggered at this point, see below).

- The ADC system must be correctly configured. This primarily consists in setting a clock source and a trigger condition. Each time it receives a triggering event, the ADC module performs a conversion on the next clock cycle. Triggering events are as follows:
  - Internal: always triggered (this means a conversion every clock cycle).
  - External input: an external digital line provides the trigger event.
  - Reference: the trigger event is generated if an input analog signal is greater than a reference value.
- The possible clock sources are:
  - Counter 2: the clock is generated by the output of counter 2.
  - Software: the driver must generate the clock by writing to a specific bit in a register.
  - External: an external digital line provides the clock.
- After a channel list has been provided and the ADC has been configured, the ADC process can start (note: if using clock source software, the driver must perform a register write for each conversion) and data is put into the AD fifo.
- The driver must be informed when there is available ADC data to be read. This can be done by polling a register, or by using interrupts.
- The driver must then acquire the data. This can be done by reading from the AD fifo, or by using DMA (involves starting a DMA operation and then reading data from main memory).

The ADC subsystem uses four main registers: global_status (BAR1 + 500), ad_fifo (BAR0 + 00), adc_cntrl (BAR1 + 300), and adc_chlist (BAR0 + 04).

global_status: global status information for the peripheral. A list of relevant bits follows:

- Bit 0, INT_ENABLE: enables/disable interrupts for the whole peripheral (individual interrupts for each subsystem are enabled if this bit is set to 1 and their individual interrupt enable bit is set).
- Bit 1, FIFO_EMPTY_STAT: read only. 1 if the AD fifo is empty, 0 if there is data (used for polling).
- Bit 2, STAT_CHLIST_DONE: set to 1 by the hardware if the ADC has reached the end of the channel list. Cleared by the driver writing 0 to it.
- Bit 4, IRQ_CHLIST_DONE: set to 1 by the hardware if the ADC has reached the end of the channel list and an interrupt has been generated (see below). Cleared by the driver writing 0 to STAT_CHLIST_DONE.

ad_fifo: reading this register pops the top value off the AD fifo (result is undefined if FIFO_EMPTY_STAT is 0).

adc_cntrl: ADC control register. A list of relevant bits follows:

- Bit0, ADC_ENABLE: enable/disable the ADC Control module.
- Bit2-1, ADC_CLK_SRC: determines the clock source. 00: software. 01: counter 2. 10: external. 11: prohibited.
- Bit4-3, ADC_TRIGGER_SRC: determines the trigger source. 00: internal. 01: reference. 10: external. 11: prohibited.
- Bit10, ADC_SW: this bit generates the ADC clock in clock source software mode.
- Bit11: if set to 1, the ADC generates additional interrupts that are unprocessed by the driver. Must be put to 0.
- Bit 12, ADC_INT_ENABLE: enables/disables interrupt generation. If enabled an interrupt is generated every time the ADC reaches the end of the channel list. This interrupt can be processed by the driver.

adc_chlist: channel list register. Writing to this register adds a new channel at the end of the channel list. Every combination of the first 14 bits is legal; it specifies a different analog input with a different gain, i.e. voltage range.

The board has a maximum conversion speed of 450,000 samples each second. If the ADC uses Counter 2 as its clock source, and for Counter 2, CNTR_SRC is set to 00 (20Mhz source) and cntr_divr2 is set to 44 or less, the board will not behave correctly.

Configuration modes. When ADC_ENABLE =1, ADC_CLK_SRC is software and ADC_TRIGGER_SRC is internal, a data acquisition is started by writing 1 to ADC_SW if the previous written value is 0. The two following configurations are prohibited when ADC_ENABLE=1: ADC_CLK_SRC is software, ADC_TRIGGERT_SRC is external; and ADC_CLK_SRC is software, ADC_TRIGGERT_SRC is reference.

If ADC_ENABLE = 1, and any other configuration of ADC_CLK_SRC and ADC_TRIGGERT_SRC is set, then the ADC is said to be active (it starts performing multiple AD conversions). The ADC can be stopped by either setting ADC_ENABLE = 0 or the ADC_CLK_SRC to software. While the ADC is active, there should be no changes to either the ADC_CNTRL register (unless the register is used to stop the ADC) or to the channel list. Furthermore, if ADC_CLK_SRC is set to Counter 2, while the ADC is active no change to Counter 2 configuration is allowed, and Counter 2 must be active and its output mode must be zero.

Channel list configuration. The channel list is filled by writing channel information to register ADC_CHLIST. The number of channels in the list is equal to the number of 16 bit words written to ADC_CHLIST, unless ADC_TRIGGER_SRC is set to reference, in which case it is equal to the number of written words minus one (the first word is used to configure the reference input). There must be at least one channel in the channel list before any data acquisition is performed. Initially the channel list is empty. The channel list is cleared whenever a 0 is written to ADC_ENABLE if the previous written value is 1 (this also clears the AD_FIFO). Channels should not be added to the list after any AD conversion is performed, unless the channel list is cleared before. A maximum of 2048 channels can be in the list.

Reading data: polling mode. This is the simplest way to get data from the ADC module. It can be used in any configuration. In polling mode, the driver reads the global_status register one or more times until FIFO_EMPTY_STAT is zero; the driver then reads a single value from the AD_FIFO register. If ADC_CLK_SRC is set to software, then the number of reads from AD_FIFO must be less than or equal to the number of AD conversions performed (determined through the ADC_SW register, and taking into account that the AD_FIFO can be cleared as explained above). In any other mode, no such easy check is possible since the number of conversions depend either on an asynchronous clock or external events.

Reading data: interrupt mode with fifo read. This mode is possible only if INT_ENABLE and ADC_INT_ENABLE are set, and ADC_CLK_SRC is not set to software. The following list of events, all of which can be intercepted by the monitor, must be repeated while the ADC is active.

- The peripheral generates an interrupt (note - interrupts could also be generated by other subsystems).
- After entering the interrupt handler, the driver reads IRQ_CHLIST_DONE set to 1. This ensures that the interrupt has indeed been generated by the ADC module.
- The driver reads N data words from the AD fifo, where N is the number of channels in the channel list (note: it is not necessary to check FIFO_EMPTY_STAT like in polling mode, since the interrupt is generated only after the entire list has been processed once).
- The driver clears the interrupt by writing 0 to STAT_CHLIST_DONE.

Note that while this mode is legal according to peripheral specification, the Linux driver does not use it.

Reading data: interrupt mode with DMA transfer. This is the highest performance mode and the one the driver uses if ADC_CLK_SRC is not set to software. Again, this mode assumes that INT_ENABLE and ADC_INT_ENABLE are set. The following list of events, all of which can be intercepted by the monitor, must be repeated while the ADC is active.

- The peripheral generates an interrupt.
- After entering the interrupt handler, the driver reads IRQ_CHLIST_DONE set to 1.
- The driver starts a DMA operation to transfer data from the AD fifo into main memory. The driver must have configured the DMA subsystem before starting the transfer. Also, the driver can only work correctly by setting the DMA INT_ENABLE_BIT to 1 since it needs to receive interrupts from it, and the DMA_BUF_SIZE must be smaller or equal to $2N$, where N is the number of channels in the channel list.
- After the DMA operation starts the driver should clear the ADC interrupt by writing 0 to STAT_CHLIST_DONE.
- From now on until the ADC is disabled, the driver receives interrupts from both the ADC and the DMA subsystem. The two interrupts can be distinguished based on the value of the IRQ_CHLIST_DONE and the DMA INT_STATUS_BIT (these bits must always be cleared by writing 0 to either STAT_CHLIST_DONE or INT_STATUS_BIT before exiting the interrupt handler).
- In both cases, the interrupt handler can start a new DMA transfer if there is enough data to be transferred. This can be checked with the following equation: if $2N(numberoftimesADCinterrupthasbeenreceived) \geq DMA_BUF_SIZE(1 + numberoftimesDMAinterrupthasbeenreceived)$ and the DMA is not currently performing another transfer, then a new transfer is initiated.

B. Counter Configuration Fault

The first driver fault is relative to counter configuration. The C user library provided with the driver exports an ADConfig function used to configure ADC Control and the associated Counter 2. The library also provides a CTConfig function to be used to configure the user counters; unfortunately, under Linux the function can also be used to change the configuration of Counter 2. This is a problem, as any user in the system could erroneously or maliciously change Counter 2 while an ADC is in progress. We now show a complete set of properties that are able to capture any invalid configuration arising from exploitation of the vulnerability.

SafeCounterModify. The example SafeCounterModify property detailed in Section I prevents the cntr_cntrl2 register from being modified while the counter is in use.

SafeDivrModify. This property is the same as SafeCounterModify , save that we are ensuring that cntr_divr2 is not modified, rather than cntr_cntrl2. These could be collapsed into one specification, but it would make recovery more complicated, because we only want to roll back the register that was actually modified (cntr_cntrl2 or cntr_divr2). The formula itself states that if cntr_divr2 has been modified and the counter has not been disabled since the last time it was enabled, than we must recover.

```
logic = PTLTL

declarations : {
  signal divrCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal divrOld : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}
```

```
event countDisable : memory write address = base1 + X"220"
                     dbyte value in "---------------0"
event divrMod : memory write address in base1 + X"228"
                {
                  divrOld <= divrCurrent;
                  divrCurrent <= value(15 downto 0);
                }
event countEnable : memory write address = base1 + X"220"
                    dbyte value in "---------------1"

formula: (divrMod) and (*)((not countDisable) S countEnable)

validation handler : {
  mem_reg <= '1';
  address_reg <= base1 + X"228";
  -- roll back to the previous cntr_divr2 value
  value_reg(15 downto 0) <= divrOld;
  divrCurrent <= divrOld;
  enable_reg <= "0011";
}
```

ConfigurationFix. The ConfigurationFix property ensures that bit 4 of cntr_cntrl2 is not set, and that bit 2-1, CNTR_SRC, is set to a valid value. We simply check setBit4, an event which corresponds to setting the 4th bit, and setBit2, and event which corresponds to an invalid CNTR_SRC configuration (invalid values of CNTR_SRC are 11 and 10). We perform recovery when either of the two events is detected by overwriting cntr_cntrl2 with the last valid value, similarly to SafeCounterModify in Section I.

```
logic = ERE

declarations : {
  signal cntrlCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal cntrlOld : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event cntrlMod : memory write address in base1 + X"220"
                 {
                   cntrlOld <= cntrlCurrent;
                   cntrlCurrent <= value(15 downto 0);
                 }
event setBit2 : memory write
                address = base1 + X"220"
                dbyte value in "-------------1--"

event setBit4 : memory write
                address = base1 + X"220"
                dbyte value in "-----------1----"

pattern: (setBit2 + setBit4)*

validation handler : {
  mem_reg <= '1';
  address_reg <= base1 + X"220";
  -- roll back to the previous cntr_cntrl2 value
  value_reg(15 downto 0) <= cntrlOld;
  cntrlCurrent <= cntrlOld;
  enable_reg <= "0011";
}
```

SafeConversionSpeed. The SafeConversionSpeed ensures that the if the ADC is using Counter 2, the conversion speed is

within the board limit. For this property we chose to show how event side effects can be used in handlers as part of checking that a property has been validated/violated. When the clkSrcSet or srcSet events are triggered, meaning that the cntr_cntrl2 or adc_cntrl registers have been modified, respectively, we store the value written to the register in monitor local registers (e.g., src <= value(15 downto 0)). The pattern specifies that the cntr_divr2 be set to a bad value (less than 45), followed by any number of updates to cntr_cntrl2 or adc_cntrl, followed by the enabling of the counter. If cntr_divr2 is set to a value larger than 44, the pattern will be violated, and the monitor will be reset. This means that the validation handler will be executed only when then value of cntr_divr2 is too low for safe conversion, but regardless of whether or not the board is actually using Counter 2. The handler then checks that it is, in fact using Counter 2, and that Counter 2 is using the 20Mhz source, before performing the recovery: setting cntr_divr2 to a valid value (45).

```
logic = ERE

declarations : {
  signal clkSrc : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal src : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event divrBad: memory write address = base1 + X"228"
               dbyte value in 0,44
event divrGood: memory write address = base1 + X"228"
                dbyte value in 45,65535
event clkSrcSet : memory write address in base1 + X"300"
                  { clkSrc <= value(15 downto 0); }
event srcSet :  memory write address in base1 + X"220"
                { src <= value(15 downto 0); }
event countEnable : memory write address = base1 + X"220"
                    dbyte value in "---------------1"

pattern : (divrBad (clkSrcSet + srcSet)* countEnable)*

validation handler : {
  if (clkSrc(2 downto 1) = "01") and (src(2 downto 1) = "00") then
    mem_reg <= '1';
    address_reg <= base1 + X"228";
    --set cntr_divr2 to 45
    value_reg(15 downto 0) <= X"2D";
    enable_reg <= "0011";
  end if;
}
```

ValidWhileConverting. The ValidWhileConverting property checks that if the ADC is using Counter 2, then Counter 2 must be active and its output mode must be 0. This could have been written in a similar manner to SafeConversionSpeed, i.e., using event side effects to store current register values and checking them in the handler. We decided to use a fully formal specification, that defines events for setting the registers to good or bad values. The formula itself specifies that, if the ADC is enabled, and clkSrc2 is good, meaning that Counter 2 is being used to time the ADC, then Counter 2 must be enabled with output mode 0. The part of the formula before the implies keyword, states that the ADC is enabled and the ADC clock source is Counter 2, the second half of the formula is the requirement that Counter 2 be in a valid configuration. The formula is true when correct behavior is exhibited, so we use a violation handler for the recovery action, which again is simply to set cntr_cntrl2 to the last valid value.

```
logic = PTLTL

declarations : {
  signal cntrlCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal cntrlOld : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event countValid : memory write address = base1 + X"220"
                   dbyte value in "------------0--1"
```

```
                    {
                      cntrlOld <= cntrlCurrent;
                      cntrlCurrent <= value(15 downto 0);
                    }
event countInvalid : memory write address = base1 + X"220"
                     dbyte value not in "------------0--1"
                     {
                       cntrlOld <= cntrlCurrent;
                       cntrlCurrent <= value(15 downto 0);
                     }
event clkSrc2Good : memory write address = base1 + X"300"
                    dbyte value in "-------------01-"
event clkSrc2Bad : memory write address = base1 + X"300"
                   dbyte value not in "-------------01-"
event adcEnable : memory write address = base1 + X"300"
                  dbyte value in "---------------1"
event adcDisable : memory write address = base1 + X"300"
                   dbyte value in "---------------0"

formula : ( ((not adcDisable) S adcEnable) and
            ((not clkSrc2Bad) S clkSrc2Good) )
          implies
          ((not countInvalid) S countValid)

violation handler : {
  mem_reg <= '1';
  address_reg <= base1 + X"220";
  -- roll back to the previous cntr_cntrl2 value
  value_reg(15 downto 0) <= cntrlOld;
  cntrlCurrent <= cntrlOld;
  enable_reg <= "0011";
}
```

As a final consideration, note that the handlers of SafeCounterModify, ConfigurationFix and ValidWhileConverting can be invoked simultaneously if an incorrect value is written to cntr_cntrl2, which results in the execution of multiple bus writes. However, this causes no problem since all handlers overwrite cntr_cntrl2 with the same valid value.

C. Channel List Fault

The second fault is relative to the way the driver handles the channel list if ADC_TRIGGER_SRC is set to reference. According to the board specification, the driver would need to write an additional value to adc_chlist (the reference value), but the driver fails to do so. The error can be detected in two cases: the driver writes a single channel to the adc_chlist register, which means that the channel list is empty, before starting a conversion; or interrupt mode with fifo read is used, and the driver tries to read more than N data words from the AD fifo, where N is the number of channels in the channel list. We now show two properties that capture the error in the two described situations.

NoZeroChannels. The NoZeroChannel property checks that the there is at least one channel in the channel list when the ADC is activated with reference trigger mode. We use an ERE pattern, and recover on validation. Recall that the channel list channel list is cleared when ADC_ENABLE is reset to zero from one. A valid trace is therefore one in which only one write to adc_chlist before the ADC is enabled again (note that the ADC can be disabled multiple times between the two activations without clearing the channel list). We check that the trigger mode is set to reference in the handler. Since there is no way to correctly start the ADC process without knowing the intended reference value, the recovery consists in disabling the ADC.

```
logic = ERE

declarations : {
  signal adcCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event adcEnabled: memory write address = base1 + X"300"
```

```
                   dbyte value in "---------------1"
                   { adcCurrent <= value(15 downto 0); }
event adcDisabled: memory write address = base1 + X"300"
                   dbyte value in "---------------1"
                   { adcCurrent <= value(15 downto 0); }
event addCh : memory write address in base0 + X"04"

pattern : (adcEnabled adcDisabled adcDisabled* addCh adcDisabled*
adcEnabled)*

validation handler : {
  if (adcCurrent(4 downto 3) = "01") then
    mem_reg <= '1';
    address_reg <= base1 + X"300";
    --disable adc
    value_reg(15 downto 0) <= adcCurrent(15 downto 1) & '0';
    enable_reg <= "0011";
  end if;
}
```

OnlyNReads. This property checks that when using the interrupt mode with fifo read, the driver does not try to read a number of data words from the AD fifo greater than the number of channels in the channel list. We express the property as a validation of a PTLTL formula. Since the property depends on the comparison between the number of reads in the AD fifo and the number of writes in adc_chlist, we need to keep two registers. The numReads register stores the number of fifo reads; it is reset every time we detect and ADC interrupt (bit 4 of global_status set, event irqAdc) and incremented by one each time we perform a read in ad_fifo (event fifoRead). The numListWrites register stores the number of writes in adc_chlist; it is reset every time ADC_ENABLE is reset to zero from one (event resetCh) and incremented in event addCh. The PTLTL formula states that trigger mode must be set to reference, a fifo read happens at the current time and we have not seen a read of FIFO_EMPTY_STAT equal to zero since the previous fifo read; the latter condition is required to prevent the formula from being validated when using polling mode. Finally, in the validation handler we recover by disabling the ADC if numReads is at least equal to numListWrites (meaning that the number of fifo reads at least one greater than the number of channels).

```
logic = PTLTL

declarations : {
  signal adcCurrent : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
  signal numListWrites : UNSIGNED := 0;
  signal numReads : UNSIGNED := 0;
}

event irqAdc: memory read address = base1 + X"500"
              dbyte value in "-----------1----"
              { numReads <= 0; }
event fifoRead: memory read address in base0
                { numReads <= numReads + 1; }
event trigRef: memory write address = base1 + X"300"
               dbyte value in "-----------01---"
               { adcCurrent <= value(15 downto 0); }
event trigNotRef: memory write address = base1 + X"300"
                  dbyte value not in "-----------01---"
                  { adcCurrent <= value(15 downto 0); }
event fifoNotEmpty: memory read address = base1 + X"500"
                    dbyte value in "--------------0-"
event addCh: memory write address in base0 + X"04"
             { numListWrites <= numListWrites + 1; }
event resetCh: memory write address = base1 + X"300"
               dbyte value in "---------------0"
              {
                if(adcCurrent(0) = '1') then
```

```
                  numListWrites <= 0;
                end if;
              }

formula : ((not trigNotRef) S trigRef)
          and fifoRead
          and (*)((not fifoNotEmpty) S fifoRead)

violation handler : {
  if (numReads >= numListWrites) then
    mem_reg <= '1';
    address_reg <= base1 + X"300";
    --disable adc
    value_reg(15 downto 0) <= adcCurrent(15 downto 1) & '0';
    enable_reg <= "0011";
  end if;
}
```

D. Additional Properties

We conclude our experimental section by showing two more properties that, while unrelated to the driver faults, are particularly instructive because they express requirements that are common for master peripherals and their drivers: correctness of DMA operations; and correctness of interrupt management. The two properties also show additional BusMOP functionalities that were not covered in the previous examples: address ranges and interrupt events.

SafeMemoryWrite. The SafeMemoryWrite property is used to check that all writes in main memory performed by the DMA engine are safe. We assume that the lowest and highest physical addresses of the main memory are stored in the base2 and base3 registers respectively. The property is then expressed as the validation of a PTLTL formula composed of the disjunction of two parts. The first part captures writes in main memory outside the buffer allocated by the driver to the peripheral, which are always invalid. The second part captures any write in the buffer while a DMA operation is not in progress; referring to the specification in Section VI-A, a DMA operation is initiating after writing a one to the DMA_REQUEST_BIT with DMA_ENABLE_BIT set to one and DIR_BIT set to zero, and we know it is finished when the driver reads DMA_DONE_BIT set to one. The base buffer address and its length are contained in the bufferAddr and bufferLen monitor registers respectively, which are set as side effects of writes to the dma_addr register and the DMA_BUF_SIZE bit field. Finally, since a validation of the property indicates a dangerous hardware faults, our recovery measure is to disconnect the peripheral from the bus.

```
logic = PTLTL

declarations : {
  signal bufferAddr : STD_LOGIC_VECTOR(31 downto 0) := X"00000000";
  signal bufferLen : STD_LOGIC_VECTOR(15 downto 0) := X"0000";
}

event addrWrite: memory write address in base0 + X"C4"
                        { bufferAddr <= value(31 downto 0); }
event lenWrite: memory write address in base0 + X"C8"
                      { bufferLen <= value(31 downto 16); }
event dmaDone: memory read address = base0 + X"C8"
                        dbyte value in "---------------1"
event dmaStart: memory write address = base1 + X"C8"
                        dbyte value in "------1----10---"
event writeBefore: memory write address in base2, bufferAddr - 1
event writeIn: memory write address in bufferAddr, bufferAddr + bufferLen - 1
event writeAfter: memory write address in bufferAddr + bufferLen, base3

formula : (writeBefore or writeAfter) or
              ( writeIn and not ((not dmaDone) S dmaStart) )

validation handler : {
  stop_reg <= '1';
```

```
}
```

AckInterrupt. The AckInterrupt property checks that whenever the driver receives an interrupt, the interrupt is acknowledged before the next interrupt is generated by the peripheral. According to the specification in Section VI-A, the peripheral can generate an interrupt in two cases: the ADC has reached the end of the channel list; or a DMA operation has finished. In the first case, the driver must acknowledge the interrupt by writing a zero to STAT_CHLIST_DONE in the bufferLen register; in the second case, it must write a zero to INT_STATUS_BIT in dma_cfg. The property, expressed as the validation of a PTLTL formula, simply states that we receive an interrupt, and that the driver did not perform any of the two acknowledging actions since the previous detected interrupt. Once again, since this is a dangerous fault either in the peripheral or the driver, we block the peripheral on validation.

```
logic = PTLTL

event perInt: interrupt
event dmaAck: memory write address = base0 + X"C8"
              dbyte value in "-------0--------"
event chlistAck: memory write address = base1 + X"500"
                 dbyte value in "-------------0--"

formula : perInt and
          (*)((not dmaAck and not chlistAck) S perInt)

validation handler : {
  stop_reg <= '1';
}
```

VII. Related Work

There are two main run-time verification approaches: 1) offline, where a log, or trace is kept, which can then be used for purposes of debugging; and 2) online, where a property is checked while the program is running. As BusMOP is an online technique, we will only describe online approaches to runtime verification.

MaC [11], PathExplorer (PaX [9], and Eagle [3] use specific verification languages which cannot be changed, while BusMOP, as an extension of MO [5], will eventually support all the logics supported in JavaMOP. Temporal Rove [6] is a commercial runtime verification tool which uses future time metric temporal logic. It provides inline specification of monitors, where the monitors are written straight in the source file. Inline specification does not make sense for BusMOP, as there is no program being monitored per se. Program Query Language (PQL [15], is an approach somewhat similar to MOP, although it also only allows one specification language. PQL can support the full generality of context free languages. Tracematches [2] is very similar to JavaMOP. The biggest difference is that its choice of regular expressions for logical formalism is hardwired. It is an extension of the AB [1] AspectJ compiler. All of the above approaches are designed to monitor specific programs, and are implemented in software. This has the effect of both adding runtime overhead, and performing a function different from that of BusMOP, which monitors COTS peripherals.

The PSL to Verilog compiler, P2 [13], is the sole attempt to perform formal runtime verification in hardware, of which we are aware. P2V is similar to BusMOP in that monitors are implemented in hardware rather than software, and that both approaches thus have no runtime overhead on the CPU. P2V, however, is more like the above approaches in that it is designed for monitoring actual programs rather than peripheral devices. Also it requires a dynamically extensible soft-core processor implemented on an FPGA, while our approach can potentially be applied to any COTS communication architecture. Further, P2V uses hardwired logic while BusMOP allows different formalisms.

VIII. Conclusions and Future Work

COTS peripherals are increasingly being adopted in the embedded market for performance reasons. However, COTS components introduce challenges in the development of critical systems, as they are unpredictable and often complete hardware specification is not publicly available. In this paper, we have proposed run-time monitoring of bus activities as a way to cope with such unpredictability. A monitoring device can be plugged on a PCI bus segment and check that all communication between peripherals and the rest of the system behaves according to specifications. Monitoring logic is automatically generated by the BusMOP framework and synthesized on FPGA, resulting in zero CPU runtime overhead. Finally, we showed the applicability of our monitoring infrastructure and recovery mechanisms on a real test case.

We plan to extend this work in two directions. From a system point of view, we plan to develop a interposing PCI/PCI-X/PCI-E monitoring device capable of executing preventive recovery actions as described in Section IV. From a formal specification point of view, we plan to extend BusMOP to support other MOP logic plugins. Most of them will require little work, with

the exception of context free grammars (CFG)[3], which would require implementing, effectively, a hardware LR(1) parser. This extension is not trivial, since memory is required for the stack, and the monitor must be able to process each event in few clock cycles. The unbounded nature of an LR(1) parser's per event memory usage makes the CFG plugin an unlikely candidate for running in a release system.

References

[1] P. Avgustinov, A. Christensen, L. Hendren, S. Kuzins, J. Lhotak, O. Lhotak, O. de Moor, D. Sereni, G. Sittampalam, and J. Tibble. ABC: an extensible AspectJ compiler. In *Proc. of the ACM Conf. on Aspect-oriented software development (ASOD'05)*, pages 87–98, 2005.

[2] P. Avgustinov, J. Tibble, and O. de Moor. Making trace monitors feasible. In *Proc. of the ACM Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'07)*, pages 589–608, 2007.

[3] H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-based runtime verification. In *Int. Conf. on Verification, Model Checking, and Abstract Interpretation (VMCAI'04)*, pages 277–306, 2004.

[4] BusMOP webpage. http://fsl.cs.uiuc.edu/BusMOP.

[5] F. Chen and G. Roşu. MOP: An Efficient and Generic Runtime Verification Framework. In *Proc. of the ACM Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'07)*, pages 569–588, 2007.

[6] D. Drusinksky. Temporal rover, 1997-2007.

[7] E.A. Emerson. *Handbook of Theoretical Computer Science*. MIT Press, 1990. Chapter 16: Temporal and modal logic.

[8] Eagle Technology. *PCI 703 Series User's Manual*. http://www.eagledaq.com/display_product_36.htm.

[9] K. Havelund and G. Rosu. Monitoring Java programs with Java pathexplorer. In *Proc. First Workshop on Runtime Verification*, 2001.

[10] K. Hoyme and K. Driscoll. Safebus(tm). *IEEE Aerospace Electronics and Systems Magazine*, pages 34–39, Mar 1993.

[11] M. Kim, M. Viswanathan, S. Kannan, I. Lee, and O. Sokolsky. Java-mac: A run-time assurance approach for Java programs. *Formal Methods in System Design*, 24(2):129–155, 2004.

[12] D. Knuth. Backus normal form vs. backus naur form. *Communications of the ACM*, 7(12):735–736, 1964.

[13] H. Lu and A. Forin. The design and implementation of p2v, an architecture for zero-overhead online verification of software programs. Technical Report MSR-TR-2007-99, Microsoft Research, 2007.

[14] M. Sipser. *Introduction to the Theory of Computation*. PWS Publishing, 1996. Chapter 1: Regular Languages.

[15] M. Martin, B. Livshits, and M. Lam. Finding application errors and security flaws using PQL: a program query language. In *Proc. of the ACM Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'05)*, pages 365–383, 2005.

[16] PCI SIG. *Conventional PCI 3.0, PCI-X 2.0 and PCI-E 2.0 Specifications*. http://www.pcisig.com.

[17] R. Pellizzoni, B. D. Buy, M. Caccamo, and L. Sha. Coscheduling of real-time tasks and PCI bus transactions. Technical report, University of Illinois at Urbana-Champaign, 2008. Available at http://netfiles.uiuc.edu/rpelliz2/www/techreps/.

[18] Xilinx, Inc. *Virtex-4 ML455 PCI/PCI-X Development Kit User Guide*. http://www.xilinx.com/support/documentation/boards_and_kits/ug084.pdf.

[3] More precisely, MOP supports DCFLs.