Using Formal Verification of Parameterized Systems in RAW Hazard Analysis in Microprocessors

Lukáš Charvát Aleš Smrčka Tomáš Vojnar
IT4Innovations Centre of Excellence, FIT, Brno University of Technology, Czech Republic

***Abstract*—Implementation of a pipeline-based execution of instructions in purpose-specific microprocessors is an error prone task, which implies a need of proper verification of the resulting design. Various techniques were proposed for this purpose, but they usually require a significant manual intervention of the developers. In this work, we propose a novel, highly automated approach for discovering RAW hazards in in-order pipelined instruction execution. Our approach combines static analysis of data paths to detect anomalies and possible hazards, followed by a transformation of detected problematic paths to a parameterised system (PS), and a subsequent formal verification to check the possibility of unhandled hazards using techniques for formal verification of PSs. We have implemented our approach and successfully applied it on multiple non-trivial microprocessors.**

I. Introduction

For many current, highly demanding embedded applications, rather complex purpose-specific microprocessors are developed, including those with pipelined execution. Development of such processors requires the designer to reason about mutual interference of sequences of instructions in the pipeline where tricky errors can easily arise. Automated or semi-automated support of the development process, including suitable verification techniques, is hence highly needed. Various techniques have been proposed for this purpose, but they are still rather limited in terms of their generality, scalability, and/or degree of automation.

Our long-term goal is to develop a set of verification techniques with formal roots, each of them specialised in checking absence of a certain kind of errors in purpose-specific microprocessors. The main idea is that, this way, a high degree of automation and scalability can be achieved since only parts of a design related to a specific error are to be investigated.

In our previous work [1], we proposed, with the above goal in mind, a fully automated approach for checking correctness of the implementation of individual instructions. In this paper, we aim at *read-after-write (RAW) hazards* in microprocessors with a single pipeline. An RAW hazard arises when an instruction writes to a storage that some later instruction reads, but it is possible for the later instruction to read an old value being rewritten by the earlier instruction.

Our approach for verifying that a single-pipeline microprocessor is free of RAW hazards starts by using a simple static analysis to find all data paths which could transfer data in a way causing an RAW hazard. Subsequently, we use SMT solving to check whether some of these paths could be enabled under some conditions. If so, we use parametric formal verification over a specially derived model of the data path to check whether there exists some sequence of instructions that could generate such conditions. If there is no such case, RAW hazards are handled properly by the processor.

Our approach concentrates mostly on the control parts of a design, for which current formal methods usually scale well, and minimizes the need of reasoning over data. We have implemented our approach in a prototype tool, and we present experimental evidence showing that our approach does indeed provide promising results in practice.

Plan of the Paper: Section II presents an overview of the related work addressing validation of single-pipeline microprocessors. Section III defines the needed notions. Section IV presents our verification approach, followed by experiments presented in Section V. Finally, Section VI concludes the paper. Due to space restrictions, a more detailed description of our approach is given in [2].

II. Related Work

Showing absence of RAW hazards is a native part of checking conformance between an RTL design and a formally encoded ISA description. The perhaps most cited approach to such checking is the so-called flushing technique [3], which has been extended, e.g., in [4], [5], [6], to handle rather complicated designs with multi-cycle execution units, exceptions, and branch prediction. The main challenge of these works is to overcome the semantic gap between the different levels of a processor description. Dealing with this issue typically requires a significant user intervention consisting in providing various additional assertions about the design or in transforming it to a purpose-specific description language.

In [7], the so-called self-consistency check that compares possible executions of each instruction in two scenarios is introduced. The comparison is made wrt. a property given by the user, e.g., a property concerning RAW hazards which deals with (i) executions of an instruction enclosed by any (random) instructions within the pipeline and (ii) executions of the same instruction surrounded by `NOP` instructions only. If the self-consistency check succeeds, conformance of the RTL and ISA descriptions of a processor can be established by separately showing conformance of the RTL/ISA descriptions of each individual instruction. The main drawback of the approach is that it requires the enclosing instructions from the first run not to violate a so-called consistent state of the microprocessor, which has to be manually defined by the user.

In [8], a formal model based on a notion of stages, parcels (instructions), and hazards has been introduced. Once the user defines predicates needed for describing the pipeline, the design can be automatically formally proven correct under

a correctness criterion given in the work. Another, a bit similar approach has been proposed in [9]. The approach introduces an abstract formal model whose components are to be linked by the user with the concrete cycle-accurate implementation through a number of mappings. Afterwards, IPC [10] is used to check several properties implying correctness of the pipeline behaviour. Again, both of the above methods require a significant manual user intervention.

Compared with the above approaches, we do not aim at full conformance checking between RTL and ISA implementations. Instead, we address one specific property—namely, absence of problems caused by RAW hazards. On the other hand, our approach is almost fully automated—the only step required from the user is to identify the architectural registers.

III. Preliminaries

A *processor structure graph* (PSG) is a tuple $G = (V, E, s, t)$. V is a finite set of vertices that is the union $V = V_s \cup V_f$ of a set V_s of *storages* and a set V_f of *Boolean circuits*, $V_s \cap V_f = \emptyset$. V_s includes five types of storages, meaning that $V_s = V_a \cup V_p \cup V_\mu \cup V_{rp} \cup V_{wp}$ where V_a, V_p, and V_μ are sets of *architectural*, *pipeline*, and *micro-architectural registers*, respectively, while V_{rp} and V_{wp} denote sets of *read* and *write ports of memories* (with all the sets being pairwise disjoint). We assume that architectural registers and memory ports are in the set of *programmer visible storages* $V_{pv} = V_a \cup V_{rp} \cup V_{wp}$. We also expect micro-architectural registers to only hold the state of processor's control logic and to be not used to carry data interchanged between programmer visible registers. Moreover, we expect all storages to have a unit write and zero read delay. Longer access times (e.g., for memory ports) can be modelled by introducing sequentially connected registers emulating the required delay. The set of Boolean circuits is the union of two types of circuits $V_f = V_{mx} \cup V_g$ where V_{mx} is a set of circuits implementing *multiplexers* and V_g is a set of the remaining (generic) circuits, $V_{mx} \cap V_g = \emptyset$.

Next, E denotes a finite set of *transfer edges*. Let $\mathbb{T} = \{\mathtt{d}, \mathtt{q}, \mathtt{st}, \mathtt{cl}, \mathtt{addr}, \mathtt{en}, \mathtt{sel}, \mathtt{m}\} \cup \{\mathtt{a}_i, \mathtt{case}_i \mid i \in \mathbb{N}\}$ be the set of *connection types* whose meaning will become clear below. Finally, $s : E \to V \times \mathbb{T}$ assigns to each edge its source vertex and its connection type, and $t : E \to V \times \mathbb{T}$ assigns to each edge its target vertex and its type of connection. The $\mathtt{d}, \mathtt{st}, \mathtt{cl}, \mathtt{addr}, \mathtt{en}$ types represent commonly used data, stall, clear, address, and enable connections of registers and memory ports with the usual semantics. The $\mathtt{q}$ type represents a data output of the given $v \in V$. The $\mathtt{m}$ type is a special type of connection that is used to interconnect write ports of the memories with their reading counterparts. The $\mathtt{a}_i$ types are argument connections of functional vertices $v_g \in V_g$. Finally, $\mathtt{sel}$ and $\mathtt{case}_i$ types are connections related to multiplexers only. The value transferred through $\mathtt{sel}$ connection selects which of $\mathtt{case}_i$ inputs is propagated to the $\mathtt{q}$ output of the multiplexer. A detailed semantics can be found in [2].

Because each vertex $v \in V$ can have at most one inbound edge for a single connection type, one can use a simpler notation to uniquely describe the edges. In particular, an edge $e \in E$ that satisfies $t(e) = (v, \mathtt{c})$, $v \in V$, can be encoded using the expression $v.\mathtt{c}$.

Given $k > 1$ and vertices $v_1, v_k \in V$ of a PSG, a walk from v_1 to v_k is an alternating sequence of vertices and edges $\langle v_1, e_1, v_2, ..., v_k \rangle$ where $v_2, ..., v_{k-1} \in V$, $e_1, ..., e_{k-1} \in E$, and every two subsequent vertices are incident with an edge present between them, i.e., $s(e_i) = (v_i, \mathtt{c_i})$, $t(e_i) = (v_{i+1}, \mathtt{c_{i+1}})$ for each $1 \leq i < k$ and $\mathtt{c_1}, ..., \mathtt{c_k} \in \mathbb{T}$. A path from v_1 to v_k is a walk where no vertex appears twice, i.e., $i \neq j \Rightarrow v_i \neq v_j$ for all $1 \leq i, j \leq k$.

Since our approach builds on analysing conditions that hold in certain stages of execution of a given instruction, we now introduce a notion of edge conditions. An *edge condition* is a pair (e, b), denoted $e \rightsquigarrow b$, meaning that the value transferred over the edge $e \in E$ is equal to some constant or variable b of bit-vector logic. By $\mathbb{E}$, we denote the set of all such edge conditions.

Our approach further uses the common notion of a parameterised system (PS) operating on a linear topology where processes (i.e., executed instructions) may perform local transitions or universally/existentially guarded global transitions [11], [12]. For our purposes, it is enough to consider global transitions only. A PS is a pair $P = (Q, \Delta)$ where Q is a finite set of states of a process and Δ is a set of transition rules over Q. A transition rule is of the form $\mathbb{Q} j < i \colon G \models q \to q'$ where $\mathbb{Q} \in \{\forall, \exists\}$, $G \subseteq Q$ and $q, q' \in Q$. A PS induces a transition system whose configurations are finite words over Q. A configuration $q_1 ... q_i ... q_n$, $1 \leq i \leq n$, changes to $q_1 ... q'_i ... q_n$ when the ith process goes from its state q_i to q'_i using some of the transition rules. The rule can be applied only if its guard is satisfied. For example, the meaning of the guard $\forall j < i \colon G$ is "for every process j to the left from i (in the linear topology), the jth process should be in a state that belongs to the set G".

We will work with the reachability problem given by a PS P, a regular set $I \subseteq Q^+$ of initial configurations, and a regular set $Bad \subseteq Q^+$ of bad configurations. In particular, we assume Bad to be given as the upward closure of a finite set $B \subseteq Q^+$ of minimal bad configurations, this is, $Bad = \{c \in Q^+ \mid \exists b \in B \colon b \sqsubseteq c\}$ where $\sqsubseteq$ is the usual sub-word relation (i.e., $u \sqsubseteq s_1 ... s_n \Leftrightarrow u = s_{i_1} ... s_{i_k}$ for some $1 \leq i_1 \leq ... \leq i_k \leq n$, $0 \leq k \leq n$). Now, let $R \subseteq Q^*$ denote the set of all reachable configurations. We say that the system P is safe wrt. I and Bad iff no bad configuration is reachable, i.e., $R \cap Bad = \emptyset$.

IV. The Proposed Verification Approach

This section describes our approach for verifying that the logic of data-flow control prevents RAW hazards, making it impossible for a later instruction to read incorrect (not yet updated) data, use them for a computation, and write them into some programmer visible storage. We expect the processor under verification to be described by a PSG, which can be easily obtained from a description of the processor on the register transfer level (RTL) written in common languages, such as VHDL or Verilog. For our experiments, we have, in particular, implemented a translation from RTL generated from the CodAL hardware description language [13] to PSGs.

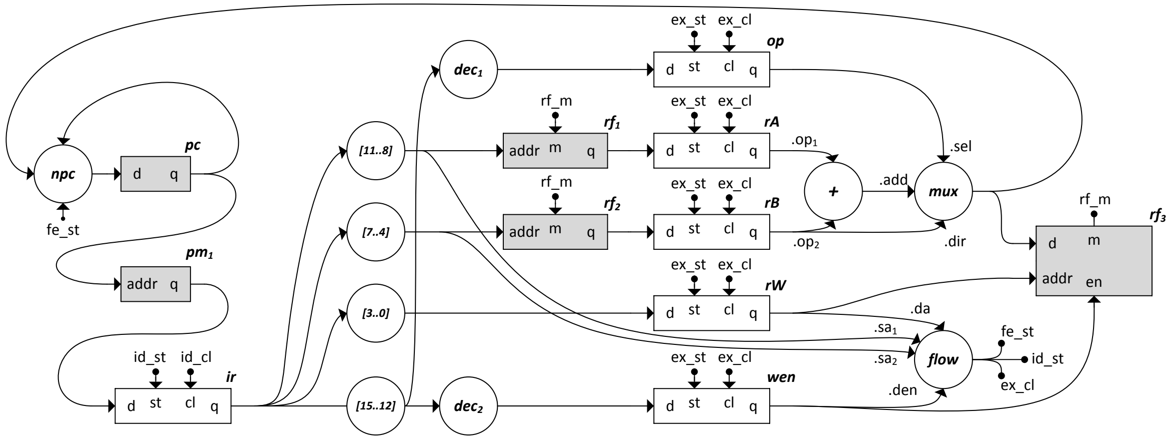

Fig. 1: A processor structure graph of a part of a RISC CPU.

Our approach consists of the following steps: (i) a data-flow analysis intended to distinguish particular stages of the pipeline, (ii) a consistency check of a correct implementation of the particular pipeline stages, (iii) a static analysis identifying constraints over data-paths of instructions that can potentially cause an RAW hazard, (iv) generation of a PS modelling mutual interaction between potentially conflicting instructions, and (v) an analysis of the constructed PS.

Example 1. Throughout the section, we will be illustrating the different steps of our approach on a running example depicted in Fig. 1. The figure shows a PSG describing a part of a simple RISC CPU. The depicted part of the CPU is used during execution of arithmetic and load/store instructions. To remain lucid, the PSG does not include any logic responsible for execution of branch instructions, and it shows only selected parts of an instruction decoder (circuits dec_1 and dec_2). Also, the write enable signal for the register pc as well as the read enable signals for the ports rf_1, rf_2 (read ports of register file rf), and pm_1 (which is a port to the program memory) are omitted because of their constant setting to "1".

The node $flow$ represents the flow logic of the controller which is responsible for dealing with RAW hazards on the register file. The flow logic implements the function $flow(sa_1, sa_2, den, da) = den \wedge (sa_1 = da \vee sa_2 = da)$ which checks the value of the enable den and the address da also connected to the write port rf_3 of an earlier instruction with the addresses sa_1, sa_2 of the read ports rf_1, rf_2 of a later instruction. In case the later instruction wants to read from the same address as the writing one, the flow logic uses stall and clear signals, transferred by edges such as $ir.\texttt{st}$ or $wen.\texttt{cl}$, to insert a pipeline bubble between the instructions. ◁

A. Data-Flow Analysis

The input of our approach (apart from the PSG) is the identification of architectural registers including the program counter. We use a simple data flow analysis to get the number of pipeline stages and the mapping of storages and logic functions into the pipeline stages. We define a *pipeline stage* as a sub-graph of a PSG responsible for executing a single-cycle step of an instruction. The pipeline stage that a vertex (representing some storage or function) of a PSG belongs to is given by the minimum number of cycles needed to propagate data from the input of the program counter (assumed to be read from a fictive stage 0) to the output of the given vertex. Hence, as a particular case, the program counter itself belongs to stage 1.

The simple data-flow analysis that we use starts from the program counter and its stage 1 and propagates the so-far computed stages forward through the PSG, always taking the minimum of values incoming to a vertex and adding one whenever a storage other than a read port (which has a zero delay) is passed. This analysis gives us the mapping $\varphi\colon V_s \to \mathbb{S}$, $\mathbb{S} = \{0, ..., n\}$, $n \in \mathbb{N}$, of storages to pipeline stages. Subsequently, by looking at all the storages from which there is a path to a given storage not passing through any further storage, we can easily get the *write stage* mapping $\varphi_{wr}\colon V_s \to 2^{\mathbb{S}}$ describing which stages directly influence the value of the given storage. In case of the program counter, we always add the fictive stage 0. Likewise, by looking at all target storages that can be reached from a given storage by a path not passing through any further storage, one can derive the *read stage* mapping $\varphi_{rd}\colon V_s \to 2^{\mathbb{S}}$ describing which stages outputs (directly connected to targeted storages) are influenced by the given storage. Pipeline stages of the storages from the PSG of Fig. 1 and the corresponding read and write stages, computed as described above, are shown in Table I.

B. Consistency Checking

The second step of our method is consistency checking which checks whether the flow logic assures a correct in-order execution of all instructions through all the identified pipeline stages. This means that all instructions which are fetched from the program memory should flow from the first stage to the last stage while maintaining their execution order with no loss or duplication of an instruction. We check whether the flow logic obeys a set of rules which express how the control connections ($\texttt{en}, \texttt{st}, \texttt{cl}$) of storages in adjacent pipeline stages should be set. In particular, we use a strengthened variant of the rules proposed in [14]. The rules are expressed as formulae in bit-vector logic and checked using an SMT solver. An example of a rule we check is the following: If some pipeline register of a stage s is stalled or not written, then all pipeline and

architectural registers and write ports of the stage s have to be stalled or not written. Due to a lack of space, a complete list of the rules, including their formal statement, is deferred to the technical report [2].

C. Static Detection of Potential RAW Hazards

In the next step, a static hazard analysis examines the PSG and the pipeline stage mappings φ, φ_{wr}, φ_{rd} determined by the data-flow analysis and identifies a finite set of so-called *hazard cases*. Each hazard case describes one possible source of an RAW hazard. In the construction of hazard cases, the most interesting step is the derivation of the so called *minimal influence path*. We define an *influence path* as a path $\langle v_1, e_1, ..., v_k\rangle$ where the value read from a programmer visible storage $v_1 \in V_{pv}$ can influence a value stored to an programmer visible storage $v_{pv} \in V_a \cup V_{wp}$ by writing to a *target* storage $v_k \in V_s$. Each influence path must fulfill the following set of properties: (i) The target storage v_k must be either (a) an architectural register or a write port, i.e., the case when $v_k = v_{pv}$, or (b) a pipeline register s.t. $t(e_{k-1}) = (v_k, \texttt{cl})$. Indeed, clearing of the pipeline register v_k will surely influence all programmer visible storages that belong to stages $s \geq \varphi(v_k)$. Next, (ii) the influence path must not traverse through stall connections of pipeline registers. Such paths cannot influence the value of any programmer visible register. Their only impact can be that they stall a stage but in such a case the previously established satisfaction of the consistency rules assures the correct conservation of all the partially executed instructions. Finally, (iii) access stages of elements along the path must conform to those obtained during the data-flow analysis and these stages have to form an increasing sequence. Otherwise, there could not be any instruction capable of a data transfer along the influence path.

An error in the RAW hazard prevention logic is manifested upon the first write of incorrect data into some programmer visible storage of the design. Therefore, it is sufficient to further work only with the minimal influence path which is an influence path where $v_i \notin V_a \cup V_{wp}$ and $t(e_{i-1}) \notin V_p \times \{\texttt{cl}\}$ for all $1 < i < k$. To discover the minimal influence paths in the given PSG, one can think of using a standard breadth-first search with the rules (i-iii) and the minimality checked on-the-fly.

A hazard case $(v_w, s_w, v_r, s_r, v_t, s_t, \pi)$ consists of (i) a programmer visible storage v_w (i.e., a register or a writing port of the memory), (ii) its write stage $s_w \in \varphi_{wr}(v_w)$, (iii) a reading storage v_r such that $v_r = v_w$ if v_r is a register or such that v_r and v_w are ports of the same memory, (iv) the read stage $s_r \in \varphi_{rd}(v_r)$ such that $s_r < s_w$ in order that the storage is read before it is written to evoke the RAW hazard, (v) a target storage v_t where the potentially incorrect value read from v_r is stored, (vi) a stage $s_t \in \varphi_{wr}(v_t)$, $s_r \leq s_t$, in which the incorrect value is stored, and (vii) a minimal influence path π describing how data are propagated from v_r to v_t between the stages s_r and s_t. Note that since the definition of a hazard case speaks about storages, their access stages, and the path along which the problematic data is transferred, it is not related to a single instruction only but to an entire class of instructions. Further, note that the case when $s_r = s_w$ is not included since the consistency checking guarantees an isolated execution of the instructions.

TABLE I: Pipeline stages and potential hazards.

Storage	Type	Stage φ	Write stages φ_{wr}	Read stages φ_{rd}	Potential hazard
pc	arch	1	$\{0,1,2,3\}$	$\{0,1\}$	✓
pm_1	port	–	$\emptyset$	$\{1\}$	×
ir	pipe	2	$\{1,2,3\}$	$\{0,1,2\}$	–
rf_1	port	–	$\emptyset$	$\{2\}$	✓
rf_2	port	–	$\emptyset$	$\{2\}$	✓
op	pipe	3	$\{2,3\}$	$\{0,3\}$	–
rA	pipe	3	$\{2,3\}$	$\{0,3\}$	–
rB	pipe	3	$\{2,3\}$	$\{0,3\}$	–
rW	pipe	3	$\{2,3\}$	$\{0,1,2,3\}$	–
wen	pipe	3	$\{2,3\}$	$\{0,1,2,3\}$	–
rf_3	port	4	$\{3\}$	$\emptyset$	✓

Example 2. Consider results of data-flow analysis computed for the PSG from Fig. 1 shown in Table I. In the table, one can see that there is a potential RAW hazard on storage rf because, for example, its read port rf_2 can be read in stage 2 ($\varphi_{rd}(rf_2) = \{2\}$), and its write port rf_3 can be written in stage 3 ($\varphi_{wr}(rf_3) = \{3\}$). Therefore, the subsequent verification will include a check whether a RAW hazard is indeed possible between two classes of instructions. The first one consists of instructions writing to rf_3 at address a in stage 3. The other includes instructions reading from rf_2 at the same address a in stage 2. In the PSG, there are several target storages, such as rf_3 and pc, where the reading instruction can store potentially incorrectly fetched data. Thus, multiple hazard cases need to be considered. For example, assume the rf_3 as a target. There are multiple minimal influence paths leading to rf_3, e.g., $\pi_1 = \langle rf_2,\ rB.\texttt{d},\ rB,\ +.\texttt{op}_2,\ +,\ mux.\texttt{add},\ mux,\ rf_3.\texttt{d},\ rf_3\rangle$ or $\pi_2 = \langle rf_2,\ rB.\texttt{d},\ rB,\ mux.\texttt{dir},\ mux,\ rf_3.\texttt{d},\ rf_3\rangle$. This means that the following two hazard cases need to be investigated: $(rf_3, 3, rf_2, 2, rf_3, 3, \pi_1)$ and $(rf_3, 3, rf_2, 2, rf_3, 3, \pi_2)$. An analogical process is then applied for pc as the target. ◁

D. Generation of a PS Modelling the Possible Hazards

We will now describe how the behaviour of the instructions given by constraints of a hazard case can be modelled and verified for correctness using a PS $P = (Q, \Delta)$. In the system P, we map n instructions in the pipeline to n processes in a linear array (with the earliest on the left). Initially, they are in a state saying that their execution has not started. Then, they proceed through individual stages of the pipeline during which they may interact with each other by the means of pipeline flow logic, e.g., an earlier instruction may force a later instruction to be stalled, or cleared. Finally, the instructions end up in a state denoting that they left the pipeline.

Let us have a hazard case $(v_w, s_w, v_r, s_r, v_t, s_t, \pi)$. In the system P, we model interactions among three classes of processes (and hence 3 types of instructions) $\mathbb{K} = \{w,\ rw,\ any\}$. The w-class includes every instruction that writes to storage v_w in stage s_w. The rw-class includes instructions that read from storage v_r in stage s_r, perform a data computation that involves the data path π, and write to storage v_t in stage s_t. The any-class instructions are used as pipeline

filler representing ANY other instruction. The set of states of a parametrized system P is then given by pairs $(k,s) \in \mathbb{K} \times \mathbb{S}$ where k given a type of an instruction and s gives the stage in which the instruction is currently executing. Hence, $Q \subseteq \mathbb{K} \times \mathbb{S}$, and we will use the notation q_s^k to denote a stage $(k,s) \in Q$. For a pipeline of length m, the sequence $q_0^k, ..., q_m^k$ records each step of a k-class instruction in the pipeline.

To simplify the following explanation of the transition relation Δ, we assume existence of a mapping ξ and predicates R^{st}, R^{cl}, R^{cf}. The mapping $\xi\colon Q \rightarrow 2^{\mathbb{E}}$ describes the behaviour of an instruction using a set of edge conditions that hold in the given state $q \in Q$. The construction of the mapping ξ is based on the hazard case and its major step is determination of edge conditions required to (i) perform a write to v_w in stage s_w, (ii) make a read from v_r in stage s_r, (iii) propagate data over the influence path π in stages $s_r, ..., s_t$[1], and (iv) make a write to v_t in stage s_t. The conditions found in (i) are used to form the mapping ξ for states of the w-class instruction while those from (ii–iv) are related to the rw-class.

Next, the predicates $R^{st}, R^{cl} \subseteq \mathbb{S} \times Q^2$, $R^{cf} \subseteq Q^2$ reflect the pipeline's control logic which disallows some instruction interleavings. The predicate $R^{st}(s, q_{s_1}^{k_1}, q_{s_2}^{k_2})$, resp. $R^{cl}(s+1, q_{s_1}^{k_1}, q_{s_2}^{k_2})$, evaluates to true for $s, s_1, s_2 \in \mathbb{S}$, $k_1, k_2 \in \mathbb{K}$ iff edge conditions $\xi(q_{s_1}^{k_1}) \cup \xi(q_{s_2}^{k_2})$ that hold in states $q_{s_1}^{k_1}$, $q_{s_2}^{k_2}$ lead to stalling, resp. clearing, of the stage s, resp. $s+1$. The predicate R^{cf} holds iff conditions $\xi(q_{s_1}^{k_1}) \cup \xi(q_{s_2}^{k_2})$ are in contradiction meaning that concurrent presence of instructions in states $q_{s_1}^{k_1}$, $q_{s_2}^{k_2}$ is impossible within the verified design. The needed reasoning over edge conditions to evaluate the predicates can be done automatically, e.g., by utilizing bit-vector solver [15].

Now, to retain in order execution, an k-class instruction yields a step $q_s^k \rightarrow q_{s+1}^k$, $s \in \mathbb{S}$, $k \in \mathbb{K}$ if there is no earlier instruction in stage $s+1$. Moreover, an instruction stays in the same state $q_{s_1}^{k_1}$, i.e., yields a step $q_{s_1}^{k_1} \rightarrow q_{s_1}^{k_1}$, $s_1 \in \mathbb{S}$, $k_1 \in \mathbb{K}$, in the case when there exists another instruction in state $q_{s_2}^{k_2}$, $s_2 \in \mathbb{S}$, $k_2 \in \mathbb{K}$, for which $R^{st}(s_1, q_{s_1}^{k_1}, q_{s_2}^{k_2})$ holds. An instruction in state $q_{s_1}^{k_1}$ is cleared if there exists another instruction in state $q_{s_2}^{k_2}$ causing $R^{cl}(s_1+1, q_{s_1}^{k_1}, q_{s_2}^{k_2})$ to be positive and simultaneously $R^{st}(s_1, q_{s_1}^{k_1}, q_{s_2}^{k_2})$ to be negative. The clearing of an instruction is represented by a transition $q_{s_1}^{k_1} \rightarrow q_{s_1+1}^{any}$, where $q_{s_1+1}^{any}$ denotes a state of an any-class instruction in stage s_1+1. As for the conflicting case, presence of an instruction in state $q_{s_1}^{k_1}$ is considered as spurious if there exists an instruction in a state $q_{s_2}^{k_2}$ so that $R^{cf}(q_{s_1}^{k_1}, q_{s_2}^{k_2})$ is positive. To avoid false alarms, a transition $q_{s_1}^{k_1} \rightarrow q_{s_1+1}^{any}$ is made in such a case. All of the above mentioned transition rules can be encoded into Δ by using existentially and universally guarded transitions. For technical details about the encoding and as well as description of the predicates R^{st}, R^{cl}, R^{cf} and the mapping ξ, we kindly refer reader to [2].

Example 3. Consider the first hazard case from Example 2. We will demonstrate the reasoning done in order to define value of the predicate $R^{st}(2, q_2^{rw}, q_3^w)$, that is, whether concurrent presence of two instructions in states q_2^{rw}, q_3^w implies stalling of stage 2. Presence of an w-class instruction in state q_3^w implies validity of the edge conditions $\xi(q_3^w) = \{rf_3.\texttt{en} \rightsquigarrow 1, rf_3.\texttt{addr} \rightsquigarrow a\}$ (meaning that a write to register a is enabled). Similarly, existence of rw-class instruction in state q_2^{rw} means that the edge conditions $\xi(q_2^{rw}) = \{rf_2.\texttt{en} \rightsquigarrow 1, rf_2.\texttt{addr} \rightsquigarrow a\}$ (enabling reading from register a) to hold.

Now, we take arbitrary representative pipeline storage from stage 2, e.g., register ir. This can be done because the requirement for stalling of all storages of the stalled stage is one of the previously checked consistency rules. Stalling of the stage 2 thus necessary means setting the value of its stall edge $ir.\texttt{st}$ to "1". The value of $ir.\texttt{st}$ is computed by *flow* circuit. Therefore, the value of the predicate $R^{st}(2, q_2^{rw}, q_3^w)$ corresponds to $\mathit{flow}(sa_1, sa_2, den, da) = den \wedge (sa_1 = da \vee sa_2 = da)$. Because $rf_3.\texttt{en} \rightsquigarrow 1 \in \xi(q_3^w)$ and both edges $\mathit{flow}.\texttt{den}$, $rf_3.\texttt{en}$ have the same source vertex wen and its connection q, one can derive that $\mathit{flow}.\texttt{den} \rightsquigarrow 1$. Similarly, because $rf_3.\texttt{addr} \rightsquigarrow a$, $rf_2.\texttt{addr} \rightsquigarrow a \in \xi(q_3^w) \cup \xi(q_2^{rw})$, and the pairs of edges $(rf_3.\texttt{addr}, \mathit{flow}.\texttt{da})$, resp. $(rf_2.\texttt{addr}, \mathit{flow}.\texttt{sa}_2)$, share common source vertex rW, resp. $[7..4]$, we can state $\mathit{flow}.\texttt{da} \rightsquigarrow a$ and $\mathit{flow}.\texttt{sa}_2 \rightsquigarrow a$. Thus, $\mathit{flow}(sa_1, a, 1, a) = 1 \wedge (sa_1 = a \vee a = a)$ and so $R^{st}(2, q_2^{rw}, q_3^w)$ surely evaluates to true. $\triangleleft$

Initially, any instruction of w, rw as well as any class can enter the pipeline. Therefore, the regular set I of initial states is $(\{w, rw, any\} \times \{0\})^+$. In order to describe bad configurations, we have to determine the number of cycles h, during which an rw-class instruction must not store a value into v_t because it would necessary mean that the written data were incorrectly fetched. From the hazard case, we know that data supposed to be written to v_t are computed in stage s_t, and the computed value is *committed* to v_t in the next cycle, i.e., in stage s_t+1. To ensure that data computed in stage s_t are correct, no write to v_w can occur for $s_t - s_r$ cycles. Otherwise, during the execution of an rw-class instruction, there would be an update of storage v_r that is read by the instruction and thus the instruction would perform computation with incorrect data. Because a RAW hazard is exhibited only after commitment of the incorrectly fetched data to v_t, the value of h is given by $(s_t+1) - s_r$. Hence, for w-class instructions, we consider h subsequent states following the one involving write to v_w, i.e., $q_{s_w+1}^w, ..., q_{s_w+h}^w$, as *hazardous*. A configuration is considered as bad if it includes an occurrence of a hazard state followed by a state $q_{s_t+1}^{rw}$ of an rw-class instruction.

The reachability problem defined by the parametrized system P and by the sets of initial and bad states can be checked using techniques described, e.g., in [12], [16].

Example 4. Consider the first hazard case described in Example 2 to be present in five staged CPU. The execution of a w-class instruction writing to register file port rf_3 is described by a process going through the sequence of states $q_0^w, q_1^w, q_2^w, q_3^w, q_4^w, q_5^w, q_6^w$ where rf_3 is written in state q_3^w and q_0, q_6 denote initial, resp. final, state. The execution of the rw-class instructions reading from port rf_2 and writing to port rf_3 passed through the sequence of states

[1]This step includes, e.g., a calculation of edge conditions that must hold for selectors of multiplexers of the path π.

TABLE II: Verification times.

Processor / variant		Data flow analysis [s]	Consistency checking [s]	Par. Mod. verif. [s]	Total [s]	Hazard Cases [#]
TinyCPU	S	8	1	12	21	9
	SF	10	3	21	34	19
	B	8	1	13	22	9
SPP8	S	28	6	62	96	40
	B	29	6	67	102	40
SPP16	S	34	7	89	130	58
	B	33	7	97	137	58
Codea2	SFH	162	18	856	1036	281
DLX5	S	77	26	378	481	43
	B	123	29	590	742	44

S stalling logic B bypassing logic F flag reg. H special reg.

$q_0^{rw}, q_1^{rw}, q_2^{rw}, q_3^{rw}, q_4^{rw}, q_5^{rw}, q_6^{rw}$. Here, the port rf_2 is read in state q_2^{rw}, and port rf_3 is written in state q_3^{rw} with its value committed in state q_4^{rw}. Because the value of rf_3 is committed in state q_4^{rw}, i.e., in stage 4, and rf_2 is read in state q_2^{rw}, i.e., in stage 2, the length h is $4-2=2$ causing states q_4^{w}, q_5^{w} to be hazard states. Thus, the set B of minimal bad configurations is $\{q_4^w q_4^{rw}, q_5^w q_4^{rw}\}$. A chosen parametric verification method can then be used to check whether a bad configuration, e.g., $q_6^{any} q_5^w q_4^{rw} q_3^{any} q_2^{any} q_1^{any} q_0^{any}$, is reachable. ◁

V. Experimental Evaluation

We have implemented the above described method in a prototype tool [17] and tested it on five processors: *TinyCPU* is a small 8-bit processor that we mainly use for testing of new verification methods. *SPP8* is an 8-bit ipcore with 3 pipeline stages, 16 general-purpose registers, and a RISC instruction set consisting of 9 instructions. *SPP16* is a 16-bit variant of the previous processor with a more complex memory model. *Codea2* is a 16-bit processor dedicated for signal processing applications [18]. It is equipped with 16 general-purpose registers, 15 special registers, a flag register, and an instruction set including 41 instructions where each may use up to 4 available addressing modes. Finally, *DLX5* is a 5-staged 32-bit processor able to execute a subset of the instruction set of the DLX architecture [19] (without floating point instructions). Some of the processors were in multiple variants that differ from each other, e.g., in the way how RAW hazards are avoided or in having some additional instruction specific registers such as a register for storing MSB bits of the product, or flag registers.

We conducted series of experiments on a PC with Intel Core i7-3770K @3.50GHz and 16 GB RAM with results presented in Table II. The columns give the verified processor, its variant, the time needed for the data flow analysis, the duration of the consistency checking, the time spent by verification of the PSs that are created based on each hazard case, and the overall verification time. The last column represents the number of hazard cases that had to be verified during the model verification phase. Note that each hazard case represents a separate task so the part of model verification can be parallelized.

By verifying the above processors, we identified a flaw in a RAW hazard resolution when accessing of the data memory in a development version of SPP8 processor.

VI. Conclusion

We have presented an approach that combines data-flow analysis and methods for formal verification of PSs in order to discover incorrectly handled RAW hazards in the RTL implementation of pipelined microprocessors. The approach was developed with the aim to be highly automated, not requiring any additional efforts from the developers (apart from specifying the architectural registers). We have implemented the approach and successfully tested it on several non-trivial microprocessors where the approach was able to discover previously unknown flaws caused by unhandled hazards.

In the future, we plan to complement the approach proposed in the paper by techniques suitable for verification of other processor features, such as write-after-write and control hazards. This is motivated by our general idea of trying to split processor verification into several simpler, specialised tasks.

Acknowledgement: This work was supported by the Czech Science Foundation under the project 14-11384S, the EU/Czech IT4Innovations Centre of Excellence project CZ.1.05/1.1.00/02.0070, and the internal BUT projects FIT-S-12-1 and FIT-S-14-2486.

References

[1] L. Charvát, A. Smrčka, and T. Vojnar, "Automatic Formal Correspondence Checking of ISA and RTL Microprocessor Description," in *Proc. of MTV'12*. IEEE, 2012.

[2] ——, "Using Formal Verification of Parameterized Systems in RAW Hazard Analysis in Microprocessors," www.fit.vutbr.cz/research/groups/verifit/tools/hades/techrep/, Brno University of Technology, technical report FIT-TR-2014-04, 2014.

[3] J. R. Burch, D. L. Dill, "Automatic Verification of Pipelined Microprocessor Control," in *Proc. of CAV'94*, LNCS 818. Springer, 1994.

[4] A. Koelbl, R. Jacoby, H. Jain, C. Pixley, "Solver Technology for System-level to RTL Equivalence Checking," in *Proc. of DATE'09*. IEEE, 2009.

[5] M. N. Velev, P. Gao, "Automatic Formal Verification of Multithreaded Pipelined Microprocessors," in *Proc. of ICCAD'11*. IEEE, 2011.

[6] K. Hao, S. Ray, and F. Xie, "Equivalence Checking for Function Pipelining in Behavioral Synthesis," in *Proc. of DATE'14*. IEEE, 2014.

[7] R. B. Jones, C. H. Seger, and D. L. Dill, "Self-consistency Checking," in *Proc. of FMCAD'96*, LNCS 1166. Springer, 1996.

[8] M. D. Aagaard, "A Hazards-based Correctness Statement for Pipelined Circuits," in *Proc. of CHARME'03*, LNCS 2860. Springer, 2003.

[9] U. Kuhne, S. Beyer, J. Bormann, and J. Barstow, "Automated Formal Verification of Processors Based on Architectural Models," in *Proc. of FMCAD'10*. IEEE, 2010.

[10] M. Ngyuen, M. Thalmaier, M. Wedler, J. Bormann, D. Stoffel, and W. Kunz, "Unbounded protocol compliance verification usign interval property checking with invariants," *IEEE Transactions on Computer-Aided Design of Integrated Circuits*, vol. 27, no. 11, 2008.

[11] E. Clarke, M. Talupur, H. Veith, "Environment Abstraction for Parameterized Verification," in *Proc. of VMCAI'06*, LNCS 3855. Springer, 2006.

[12] P. A. Abdulla, F. Haziza, and L. Holík, "All for the Price of Few (Parameterized Verification through View Abstraction)," in *Proc. of VMCAI'13*, LNCS, 7737. Springer, 2013.

[13] *CodAL Architecture Description Language*, www.codasip.com/products/codal/, Codasip Ltd., 2013.

[14] P. Mishra, H. Tomiyama, N. Dutt, A. Nicolau, "Automatic Verification of In-order Execution in Microprocessors with Fragmented Pipelines and Multicycle Functional Units," in *Proc. of DATE'02*. IEEE, 2002.

[15] R. Brummayer, A. Biere, "Boolector: An Efficient SMT Solver for Bit-vectors and Arrays," in *Proc. of TACAS'09*, LNCS 5505. Springer, 2009.

[16] A. Bouajjani, P. Habermehl, and T. Vojnar, "Abstract Regular Model Checking," in *Proc. of CAV'04*, LNCS 3114. Springer, 2004.

[17] "Hades Hardware Verification Tool," www.fit.vutbr.cz/research/groups/verifit/tools/hades/, 2014.

[18] "Codea2 Core IP in Codasip Studio," www.codasip.com/products/codea2/, Codasip Ltd., 2013.

[19] D. A. Patterson, J. L. Hennessy, *Computer Organization and Design: The Hardware/Software Interface*. Boston: Morgan Kaufmann, 2012.