Higher Order Logic Theorem Proving and Its Applications
Melham, T. and Camilleri, J. (Eds.)
Lecture Notes in Computer Science 854
Springer Verlag, 1994, pp. 424-439

Implementational Issues for Verifying RISC-Pipeline Conflicts in HOL

Sofiène Tahar[1] and Ramayya Kumar[2]

[1] University of Karlsruhe, Institute of Computer Design and Fault Tolerance (Prof. D. Schmid), P.O. Box 6980, 76128 Karlsruhe, Germany
e-mail: tahar@ira.uka.de

[2] Forschungszentrum Informatik, Department of Automation in Circuit Design, Haid-und-Neu Straße 10-14, 76131 Karlsruhe, Germany
e-mail: kumar@fzi.de

Abstract. We outline a general methodology for the formal verification of instruction pipelines in RISC cores. The different kinds of conflicts, i. e. resource, data and control conflicts, that can occur due to the simultaneous execution of the instructions in the pipeline have been formally specified in HOL. Based on a hierarchical model for RISC processors, we have developed a constructive proof methodology, i.e. when conflicts at a specific abstraction level are detected, the conditions under which these occur are generated and explicitly output to the designer, thus easing their removal. All implemented specifications and tactics are kept general, so that the implementation could be used for a wide range of RISC cores. In this paper, the described formalization and proof strategies are illustrated via the DLX RISC processor.

1 Introduction

In this paper, we concentrate on the formalization and the correctness proofs of instruction pipelines in RISC cores. The previous work (from other researchers) in verifying RISC processors could verify only parts of processors, at certain levels of abstraction [2, 11]. Our work is distinguished by the fact that an overall methodology for the verification of RISC cores, at all levels of abstraction, is being developed [13, 14, 15].

There are three kinds of conflicts which occur in instruction pipelines, namely, resource, data and control conflicts. These conflicts occur due to the data and control dependencies and resource contentions when many instructions are simultaneously executed in the pipeline. The formalization of all possible conflicts and the description of the automated proof techniques for conflict detections, is given in this paper. The proof techniques that are given are constructive, i.e. the conditions under which the conflicts occur are explicitly stated, so that the designer can easily formulate the conflict resolution mechanisms either in hardware or generate software constraints which have to be met.

The organization of this paper is as follows. Section 2 briefly presents the hierarchical model on which the formal verification methodology is based. Section 3 gives an overview of the overall verification process. Section 4 includes some formal definitions of types and functions used in the formal specifications that follow. Sections 5 through 7 define the conflicts arising due to the pipelined nature of RISCs and describe the correctness proofs for resource, data and control conflicts, respectively. Section 8 contains some experimental results and section 9 concludes the paper. It is to be noted, that all the methods and techniques presented in this paper are illustrated by means of a RISC example — DLX[1] [6].

2 RISC Verification Model

Our RISC interpreter model has been adapted from the interpreter model of Anceau [1], which has been effectively used for the design, description and verification of microprogrammed processors [7, 17].

The adapted RISC interpreter model comprises the instruction, class, stage and phase levels, each of which corresponds to an interpreter at different abstraction levels, and the lowest level which corresponds to the circuit implementation — EBM (figure 1). Each interpreter consists of a set of visible states and a state transition function which defines the semantics of the interpreter at that level of abstraction.

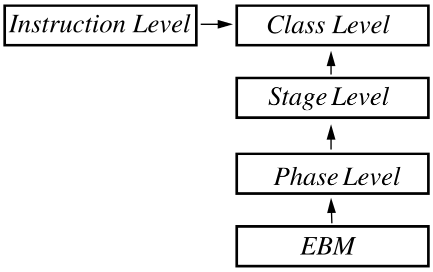

Fig. 1. RISC Interpreter Model

These interpreter levels incorporate structural and temporal abstractions, which are exploited in the proof process. Table 1 lists the set of transfers which occur at the stage and phase levels of the DLX processor, where the rows and columns represent the pipeline stages and the instruction classes, respectively[2].

Table 1. DLX Pipeline Structure

	ALU	ALU_I	LOAD	STORE	CONTROL
IF	IR ← M[PC] PC ← PC+4	IR ← M[PC] PC ← PC+4	IR ← M[PC] PC ← PC+4	IR ← M[PC] PC ← PC+4	IR← M[PC] PC ← PC+4
ID	A $\underset{\phi_2}{\leftarrow}$ RF[rs1] B $\underset{\phi_2}{\leftarrow}$ RF[rs2] IR1 ← IR	A $\underset{\phi_2}{\leftarrow}$ RF[rs1] B $\underset{\phi_2}{\leftarrow}$ RF[rs2] IR1 ← IR	A $\underset{\phi_2}{\leftarrow}$ RF[rs1] B $\underset{\phi_2}{\leftarrow}$ RF[rs2] IR1 ← IR	A $\underset{\phi_2}{\leftarrow}$ RF[rs1] B $\underset{\phi_2}{\leftarrow}$ RF[rs2] IR1 ← IR	BTA $\underset{\phi_1}{\leftarrow}$ f (PC) PC $\underset{\phi_2}{\leftarrow}$ BTA
EX	Aluout ← A op B	Aluout←A op (IR1)	MAR ← A+(IR1)	MAR← A+(IR1) SMDR ← B	
MEM	Aluout ← Aluout1	Aluout ← Aluout1	LMDR←M[MAR]	M[MAR]←SMDR	
WB	RF[rd] $\underset{\phi_1}{\leftarrow}$ Aluout1	RF[rd] $\underset{\phi_1}{\leftarrow}$ Aluout1	RF[rd] $\underset{\phi_1}{\leftarrow}$ LMDR		

[1] DLX is an hypothetical RISC which includes the most common features of existing RISCs such as Intel i860, Motorola M88000, Sun SPARC or MIPS R3000.

[2] For more details about this model, please refer to [13, 14, 15].

3 Verification Tasks

Starting from the instruction set of a RISC processor, we have to show that this instruction set is executed correctly by the EBM, in spite of the pipelined architecture. A RISC processor executes n_s instructions in parallel (see figure 2), in n_s different pipeline stages. This parallel execution increases the overall throughput of the processor; however, it should be noted that no single instruction runs faster. Hence, we have to prove, on the one hand, that the sequential execution of each instruction is correctly implemented by the EBM and on the other hand, to show that the pipelined sequencing of the instructions is executed correctly. Thus the correctness proof is split into two independent steps as follows:

1. given some software constraints on the actual architecture and given the implementation EBM, we prove that any sequence of instructions is correctly pipelined, i.e.:

$$\text{SW\_Constraints, EBM} \vdash \text{Correct\_Instr\_Pipelining} \qquad \textbf{(1)}$$

2. we prove that the EBM implements the semantic of each single architectural instruction correctly, i.e.:

$$\vdash \text{EBM} \Rightarrow \textit{Instruction Level} \qquad \textbf{(2)}$$

The software constraints in (1) are those conditions which are to be met for designing the software so as to avoid conflicts, e.g. the number of delay slots to be introduced between the instructions while using a software scheduling technique. Additionally, it is also assumed that the EBM includes some conflict resolution mechanisms in hardware.

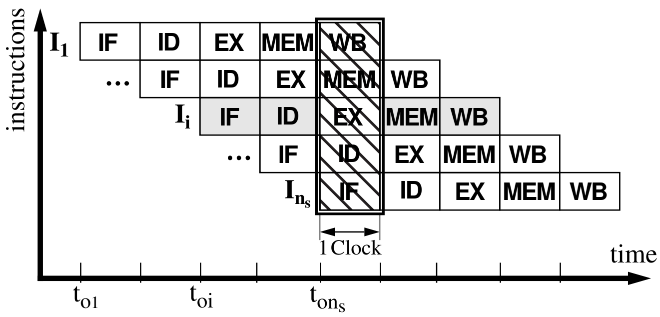

Fig. 2. Pipelined Execution

The proof of step (2) has been discussed and reported in our previous work [14, 15] and we recapitulate it briefly in the next subsection. Step (1) is the main topic of this paper and will be discussed in detail in the rest of the paper.

3.1 Semantical Correctness

The higher-order logic specifications and implementations at the various levels of abstraction are used and the following is proved — *EBM* ⇒ *Phase Level* ⇒ *Stage Level* ⇒ *Class Level*. This proof is similar to the one done by Windley in proving the correctness of microprogrammed processors [17]. Each obtained theorem is then instantiated for each instruction at the instruction level. We have implemented tactics in the HOL theorem prover [4], using the hardware verification environment —

MEPHISTO [9], which automates the entire process, given the formal definitions of the specifications and implementations at each abstraction level [15].

3.2 Pipeline Correctness

Step (1) can be proved by showing that all possible combinations of n_s instructions are executed correctly. This implies that the parallel execution of n_s instructions does not lead to any *conflicts*. There are three kinds of conflicts (also called *hazards*) that can occur during the pipelined execution of any RISC machine namely, resource, data and control conflicts[3]. The predicate Correct_Instr_Pipelining in (1), which implies the absence of these conflicts, is defined as the conjunction of predicates for each conflict kind. Formally:

$$\vdash_{def} \text{Correct\_Instr\_Pipelining:=} \quad \neg\text{Res\_Conflict} \wedge \neg\text{Data\_Conflict} \wedge \neg\text{Control\_Conflict}$$

and the correctness statement (1) can be rewritten as:

$$\text{SW\_Constraints, EBM} \vdash \neg\text{Res\_Conflict} \wedge \neg\text{Data\_Conflict} \wedge \neg\text{Control\_Conflict}$$

thus the whole correctness proof is tackled by splitting it into three independent parts, each corresponding to one kind of conflict. These parts are elaborated in the sections to follow.

In proving step (1), we have to ensure that all possible combinations of instructions occurring in n_s stages are executed correctly. This large number can be drastically reduced by exploiting the notion of classes which results in a combination of a smaller number. When conflicts are formalized at the class, stage and phase levels, the existence of multiple instructions in the pipeline can be formalized by predicates which we call the *multiple conflict predicates*. Their proofs are further simplified by constructing conflict predicates between pairs of instructions which are called the *dual conflict predicates*. These notions will be clarified in the sections to follow.

3.3 Different Approaches for Verification

The hierarchical model described above has been used in our previous work in hierarchically structuring the proof of the semantical correctness [15]. While formalizing this process, we discovered that the proofs can be hierarchically managed in a top-down or bottom-up manner, so that a *verification-driven design* or a *post-design verification* can be performed. By a top-down verification-driven design, we mean that the verification and design process are interleaved, so that the verification of a current design status, against the specification, yields the necessary constraints for the future design steps. A post-design verification is verification in the normal sense which is performed in a top-down or bottom-up manner after the entire design is completed at all levels. In a related work [16], we handle the verification of pipeline conflicts in a post-design manner, in that the verification of each kind of conflict is done more or less in a single step. Within the scope of this paper, the correctness proofs will be handled by means of the top-down verification-driven design methodology.

[3] A formal definition of these conflicts is given in later sections.

4 Formal Definitions

In our previous papers, we have given a detailed formalization of instructions at all abstraction levels [13, 15]. In this section, we briefly introduce some new types and functions that are useful for handling conflict predicates. According to our hierarchical model, we define for each abstraction level a set of enumeration types for corresponding instructions, resources and pipeline characteristics, i.e. pipeline stages or clock phases. Referring to the pipeline structure of table 1, where the columns represent the five instruction classes — ALU, ALU_I, LOAD, STORE and CONTROL and the rows correspond to the related stage and phase transfers at the five pipeline stages — IF, ID, EX, MEM and WB and the two clock phases ϕ_1 and ϕ_2, respectively, we define the following enumeration types for the DLX example:

- types for pipeline stages and clock phases:

  `define_type` "pipeline_stage = IF | ID | EX | MEM | WB"
  `define_type` "clock_phase = Ph1 | Ph2"

- types for the set of all instructions at each abstraction level:

  `define_type` "class_instruction = ALU | ALU_I | LOAD | STORE | CONTROL"
  `define_type` "stage_instruction = IF_X | ID_X | ID_C | EX_A | EX_I | … | WB_L"
  `define_type` "phase_instruction = Ph1_IF_X | … | Ph1_EX_A | … | Ph2_WB_L"

- types for resources (related to the structural abstraction):

  `define_type` "CL_resource = PC | RF of RF_addr | I_MEM | D_MEM | IAR"
  `define_type` "SL_resource = PC | RF of RF_addr | I_MEM | … | IR | A | B | Aluout | … "
  `define_type` "PL_resource = PC | RF of RF_addr | … | IR | A | B | Aluout | … | BTA"

For the specification of the conflict predicates, we have defined several functions and predicates as follows:

- abstraction functions, which either compute higher level instructions from lower ones or extract lower level instructions from higher ones:

  ClassToStage : ((pipeline_stage, class_instruction) $\rightarrow$ stage_instruction)
  StageToClass : (stage_instruction $\rightarrow$ class_instruction)
  StageToPhase : ((clock_phase, stage_instruction) $\rightarrow$ phase_instruction)
  PhaseToStage : (phase_instruction $\rightarrow$ stage_instruction)

  e.g. ClassToStage (ID, CONTROL) = ID_C, PhaseToStage (Ph2_ID_C) = ID_C.

- functions that compute the logical pipeline stage or clock phase types from a stage or a phase instruction, respectively:

  Stage_Type : (stage_instruction $\rightarrow$ pipeline_stage)
  Phase_Type : (phase_instruction $\rightarrow$ clock_phase)

- functions which compute the ordinal values of a given pipeline stage and clock phase, respectively[4]:

  Stage_Rank : (pipeline_stage $\rightarrow$ num)
  Phase_Rank : (clock_phase $\rightarrow$ num)

  e.g. Stage_Rank (ID) = 2, Phase_Rank (Ph1) = 1.

- predicates, which are true if a given resource is used by a given stage and phase instruction, respectively:

[4] These functions are needed to express the sequential order of the execution of stage and phase instructions.

Stage_Used : ((stage_instruction, SL_resource) $\rightarrow$ bool)
Phase_Used : ((phase_instruction, CL_resource) $\rightarrow$ bool)

e.g. Stage_Used (ID_C, PC) = *True*. These Predicates are extracted from the formal specifications of the stage and phase levels, respectively (refer to [15]).

- predicates that imply that a given resource is read (*domain*) or written (*range*) by a given class or stage instruction at a given pipeline stage or clock phase, respectively:

Stage_Domain : ((class_instruction, pipeline_stage, CL_resource) $\rightarrow$ bool)
Stage_Range : ((class_instruction, pipeline_stage, CL_resource) $\rightarrow$ bool)
Phase_Domain : ((stage_instruction, clock_phase, CL_resource) $\rightarrow$ bool)
Phase_Range : ((stage_instruction, clock_phase, CL_resource) $\rightarrow$ bool)

e.g. Stage_Domain (ALU, ID, RF) = *True*, Phase_Range (ID_C, Ph2, D_MEM) = *False* (refer also to table 1). These predicates are extracted from the specification of the class and stage level instructions at the clock cycle and clock phase time granularities, respectively (refer to [15]).

5 Resource Conflicts

Resource conflicts (also called *structural hazards* [6, 8, 10, 12] or *collisions* [8, 12]) arise when the hardware cannot support all possible combinations of instructions during simultaneous overlapped execution. This occurs when some resources or functional units are not duplicated enough and two or more instructions attempt to use them simultaneously. A resource could be a register, a memory unit, a functional unit, a bus, etc. The use of a resource is a write operation for storage elements and an allocation for functional units. In the sections to follow, we will first formally specify the resource conflicts and then discuss the correctness proof issues.

5.1 Resource Conflict Specification

Referring to the hierarchical model (see section 2), the formal specifications of resource conflicts are handled according to the different abstract levels. Further, related to each abstraction level only the visible resources are considered by the corresponding resource conflict predicates.

Class Level Conflicts. The resource conflict predicate Res_Conflict, as mentioned in section 3.2, is equivalent to a multiple conflict between the maximal number of class instructions that occur in the pipeline. This Multiple_Res_Conflict predicate is true if any pair of the corresponding stage instructions compete for one resource (see hatched box in figure 2). Formally, Multiple_Res_Conflict is defined in terms of disjunctions over all possible stage instruction pair conflicts (assigned by Dual_Stage_Conflict), which correspond to the class instructions $I_1 \ldots I_{n_s}$:

$$\text{Res\_Conflict} = \vdash_{def} \text{Multiple\_Res\_Conflict}\,(I_1, \ldots, I_{n_s}) := \bigvee_{\substack{i,j \\ (i,j = 1 \ldots n_s) \\ (i < j)}} \text{Dual\_Stage\_Conflict}\,(\text{ClassToStage}\,(\psi_i{}^{5}, I_{n_s - i+1}), \text{ClassToStage}\,(\psi_j, I_{n_s - j+1}))$$

[5] Depending on the index i, ψ_i represents the related pipeline stage, i.e. ψ_1 = IF, ψ_2 = ID, ψ_3 = EX, etc.

Stage Level Conflicts. A dual resource conflict happens when two stage instructions attempt to use the same resource. Further, since only stage instructions of different types could be executed simultaneously in the pipeline (see hatched box in figure 2), we should ensure that the checked stages are of different logical types. Formally, the Dual_Stage_Conflict predicate is specified using the function Stage_Type and the predicate Stage_Used as follows:

$$
\begin{aligned}
&\vdash_{def}\ \text{Dual\_Stage\_Conflict}\ (S_i, S_j) :=\\
&\qquad \exists\, r\!: \text{SL\_resource}.\\
&\qquad\quad \text{Stage\_Type}\ (S_i) \neq \text{Stage\_Type}\ (S_j) \wedge\\
&\qquad\quad \text{Stage\_Used}\ (S_i, r) \wedge \text{Stage\_Used}\ (S_j, r)
\end{aligned}
$$

Looking closer, even when the predicate Dual_Stage_Conflict is true, a conflict occurs only if the stage instructions S_i and S_j use the resource r at the same phase of the clock, since a multi-phased non-overlapping clock is used (figure 3).

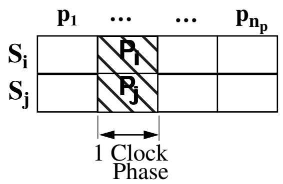

Fig. 3. Phase Parallel Execution

Having an implementation of the stage instructions at the phase level and considering all combinations of phase instructions for any two stage instructions, the dual stage conflict is defined at this lower level in terms of a multiple phase conflict predicate. Formally, Multiple_Phase_Conflict is defined in form of disjunctions over all possible phase instruction pair conflicts (assigned by Dual_Phase_Conflict):

$$
\begin{aligned}
&\text{Dual\_Stage\_Conflict} =\\
&\vdash_{def}\ \text{Multiple\_Phase\_Conflict}\ (S_i, S_j) :=\\
&\qquad \bigvee_{\substack{k\\(k=1\ldots n_p)}} \text{Dual\_Phase\_Conflict}\ (\text{StageToPhase}\ (\phi_k{}^{6}, S_i), \text{StageToPhase}\ (\phi_k, S_j))
\end{aligned}
$$

Phase Level Conflicts. A dual resource conflict at the phase level occurs only when any two phase instructions that compete on the same resource, are of the same phase type, i.e. the same clock phase is involved (see figure 3) and belong to stage instructions of different types. Using the functions Phase_Type, Stage_Type and StageToPhase and the predicate Phase_Used, this is formally given as follows:

$$
\begin{aligned}
&\vdash_{def}\ \text{Dual\_Phase\_Conflict}\ (P_i, P_j) :=\\
&\qquad \exists\, r\!: \text{PL\_resource}.\\
&\qquad\quad \text{Phase\_Type}\ (P_i) = \text{Phase\_Type}\ (P_j) \wedge\\
&\qquad\quad \text{Stage\_Type}\ (\text{PhaseToStage}\ (P_i)) \neq \text{Stage\_Type}\ (\text{PhaseToStage}(P_j)) \wedge\\
&\qquad\quad \text{Phase\_Used}\ (P_i, r) \wedge \text{Phase\_Used}\ (P_j, r)
\end{aligned}
$$

[6] According to the index k, ϕ_k represents a clock phase, i.e. ϕ_1 = Ph1, ϕ_2 = Ph2.

5.2 Resource Conflict Proof

Our ultimate goal is to show that for all class instruction combinations, no resource conflicts occur, i.e. the predicate Multiple_Res_Conflict is never true:

set_goal "$\forall I_1 \ldots I_{n_s}$: class_instruction . $\neg$Multiple_Res_Conflict ($I_1, \ldots, I_{n_s}$)"

Using the definition of Multiple_Res_Conflict (cf. section 5.1), the expansion of this goal at the stage level yields to a case explosion. In order to bypass this intractable goal, we do the proof in two steps as follows:

1. we prove the non happening of dual conflicts:

set_goal "$\forall$ S_i S_j : stage_instruction . $\neg$Dual_Stage_Conflict (S_i, S_j)"

2. we conclude the negation of the multiple conflict predicate from the first step:

set_goal "($\forall$ S_i S_j : stage_instruction . $\neg$Dual_Stage_Conflict (S_i, S_j)) $\Rightarrow$ ($\forall I_1 \ldots I_{n_s}$: class_instruction. $\neg$Multiple_Res_Conflict ($I_1, \ldots, I_{n_s}$))"

Since the dual conflict predicate is a generalization of the multiple conflict predicate, the proof of the second step is straightforward; we even do not need to expand the dual conflict predicate. The proof of the first step, without any assumptions, leads either to *True* or to a number of subgoals which explicitly include a specific resource and the specific stage instructions which conflict. For example a conflict due to the resource PC between the common IF-stage instruction (IF_X) and the ID-stage instruction (ID_C) of control class is output as follows:

(S_i = IF_X), (S_j = ID_C), (r = PC) $\vdash$ F

As mentioned above, a resource conflict is only ensured when the given phase level implementation still confirm it. Referring to the last example, the simultaneous use of the resource PC at the phase level is checked by explicitly setting the following goal using the Multiple_Phase_Conflict predicate:

set_goal "$\neg$Multiple_Phase_Conflict (IF_X, ID_C, PC)"

If the goal is proven true, then the conflict freedom is shown, else the resource conflict remains and in order to resolve it, the implementation EBM has to be changed appropriately, e.g. by using an additional buffer.

6 Data Conflicts

Data conflicts (also called *data hazards* [6, 8, 12], *timing hazards* [10], *data dependencies* [5] or *data races* [3]) arise when an instruction depends on the results of a previous instruction. Such data dependencies could lead to faulty computations when the order in which the operands are accessed is changed by the pipeline.

Data conflicts are of three types called, read after write (RAW), write after read (WAR) and write after write (WAW) [6, 8, 12, 5] (also called destination source (DS), source destination (SD) and destination destination (DD) conflicts [10]). Given that an instruction I_j is issued after I_i, a brief description of these conflicts is:

RAW conflict	– I_j *reads* a source before I_i *writes* it
WAR conflict	– I_j *writes* into a destination before I_i *reads* it
WAW conflict	– I_j *writes* into a destination before I_i *writes* it

6.1 Data Conflict Specification

Data conflicts include temporal aspects that are related to the temporal abstractions of our hierarchical model. Therefore, similarly to resource conflicts, the formal specifications of data conflicts are considered at different abstraction levels.

Class Level Conflicts. Considering a full pipeline (see figure 2), the data conflict predicate, i.e. Data_Conflict, should involve the maximal number n_s of instructions that could lead to data conflicts. The predicate Data_Conflict is defined in terms of a multiple data conflict predicate, which includes n_s instructions $I_1 \ldots I_{n_s}$ with corresponding sequential issue times $t_{o1} \ldots t_{on_s}$[7]. The predicate Multiple_Data_Conflict is true whenever any two class instructions conflict on some data. Hence, we define Multiple_Data_Conflict as the disjunction of all possible dual data conflicts (represented by Dual_Data_Conflict) as follows:

$$
\begin{array}{l}
\text{Data\_Conflict} = \\
\quad \vdash_{def} \text{Multiple\_Data\_Conflict}\ (I_1, \ldots, I_{n_s}) := \\
\qquad \exists\, t_{o1} \ldots t_{on_s} : \text{time} . \\
\qquad\quad \bigvee_{\substack{i,j \\ (i,j = 1 \ldots n_s) \\ (i < j)}} \text{Dual\_Data\_Conflict}\ ((I_i, t_{oi}), (I_j, t_{oi} + j - 1))
\end{array}
$$

The predicate Dual_Data_Conflict is true, if there exists a resource of the programming model (class level) for which two class instructions I_i and I_j issued at time points t_{oi} and t_{oj}, respectively, conflict. Further, according to our hierarchical model, the Dual_Data_Conflict is handled hierarchically, first at the stage then at the phase level. Formally, Dual_Data_Conflict is defined in terms of a Stage_Data_Conflict predicate, as follows:

$$
\begin{array}{l}
\vdash_{def} \text{Dual\_Data\_Conflict}\ ((I_i, t_{oi}), (I_j, t_{oj})) := \\
\qquad \exists\, r : \text{CL\_resource} . \\
\qquad\quad \text{Stage\_Data\_Conflict}\ ((I_i, t_{oi}), (I_j, t_{oj}), r)
\end{array}
$$

Stage Level Conflicts. Let I_i be an instruction that is issued into the pipeline at time t_{oi} and *writes* a given resource r at t_{ui} ($t_{oi} \le t_{ui}$). Let I_j be another instruction that is issued at later time t_{oj}, i.e. ($t_{oi} < t_{oj}$) and *reads* the same resource r at t_{uj}. A RAW data conflict occurs when the resource r is read by I_j *before* (and not *after*) this resource is written by the previous sequential instruction I_i (figure 4). Let s_i and s_j be the related pipeline stages in which the resource r is written and read, respectively. Assuming a linear pipeline execution of instructions, i.e. no pipeline freeze or stall happen, the use time points t_{ui} and t_{uj} are equal to ($t_{oi} + \Psi$[8](s_i)) and ($t_{oj} + \Psi(s_j)$), respectively. Hence, the timing condition for the RAW conflict, i.e. ($t_{uj} \le t_{ui}$), is equivalent to $(t_{oj} - t_{oi}) \le (\Psi(s_i) - \Psi(s_j))$.

[7] We assume a linear pipelining of instructions, i.e. no pipeline freeze or stall exist.

[8] The symbol Ψ represents the function Stage_Rank, which computes the ordinal value of a given pipeline stage (cf. section 4).

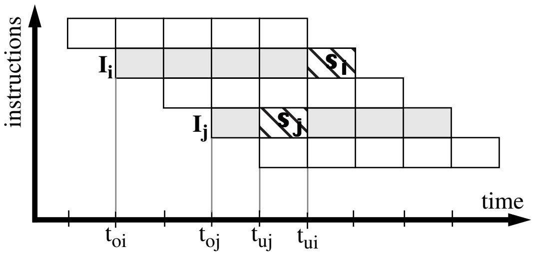

Fig. 4. RAW Data Conflict

Using the function Stage_Rank (represented by the symbol Ψ) and the predicates Stage_Range and Stage_Domain, the formal specification of the stage RAW data conflict is thus given as follows:

$\vdash_{def}$ Stage_RAW_Conflict $((I_i, t_{oi}), (I_j, t_{oj}), r)$:=
$\exists\, s_i\, s_j$: pipeline_stage .
$(0 < (t_{oj} - t_{oi})) \wedge$
$(t_{oj} - t_{oi}) \leq (\Psi\ (s_i) - \Psi\ (s_j))) \wedge$
Stage_Range $(I_i, s_i, r) \wedge$
Stage_Domain (I_j, s_j, r)

Similarly, the WAR and WAW predicates are defined as follows, where the predicate parameters are the same as by the definition of Stage_RAW_Conflict and the semantics of the data conflict is reflected by the order of the Stage_Range and Stage_Domain predicates:

$\vdash_{def}$ Stage_WAR_Conflict (…):=
$\exists\, s_i\, s_j$: pipeline_stage .
$(0 < (t_{oj} - t_{oi})) \wedge$
$(t_{oj} - t_{oi}) \leq (\Psi\ (s_i) - \Psi\ (s_j))) \wedge$
Stage_Domain $(I_i, s_i, r) \wedge$
Stage_Range (I_j, s_j, r)

$\vdash_{def}$ Stage_WAW_Conflict (…):=
$\exists\, s_i\, s_j$: pipeline_stage .
$(0 < (t_{oj} - t_{oi})) \wedge$
$(t_{oj} - t_{oi}) \leq (\Psi\ (s_i) - \Psi\ (s_j))) \wedge$
Stage_Range $(I_i, s_i, r) \wedge$
Stage_Range (I_j, s_j, r)

Phase Level Conflicts. A special case of the data conflict timing condition arises when a resource is simultaneously used by the instructions I_i and I_j, i.e. $t_{ui} = t_{uj}$. In this situation, the data conflict should be examined at the phase level. Let S_i and S_j be any two stage instructions, where the rank of S_i is greater than that of S_j, e.g. S_i = WB_L and S_j = ID_C. According to figure 5, a RAW data conflict at the phase level happens when the resource r is *written* by the stage instruction S_i at a clock phase p_i that occurs *after* clock phase p_j, where it is *read* by S_j, i.e. $(\tau_{ui} \geq \tau_{uj})$. Since instructions at the phase level are executed purely in parallel, they all have the same issue time t_o (figure 5), the timing condition $(\tau_{ui} \geq \tau_{uj})$ is equivalent to $(\tau_o + \xi^9(p_i)) \geq \tau_o + \xi(p_j)) = (\xi(pi) \geq \xi(p_j))$. Using the functions Stage_Rank, Phase_Rank

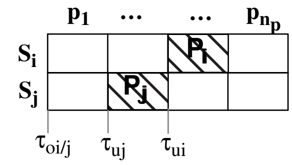

Fig. 5. Phase Level RAW Data Conflict

[9] The symbol ξ represents the function Phase_Rank, which computes the ordinal value of the clock phase (cf. section 4).

and Stage_Type and the predicates Phase_Domain and Phase_Range, the phase level RAW data conflict predicate is formally given as follows:

$$
\begin{array}{l}
\vdash_{def} \text{Phase\_RAW\_Conflict}\ (S_i, S_j, r) := \\
\quad \exists\, p_i\, p_j : \text{clock\_phase}\,. \\
\qquad (\xi\,(p_i) > \xi\,(p_j)) \wedge \\
\qquad (\Psi\,(\vartheta^{10}(S_i)) > (\Psi\,(\vartheta\,(S_j)) \wedge \\
\qquad \text{Phase\_Range}\ (S_i, p_i, r) \wedge \\
\qquad \text{Phase\_Domain}\ (S_j, p_j, r)
\end{array}
$$

In a similar manner, the formal definitions of the phase level WAR and WAW data conflict predicates are given as follows:

$$
\begin{array}{l}
\vdash_{def} \text{Phase\_WAR\_Conflict}\ (S_i, S_j, r) := \\
\quad \exists\, p_i\, p_j : \text{clock\_phase}\,. \\
\qquad (\xi\,(p_i) > \xi\,(p_j)) \wedge \\
\qquad (\Psi\,(\vartheta(S_i)) > (\Psi\,(\vartheta\,(S_j)) \wedge \\
\qquad \text{Phase\_Domain}\ (S_i, p_i, r) \wedge \\
\qquad \text{Phase\_Range}\ (S_j, p_j, r)
\end{array}
$$

$$
\begin{array}{l}
\vdash_{def} \text{Phase\_WAW\_Conflict}\ (S_i, S_j, r) := \\
\quad \exists\, p_i\, p_j : \text{clock\_phase}\,. \\
\qquad (\xi\,(p_i) > \xi\,(p_j)) \wedge \\
\qquad (\Psi\,(\vartheta(S_i)) > (\Psi\,(\vartheta\,(S_j)) \wedge \\
\qquad \text{Phase\_Range}\ (S_i, p_i, r) \wedge \\
\qquad \text{Phase\_Range}\ (S_j, p_j, r)
\end{array}
$$

6.2 Data Conflict Proof

Our ultimate goal in the proof of the non existence of data conflicts relies in showing that none of the data conflicts (RAW, WAR and WAW) occurs. This proof is split into three independent parts each corresponding to one data conflict type. These proofs are similar and in the following we will handle RAW conflicts for illustration purposes.

At the top-most level, the goal to be proven for RAW data conflicts is given in terms of the multiple RAW data conflict predicate as follows:

$$
\begin{array}{l}
\texttt{set\_goal}\ \text{"}\forall\, I_1 \ldots\, I_{n_s} : \text{class\_instruction}\,. \\
\qquad\qquad \neg\text{Multiple\_RAW\_Conflict}\ (I_1, \ldots, I_{n_s})\text{"}
\end{array}
$$

This goal includes a quantification over all possible conflict combinations that could occur between all permutations of n_s instruction within the pipeline. As the case by resource conflicts, the direct proof of this goal yields to a case explosion. Hence, we manage the proof in two steps in that we first prove the non happening of dual conflicts and then conclude the negation of the multiple conflict predicate from the first step. The proof of step 2 is done straightforwardly. The goal for the second step is given as follows:

$$
\begin{array}{l}
\texttt{set\_goal}\ \text{"}\forall\, I_i\, I_j : \text{class\_instruction}\,. \\
\qquad \forall\, t_{oi}\, t_{oj} : \text{time}\,. \\
\qquad\quad \forall\, r : \text{CL\_Resource}\,. \\
\qquad\qquad \neg\text{Stage\_RAW\_Conflict}\ ((I_i, t_{oi}), (I_j, t_{oj}), \mathrm{r})\text{"}
\end{array}
$$

The proof of this goal yields either to *True* or to a number of subgoals, which include the specific resource and class instructions that conflict. For example, a data conflict that occurs between LOAD and ALU-instructions due to the resource register file RF, which is written at the WB-stage by the LOAD-instruction and read at the ID-stage by the

[10] The symbol ϑ represents the function Stage_Type, which computes the logical stage type of a stage instruction (cf. section 4).

ALU-instruction is detected and output as follows, where the number "3" corresponds to the difference $\Psi(s_i) - \Psi(s_j) =$ "Ψ (WB) - Ψ (ID)":

$$(I_i = \text{LOAD}), (I_j = \text{ALU}), (s_i = \text{WB}), (s_j = \text{ID}), (r = \text{RF}), (0 < (t_{oj} - t_{oi})) \vdash \neg((t_{oj} - t_{oi}) \le 3)$$

This result is interpreted as follows: as long as the issue times of the conflicting LOAD and ALU-instructions satisfy the condition "$(t_{oj} - t_{oi}) \le 3$", there exists a data conflict. In order to resolve this conflict, we should neutralize this timing condition. This can be done by considering the following two cases:

1. "$(t_{oj} - t_{oi}) = 3$": the data conflict should be explored at the lower time granularity, i.e. phase level, by setting the following goal:

```
set_goal "¬Phase_RAW_Conflict
              (ClassToStage (WB, LOAD), ClassToStage (ID, ALU), RF)"
```

If the goal is proven correct, no data conflict happens, else either the hardware EBM should be changed, e.g. though inclusion of more clock phases, or one uses the software scheduling technique [6].

2. "$(t_{oj} - t_{oi}) < 3$": The timing information gives an exact reference for the maximum number of pipeline slots or bypassing paths that have to be provided by the software scheduling technique or the implementation EBM, respectively, namely (3-1 = **2**).

Using instruction scheduling, the needed software constraint (assumption) that leads to the proof of the dual data conflict goal, could then be defined as:

$$\vdash_{def} \text{SW\_Constraint} := (I_i = \text{LOAD}) \wedge (I_j = \text{ALU}) \wedge (r = \text{RF}) \wedge (0 < (t_{oj} - t_{oi})) \Rightarrow ((t_{oj} - t_{oi}) > 3)$$

Using the *bypassing* (also called *forwarding*) technique [6], we ensure that the needed data is forwarded as soon as it is computed (end of the EX-stage) to the next instruction (beginning of the EX-stage). This behaviour is implemented in hardware by using some registers and corresponding feedback paths that hold and forward this data, respectively. The existence of these registers and corresponding forward paths can be easily formalized as follows:

$$\vdash_{def} \text{Bypassing} := \exists rb.\ (rb = \text{RF}) \wedge \text{Stage\_Range}\ (I_i, \text{EX}, rb) \wedge \text{Stage\_Domain}\ (I_j, \text{EX}, rb)$$

Assuming this bypassing condition in the dual data conflict goal, the existentially quantified pipeline stage variables s_i and s_j in the definition of Stage_RAW_Conflict (cf. section 6.1) are set to EX and the timing condition is hence reduced to:

$$\ldots, (0 < (t_{oj} - t_{oi})) \vdash \neg((t_{oj} - t_{oi}) \le 0)$$

which is always true.

To summarize, given some specific software constraints in form of instruction scheduling timing conditions and given the implementation EBM, which includes some bypassing paths with appropriate logic, we are able to prove that none of the data conflicts (RAW, WAR and WAW) happens, i.e. formally:

$$\text{SW\_Constraints, EBM} \vdash (\neg\text{RAW\_Data\_Conflict} \wedge \neg\text{WAR\_Data\_Conflict} \wedge \neg\text{WAW\_Data\_Conflict})$$

7 Control Conflicts

Control conflicts (also called *control hazards* [6, 12], *branch hazards* [6], *sequencing hazards* [10] or *branch dependencies* [5]) arise from the pipelining of branches and other instructions that change the program counter *PC*, i.e. interruption of the linear instruction flow.

Control conflicts involve all kinds of interrupts to the linear pipeline flow. This includes the software branches, jumps, traps, etc. and hardware interrupts, stalls, etc. Using this classification, we define *SW control conflicts* and *HW control conflicts* as those conflicts which are caused by a software instruction and a hardware event, respectively. The correctness proof of *Control_conflict* is therefore split into two independent parts as follows:

7.1 SW Control Conflict

Let $A_f\,(I_i,\ t_{oi})$ be the fetch address of an instruction I_i issued at the time point t_{oi} and $A_n(I_i,\ t_{oi})$ be the fetch address of the next instruction (also called next address of I_i). Let *PC* be the program counter. In a pipelined instruction execution, at each clock cycle a new instruction is fetched, i.e. $\forall\ t.\ PC\ (t{+}1) = Inc\ PC\ (t)$[11]. The fetch and next addresses of an instruction I_i are equal to $A_f\ (I_i,\ t_{oi}) = PC\ (t_{oi})$ and $A_n\ (I_i,\ t_{oi}) = PC\ (t_{oi}{+}1)$, respectively. During a sequential program execution, we have $A_n\ (I_i,\ t_{oi}) = A_f(I_i{+}1,t_{oi}{+}1)$ and we obtain the following:

$$\forall\ t_{oi}\ .\ A_n\ (I_i,\ t_{oi}) = A_f\ (I_i{+}1,\ t_{oi}{+}1) = PC\ (t_{oi}{+}1) = Inc\ PC\ (t_{oi})$$

If I_i is a control instruction, then the next address is its target address, i.e. $A_n\ (I_i,\ t_{oi}) = A_t\,(I_i,\ t_{oi})$. Due to the sequential execution of a single instruction, the target instruction can only be fetched after the instruction I_i is fetched, decoded and the target address has been calculated. Since all this cannot happen in one clock cycle, the target address $A_t\,(I_i,t_{oi}) = PC\ (t_{oi}{+}n)$, where $n > 1$. Hence, the next address is not equal to the target address, i.e.:

$$A_n\ (I_i,\ t_{oi}) = Inc\ PC\ (t_{oi}) \neq A_t\ (I_i,\ t_{oi}) = PC\ (t_{oi}{+}n)$$

and the wrong instruction is fetched next.

A closer look at this situation shows that a software control conflict occurs when an instruction attempts to read the resource *PC* that is not yet updated (written) by a previous instruction. This complies with the definition of RAW data conflict in *PC* [8] and the software control conflict could be defined as follows:

$\vdash_{def}$ SW_Control_Conflict := Stage_RAW_Conflict ((CONTROL, t_{oi}), (I_j, t_{oj}), PC)

The conflict freedom proof is therefore only a special case of the data conflict proofs and the goal to be proven is set as follows:

```
set_goal "∀ I_j : class_instruction.
           ∀ t_oi t_oj : time.
             ¬Stage_RAW_Conflict ((CONTROL, t_oi), (I_j, t_oj), PC) "
```

and we yield for the DLX processor example 5 subgoals of the following form:

$$(I_j = \text{ALU}),\ (s_i = \text{ID}),\ (s_j = \text{IF}),\ (0 < (t_{oj} - t_{oi}))\ \vdash \neg((t_{oj} - t_{oi}) \leq 1)$$

[11] *Inc* is an increment function which is independent of the increment value, e.g. PC+4, PC+8, etc.

For conflict resolution no bypassing is possible, since the calculation of the target address cannot be done earlier. The commonly used technique is software scheduling [6]. In the DLX RISC processor, we just need one delay slot ($(t_{oj} - t_{oi}) \leq 1$) to ensure that control instructions are executed correctly. The software constraint needed in this case is defined as follows:

$$\vdash_{def} \text{SW\_Constraint} := (0 < (t_{oj} - t_{oi})) \Rightarrow ((t_{oj} - t_{oi}) > 1)$$

7.2 HW Control Conflicts

In contrast to the previous conflict kind, this kind of conflict cannot be handled in general since the interrupt behaviour and conflict resolution are fully hardware implementation dependent. Different RISC processors handle interrupts in different ways, i.e where the forced branch is to be inserted, how to restart exactly the interrupted program, etc.

In general, one has to formally specify the interrupt behaviour in the form of a predicate, whose implication has to be proven correct from the implementation EBM, i.e.:

$$\vdash \text{EBM} \Rightarrow \text{INTR\_SPEC}$$

The predicate INTR_SPEC[12] describes the behaviour of the implemented hardware interrupt and ensures that no conflict occurs when the linear pipeline flow is interrupted, e.g. proper saving and recovery of the processor state before and after the interrupt handling. Formally:

$$\vdash \text{INTR\_SPEC} \Rightarrow \neg\text{HW\_Control\_Conflict}$$

and consequently we yield:

$$\text{EBM} \vdash \neg\text{HW\_Control\_Conflict}$$

Summary

Our ultimate goal in proving the non-existence of pipeline conflicts is concluded from the theorems yielded in sections 5.2, 6.2, 7.1 and 7.2, i.e.:

$$\text{SW\_Constraints, EBM} \vdash \begin{array}{l}(\neg\text{Res\_Conflict} \wedge \\ \neg\text{Data\_Conflict} \wedge \\ \neg\text{Control\_Conflict})\end{array}$$

8 Experimental Results

The methodology presented so far, has been validated by using the DLX processor as an example. This processor contains a 5 stage pipeline with a 2 phased clock, and its architecture includes 51 basic instructions (integer, logic, load/store and control). All these instructions are grouped into 5 classes according to which the stage and phase instructions are defined (refer to table 1). The EBM for the DLX core has been implemented in a commercial VLSI design environment using a 1.0 µm CMOS technology. The design has approximately 150 000 transistors which occupies a silicon area of about 60.34 mm^2, it has 172 I/O pads and currently runs at a clock rate of 12.5 MHz.

[12] A formal specification of INTR_SPEC for DLX is beyond the scope of this paper and is not presented here.

The formalization of the specifications for DLX have been done within the HOL verification system [4]. The implementation of the interpreter model specifications and the proof of step (2) is reported in [15]. The overall specification for the DLX core is about 4760 lines long and the proof of step (2) took about 457.66 secs on a SPARC10, with a 96 MB main memory.

In proving step (1), the specification and tactics for each kind of conflict have been implemented in HOL90.6. These specifications and tactics are general enough to be applicable to a wide range of other RISC processors. The hierarchical structuring of the proofs resulted in parameterized tactics that are used for more than one kind of conflict. All proofs have been mainly done using 4 proof tactics — MULTIPLE_CONF_TAC, RES_CONF_TAC, DATA_CONF_TAC and CONTROL_CONF_TAC — with a total proof script text of 850 lines. The overall definition and specification text of step (1) is about 2690 lines. The run times and the number of created inferences, on a SPARC10 with a 96 MB main memory, for the theorem generation of the Used and Range/Domain predicates, for the pipeline verification and for the semantical correctness proofs, respectively, are given for the DLX processor example in the following table:

Table 2. Verification Results for the DLX processor

Verification Goal	**Time (in sec)**	**# Inferences**	**Comments**
Stage_Used Gen.	247.27	768226	208 theorems generated
Phase_Used Gen.	1718.03	3756157	422 theorems generated
Stage Ran. & Dom. Gen.	324.44	534283	250 theorems generated
Phase Ran. & Dom. Gen.	272.27	267926	260 theorems generated
Stage Dual Res. Conf.	1691.64	3330880	0 conflicts
Dual $\Rightarrow$ Multiple (Res.)	1.30	2434	-
Stage RAW Data Conf.	570.64	1669029	15 conflict cases (3 slots) by RF and 5 conflict cases (1 slot) by PC
(using Bypassing)	1.89	5438	0 conflicts
(using SW-scheduling)	458.96	1344577	0 conflicts
Stage WAR Data Conf.	572.23	1669069	0 conflicts
Stage WAW Data Conf.	572.50	1669047	0 conflicts
Dual $\Rightarrow$ Multiple (Data)	0.27	389	-
SW Control Conflicts	32.44	91403	5 conflict cases (1 slot)
Σ Pipeline Correctness	6464.18	15108496	–

9 Conclusions and Future Work

In this paper, we have shown the formalization and the verification of the different kinds of conflicts that occur in the instruction pipelines of RISC cores. The employment of the hierarchical RISC interpreter model and in particular the exploitation of the class level, empowers us to derive compact specifications of the conflicts. We have implemented constructive proofs for these conflicts and hence the designer gets invaluable feedback for resolving these conflicts, either by making appropriate modifications to the hardware or by generating the required software constraints.

The interpreter model, the formal specifications and the proof techniques are generic and hence applicable to any RISC core. Furthermore, the RISC interpreter model closely reflects the RISC design hierarchy thus the specifications in higher-order logic are easy to derive. Given such specifications and an implementation, the proof process has been automated by using parametrizable tactics. These tactics are independent of the underlying implementation and can be used for a large number of RISC cores. The entire methodology has been validated by using the DLX core.

The feasibility of formal verification techniques, when applied cleverly to specific classes of circuits is illustrated by the runtimes shown in table 2. In our future work, we shall extend the layer of the core to include pipelined functional units, floating point processor, etc.

References

1. Anceau, F.: The Architecture of Microprocessors; Addison-Wesley Publishing Company, 1986.
2. Buckow, O.: Formale Spezifikation und (Teil-) Verifikation eines SPARC-kompatiblen Prozessors mit LAMBDA; Diplomarbeit, Universität-Gesamthochschule Paderborn, Fachbereich Mathematik-Informatik, Oktober 1992.
3. Furber, S.: VLSI RISC Architecture and Organization; Electrical Engineering and Electronics, Dekker, New York, 1989.
4. Gordon, M.; Melham, T.: Introduction to HOL: A Theorem Proving Environment for Higher Order Logic; Cambridge, University Press, 1993.
5. Van De Goor, A.: Computer Architecture and Design; Addison-Wesley, 1989.
6. Hennessy, J.; Patterson, D.: Computer Architecture: A Quantitative Approach; Morgan Kaufmann Publishers, Inc., San Mateo, California, 1990.
7. Joyce, J.: Multi-Level Verification of Microprocessor-Based Systems; Ph.D. Thesis, Computer Laboratory, Cambridge University, December 1989.
8. Kogge, P.: The Architecture of Pipelined Computers; McGraw-Hill, 1981.
9. Kumar, R.; Schneider, K.; Kropf, Th.: Structuring and Automating Hardware Proofs in a Higher-Order Theorem-Proving Environment; Journal of Formal Methods in System Design, Vol.2, No. 2, 1993, pp. 165-230.
10. Milutinovic, V.: High Level Language Computer Architecture; Computer Science Press, Inc., 1989.
11. Srivas, M.; Bickford, M.: Formal Verification of a Pipelined Microprocessor; IEEE Software, September 1990, pp. 52-64.
12. Stone, H.: High-Performance Computer Architecture; Addison-Wesley Publishing Company, 1990.
13. Tahar, S.; Kumar, R.: A Formalization of a Hierarchical Model for RISC Processors; In: Spies, P. (Ed.), Proc. European Informatics Congress Computing Systems Architecture (Euro-ARCH93), Munich, October 1993, Informatik Aktuell, Springer Verlag, pp. 591-602.
14. Tahar, S.; Kumar, R.: Towards a Methodology for the Formal Hierarchical Verification of RISC Processors; Proc. IEEE International Conference on Computer Design (ICCD93), Cambridge, Massachusetts, October 1993, pp. 58-62.
15. Tahar, S.; Kumar, R.: Implementing a Methodology for Formally Verifying RISC Processors in HOL; Proc. International Meeting on Higher Order Logic Theorem Proving and its Applications (HUG93), Vancouver, Canada, August 1993, pp. 283-296.
16. Tahar, S.; Kumar, R.: Formal Verification of Pipeline Conflicts in RISC Processors; to appear in Proc. European Design Automation Conference (EURO-DAC94), Grenoble, France, September 1994.
17. Windley, P.: The Formal Verification of Generic Interpreters; Ph.D. Thesis, Division of Computer Science, University of California, Davis, July 1990.