

# Towards Provably-Correct Hardware Compilation Tools Based on Pass Separation Techniques

Steve McKeever and Wayne Luk

Department of Computing, Imperial College, 180 Queen's Gate, London, UK  
{swm2, wl}@doc.ic.ac.uk

**Abstract.** This paper presents a framework for verifying compilation tools based on parametrised hardware libraries expressed in Pebble, a simple declarative language. An approach based on pass separation techniques is described for specifying and verifying Pebble abstraction mechanisms, such as the loop statement. We show how this approach can be used to verify the correctness of the flattening procedure in the Pebble compiler, which also results in a more efficient implementation than a non-verified version. The approach is useful for guiding compiler implementations for Pebble and related languages such as VHDL; it may also form the basis for automating the generation of provably-correct tools for hardware development.

## 1 Introduction

Advance in integrated circuit technology leads to an increasing emphasis on building designs from hardware libraries. A single parametrised library can be used to generate many implementations supporting multiple architectures, variable bit widths and trade-offs in speed and size. Such libraries enable effective hardware utilisation by exploiting technology-specific features whenever desirable, allowing designs with optimal performance and resource usage while minimising the need for knowing low-level details.

This paper describes an approach for developing provably-correct compilation tools for the Pebble language [12], which has been used to produce hardware libraries in VHDL, an industry standard language. While it is desirable to have hardware libraries in industry standard languages, there are, however, two major difficulties with developing VHDL libraries. First, VHDL is a versatile but complex language, and it takes much effort to write good parametrised code and to check its behaviour by simulation or other means – even when the subset used for realistic hardware libraries is small [13]. Second, most vendors have their own VHDL dialect; for instance not all VHDL tools support multi-dimensional vectors. It is unattractive to develop and maintain the same set of library elements in various vendor-specific dialects.

Our aim for Pebble is to enable application builders and library developers to work at a higher level of abstraction than that provided by VHDL, while ensuring that the resulting libraries are as flexible and efficient as those produced by hand. Our Pebble compiler targets various description formats, including parametrised and flattened VHDL and EDIF. It enables designers to compose and instantiate library elements, and it has been used to develop many designs for applications such as speech processing, data compression, and special video and graphics effects in augmented reality [14].

An important component of the Pebble compiler is the flattening procedure, which produces flattened descriptions from hierarchical descriptions. Flattened descriptions are required by many tools, such as place-and-route programs for design implementation, and model checkers for design verification [18]. Interestingly, our proof of the flattening procedure not only offers users greater confidence in its correctness, it also leads to a more efficient implementation. Although this paper focuses on a specific proof, recent work indicates that similar techniques can be applied to verify other abstraction mechanisms, such as polymorphic variables and higher-order functions.

Pass separation [11] provides a framework for verifying abstraction mechanisms for generic descriptions, such as hierarchical blocks and `GENERATE-FOR` loops in Pebble, with respect to a Structural Operational Semantics [16]. For instance, the key to our proof is an environment invariant – Equation (1) – inspired by pass separation (Sect. 3.2). Such proofs are rare, but we feel that they are well-suited to verifying development tools for domain-specific languages containing multiple evaluation phases.


While much has been published on formal methods and tools for hardware design, it appears that most researchers focus on tools for producing correct designs [2],[18] rather than on the correctness of the tools themselves. Some researchers study embedding of synthesis algorithms in a theorem prover [4], or correctness condition generators for designs generated by synthesis tools [15]. Our work is complementary to their efforts, and is in a similar spirit to research on verifying compiler correctness for imperative descriptions for hardware [6],[7] and software [8] implementations.

## 2 Overview of Pebble

Pebble can be regarded as a much simplified variant of Structural VHDL. It provides a means of representing block diagrams hierarchically and parametrically [12]. A Pebble program is a block, defined by its name, parameters, interfaces, local definitions, and its body. The block interfaces are given by two lists, usually interpreted as the inputs and outputs. An input or an output can be of type `WIRE`, or it can be a multi-dimensional vector of wires. A wire can carry integer, boolean or other primitive data values.

A primitive block has an empty body; a composite block has a body containing the instantiation of composite or primitive blocks in any order. Blocks connected to each other share the same wire in the interface instantiation. For hardware designs, the primitive blocks can be bit-level logic gates and registers, or they can, like an adder, process word-level data such as integers or fixed-point numbers; the primitives depend on the availability of corresponding components in the domain targeted by the Pebble compiler. The `GENERATE-IF` statement enables conditional compilation and recursive definition, while the `GENERATE-FOR` statement allows the concise description of regular circuits.

Pebble has a simple, block-structured syntax. As an example, Fig. 2 describes the multiplexor array in Fig. 1, provided that the size parameter  $n$  is 4. In more complex descriptions, the parameters in a Pebble program can include the number of pipeline stages or the pitch between neighbouring interface connections [12]. Different network structures, such as tree- or butterfly-shaped circuits, can be described parametrically by indexing the components and wires.



**Fig. 1.** An array of multiplexers described by the Pebble program in Fig. 2.

```

BLOCK muxarray (n:GENERIC)
    [c:WIRE, x,y:VECTOR (n-1..0) OF WIRE]
    [z:VECTOR (n-1..0) OF WIRE]
    VAR i
    BEGIN
        GENERATE FOR i = 0..(n-1)
        BEGIN
            mux [c,x(i),y(i)] [z(i)]
        END
    END;

```

**Fig. 2.** A description of an array of multiplexers (Fig. 1) in Pebble. The external input  $c$  is used to provide a common control input for each multiplexer.

The semantics of Pebble depends on the behaviour of the primitive blocks and their composition in the target technology. Currently a synchronous circuit model is used in our tools (Sect. 3), and special control components for modelling run-time reconfiguration are also supported [12]. However, other models can be used if desired. Indeed Pebble can model any block-structured systems, not just electronic circuits.

Advanced features of Pebble include support for annotations and for modules. Such features improve design efficiency and reusability, and facilitate interface to components in other languages, including behavioural descriptions. Discussions about these features are beyond the scope of this paper.

### 3 Program Staging and Pass Separation

This section introduces a framework in which abstraction mechanisms for Pebble can be specified and verified. Our approach consists of three steps. The first is to provide a semantics for a flattened version of Pebble. The second is to characterise an abstraction mechanism in two ways: (a) specify how designs exploiting the abstraction mechanism can be transformed into flattened Pebble, and (b) provide the semantics of the abstraction mechanism directly. The third is to show that the semantics of a design produced by (a) is consistent with (b).

In the following, we specify and verify two Pebble abstraction mechanisms: hierarchical blocks, and `GENERATE-FOR` loops. The insight is to recognise that pass separation provides a framework for the above approach, so that the correctness of the two abstraction mechanisms can be demonstrated with respect to a Structural Operational Semantics [16] for Pebble. Only elementary understanding of such semantics is required to follow our work. We shall first present an overview of pass separation, and then explain how it can be used in verifying Pebble abstraction mechanisms. The key to our proof is an environment invariant – Equation (1) – inspired by pass separation.


Consider a repeated computation, part of whose input context remains invariant across all repetitions. Program staging is a technique which improves performance by separating the computation into two phases. We follow this approach to separate the task of interpreting a Pebble program on a variety of inputs: an early phase flattens the parametrised description into a collection of primitive gate calls, and a late phase completes the task given the varying inputs.

Two popular methods for separating a computation into stages are partial evaluation and pass separation [11]. We have used both methods in our study of Pebble and in tool development; in the following we shall focus on pass separation as a means to study abstraction mechanisms for Pebble.

Pass separation constructs, from a program  $p$ , two programs  $p_1, p_2$  such that:

$$\llbracket p \rrbracket (x, y) = \llbracket p_2 \rrbracket (\llbracket p_1 \rrbracket x, y)$$

for all  $x$  and  $y$ , where  $\llbracket p \rrbracket$  is the function mapping the program  $p$  to its meaning. The equation indicates that  $\llbracket p \rrbracket$  can be split into two stages: computing  $v = \llbracket p_1 \rrbracket x$  and  $\llbracket p_2 \rrbracket (v, y)$ . The intention in performing pass separation is to “move” as many computations from  $p$  to  $p_1$  as possible, given only input  $x$ . In our framework, let  $p$  be the semantics ( $\mathcal{PS}$ ) of Pebble,  $x$  be a parametrised circuit description  $C$ , and  $y$  be some input data. Then  $p_1$  corresponds to the abstraction mechanism, which in this case can be described using a flattening procedure ( $\mathcal{FP}$ ). Similarly  $p_2$  corresponds to the semantics ( $\mathcal{FS}$ ) of the flattened description on the data as shown in Fig. 3.



**Fig. 3.** Commuting diagram describing the pass separation equation.

We shall restrict our attention to a subset of Pebble which does not include vectors or `GENERATE-IF` statements. We begin by presenting the semantics of Flattened Pebble before adding the necessary structure to create Hierarchical Pebble. We then display a

procedure for instantiating the generics. We show that for a set of input values, a flattened description produces the same results as those from the hierarchical description.

### 3.1 Flattened Pebble

In its most basic form, a circuit consists of a collection of primitive block calls mapping input wires to output wires. Intermediate wires link these primitive block calls together. A circuit description is enclosed within a `main` block and primitive block identifiers are denoted by  $id_{\mathcal{P}}$ , as shown in the following syntax:

```

circuit ::= BLOCK main
           [idin1 : typein1, . . . , idinn : typeinn]
           [idout1 : typeout1, . . . , idoutm : typeoutm]
           dec1 ; . . . ; decj
           BEGIN stmts END
type    ::= WIRE
dec     ::= VAR id : type
stmts   ::= stmt1 ; . . . ; stmtk
stmt    ::= id $\mathcal{P}$  [id1, . . . , idn] [id1, . . . , idm]

```

Note that the language is applicative, as each wire is given an attribute only once. Input wires are defined and given the appropriate values; the purpose of the semantics is to find suitable definitions for the output wires. The semantic domain for wires is parametrised by the metavariable  $a$ , so that primitive objects can be of any type; it allows us to deal with both bit-level descriptions and word-level descriptions.

**data** wire  $a = \text{Undefined} \mid \text{Defined } a$

To deal with sequential circuits, the datatype  $a$  can be lifted to the stream domain [9].

The Structural Operational Semantics rules for Flattened Pebble, defined by  $\rightarrow$  transitions [16], are given in Fig. 4. A local environment  $\rho$  maps identifiers to their wires. Primitive logic operators that map boolean pairs to booleans, such as XOR, are held in an environment labeled  $\delta$ . Such operators can only be applied to wires that are defined.

Two rules provide the meaning of primitive gate calls. If one or more inputs are Undefined, then the statement is returned unevaluated, as the gate call cannot be completed. The second rule applies the primitive function to the gate's parameters. Statements can be evaluated in any order; those that complete update the local environment  $\rho$ . When all statements have been reduced, the final environment is returned.

The output and intermediate wires of the `main` block are initially declared as Undefined. The block's statements are evaluated to calculate the final environment  $\rho'$ , from which the output wires are extracted.

### 3.2 Hierarchical Pebble

Parametrised designs require the addition of a separate parameter list for generic values. Blocks other than `main` can receive values that define the bounds of loops or that can

$$\begin{array}{c}
\frac{\exists j \cdot 1 \leq j \leq n \wedge (\rho \text{ id}_j) = \text{Undefined}}{\delta \vdash \langle \text{id}_{\mathcal{P}} [id_1, \dots, id_n] [id'_1, \dots, id'_m], \rho \rangle \rightarrow_{stmt} \langle \text{id}_{\mathcal{P}} [id_1, \dots, id_n] [id'_1, \dots, id'_m], \rho \rangle} \\
\frac{(\rho \text{ id}_1) = \text{Defined } v_1 \wedge \dots \wedge (\rho \text{ id}_n) = \text{Defined } v_n \quad (\delta \text{ id}_{\mathcal{P}}) (v_1, \dots, v_n) = (v'_1, \dots, v'_m)}{\delta \vdash \langle \text{id}_{\mathcal{P}} [id_1, \dots, id_n] [id'_1, \dots, id'_m], \rho \rangle \rightarrow_{stmt} \rho \oplus \{id'_1 \mapsto \text{Defined } v'_1, \dots, id'_m \mapsto \text{Defined } v'_m\}} \\
\frac{\delta \vdash \langle \text{stmt}_i, \rho \rangle \rightarrow_{stmt} \rho'}{\delta \vdash \langle \text{stmt}_1; \dots; \text{stmt}_{i-1}; \text{stmt}_i; \text{stmt}_{i+1}; \dots; \text{stmt}_k, \rho \rangle \rightarrow_{stmts} \langle \text{stmt}_1; \dots; \text{stmt}_{i-1}; \text{stmt}_{i+1}; \dots; \text{stmt}_k, \rho' \rangle} \\
\frac{\delta \vdash \langle \text{stmt}_i, \rho \rangle \rightarrow_{stmt} \langle \text{stmt}_i, \rho \rangle}{\delta \vdash \langle \text{stmt}_1; \dots; \text{stmt}_{i-1}; \text{stmt}_i; \text{stmt}_{i+1}; \dots; \text{stmt}_k, \rho \rangle \rightarrow_{stmts} \langle \text{stmt}_1; \dots; \text{stmt}_{i-1}; \text{stmt}_i; \text{stmt}_{i+1}; \dots; \text{stmt}_k, \rho \rangle} \\
\delta \vdash \langle [], \rho \rangle \rightarrow_{stmts} \rho \\
\rho_1 = \{id_{in_1} \mapsto \text{Defined } v_1, \dots, id_{in_n} \mapsto \text{Defined } v_n\} \\
\rho_2 = \{id_{out_1} \mapsto \text{Undefined}, \dots, id_{out_m} \mapsto \text{Undefined}\} \\
\rho_3 = \{id_1 \mapsto \text{Undefined}, \dots, id_j \mapsto \text{Undefined}\} \\
\rho = \rho_1 \oplus \rho_2 \oplus \rho_3 \\
\delta \vdash \langle \text{stmts}, \rho \rangle \rightarrow_{stmts} \rho' \\
\delta \vdash \left\langle \begin{array}{l} \text{BLOCK main} \\ [id_{in_1} : \text{WIRE}, \dots, id_{in_n} : \text{WIRE}] \\ [id_{out_1} : \text{WIRE}, \dots, id_{out_m} : \text{WIRE}] \\ \text{VAR } id_1 : \text{WIRE}; \dots \text{ VAR } id_j : \text{WIRE} \\ \text{BEGIN stmts END} \\ \rightarrow_{main} [(\rho' id_{out_1}), \dots, (\rho' id_{out_m})] \end{array} \right\rangle, [v_1, \dots, v_n] \rangle
\end{array}$$

**Fig. 4.** Semantics of Flattened Pebble, based on  $\rightarrow$  rules for *main*, *stmts* and *stmt*.

be passed to subsequent gate calls as defined in the syntax for Hierarchical Pebble:

```

circuit ::= main ; block1 ; ... ; blocki
main   ::= BLOCK main
          [idin1 : typein1, ..., idinn : typeinn]
          [idout1 : typeout1, ..., idoutm : typeoutm]
          dec1 ; ... ; decj
          BEGIN stmts END
block  ::= BLOCK id (idgen1, ..., idgenq)
          [idin1 : typein1, ..., idinn : typeinn]
          [idout1 : typeout1, ..., idoutm : typeoutm]
          dec1 ; ... ; decj
          BEGIN stmts END

```

```

type ::= WIRE
dec  ::= VAR id : type | VAR id : NUM
stmts ::= stmt1; ... ; stmtk
stmt  ::= idP [id1, ..., idn] [id1, ..., idm]
        | id (exp1, ..., expq) [id1, ..., idn] [id1, ..., idm]
        | GENERATE FOR id = exp1 .. exp2
          BEGIN stmts END
exp   ::= id | n | exp1 bop exp2 | uop exp

```

The semantic rules for Hierarchical Pebble statements, defined by  $\Rightarrow$  transitions, are given in Fig. 5. Two new environments are introduced:  $\Gamma$  maps block names to their bodies, while  $\sigma$  maps generic variables and loop indices to their values. Arithmetic expressions are evaluated by the valuation function  $\mathcal{E}$  in an appropriate value environment. The rules for primitive gate calls and statement lists remain essentially unchanged except for the additional environments.

$$\begin{array}{c}
\frac{\mathcal{E}_\sigma[\text{exp}_1] > \mathcal{E}_\sigma[\text{exp}_2]}{\Gamma, \delta, \sigma \vdash \left\langle \left( \begin{array}{l} \text{GENERATE FOR } id_{index} = \text{exp}_1 .. \text{exp}_2 \\ \text{BEGIN } stmts \text{ END} \end{array} \right), \rho \right\rangle \Rightarrow_{stmt} \rho} \\
\\
\frac{\mathcal{E}_\sigma[\text{exp}_1] \leq \mathcal{E}_\sigma[\text{exp}_2]}{\Gamma, \delta, \sigma \oplus \{id \mapsto \mathcal{E}_\sigma[\text{exp}_1]\} \vdash \langle stmts, \rho \rangle \Rightarrow_{stmts} \rho'} \\
\Gamma, \delta, \sigma \vdash \left\langle \left( \begin{array}{l} \text{GENERATE FOR } id_{index} = (\text{exp}_1+1) .. \text{exp}_2 \\ \text{BEGIN } stmts \text{ END} \end{array} \right), \rho' \right\rangle \Rightarrow_{stmt} \rho'' \\
\hline
\Gamma, \delta, \sigma \vdash \left\langle \left( \begin{array}{l} \text{GENERATE FOR } id_{index} = \text{exp}_1 .. \text{exp}_2 \\ \text{BEGIN } stmts \text{ END} \end{array} \right), \rho \right\rangle \Rightarrow_{stmt} \rho'' \\
\\
\vdots \\
(\Gamma id) = \left( \begin{array}{l} \text{BLOCK } id (id_{gen_1}, \dots, id_{gen_q}) [id_{in_1} : \text{WIRE}, \dots, id_{in_n} : \text{WIRE}] \\ \quad [id_{out_1} : \text{WIRE}, \dots, id_{out_m} : \text{WIRE}] \\ \text{VAR } id_{index_1} : \text{NUM}; \dots \text{VAR } id_{index_j} : \text{NUM}; \\ \text{VAR } id_{local_1} : \text{WIRE}; \dots \text{VAR } id_{local_k} : \text{WIRE} \\ \text{BEGIN } stmts \text{ END} \\ \sigma_1 = \{id_{gen_1} \mapsto \mathcal{E}_\sigma[e_1], \dots, id_{gen_q} \mapsto \mathcal{E}_\sigma[e_q]\} \\ \rho_1 = \{id_{in_1} \mapsto (\rho id_1), \dots, id_{in_n} \mapsto (\rho id_n)\} \\ \rho_2 = \{id_{out_1} \mapsto \text{Undefined}, \dots, id_{out_m} \mapsto \text{Undefined}\} \\ \rho_3 = \{id_{local_1} \mapsto \text{Undefined}, \dots, id_{local_k} \mapsto \text{Undefined}\} \\ \rho' = \rho_1 \oplus \rho_2 \oplus \rho_3 \end{array} \right) \\
\Gamma, \delta, \sigma_1 \vdash \langle stmts, \rho' \rangle \Rightarrow_{stmts} \rho'' \\
\hline
\Gamma, \delta, \sigma \vdash \langle id (e_1, \dots, e_q) [id_1, \dots, id_n] [id'_1, \dots, id'_m], \rho \rangle \\
\Rightarrow_{stmt} \rho \oplus \{id'_1 \mapsto (\rho'' id_{out_1}), \dots, id'_m \mapsto (\rho'' id_{out_m})\}
\end{array}$$

**Fig. 5.** Semantics of Hierarchical Pebble. The  $\Rightarrow$  rules dealing with primitive statements are similar to the corresponding  $\rightarrow$  rules and are not shown.

Two rules are required for loops. The first rule deals with the situation when the loop has terminated. A loop has terminated when the expression representing the lower bound denotes a value that is greater than the upper bound. In such cases the current  $\rho$  environment is returned unchanged. Alternatively, the loop's statements are evaluated in a value environment where the loop index is bound to the lower bound value. Once completed, the loop is re-evaluated with the updated  $\rho'$  and a lower bound expression that has been incremented by one. The rule for enacting new block calls retrieves the block definition from the environment  $\Gamma$ ; it creates a value environment  $\sigma_1$  by mapping the generic variable names to their values calculated in the outer block's value environment  $\sigma$ ; it creates an initial wire environment  $\rho'$  by mapping the input variable names to their wire values extracted from  $\rho$  and coalesced with Undefined bindings for output and intermediate variable names; and it evaluates the called block's statements to create a final environment  $\rho''$ , from which the output wires are extracted. The rule for the main block creates the initial environment  $\rho$  in much the same manner as with Flattened Pebble descriptions, and is shown in Fig. 6. The blocks statements are evaluated in an initial empty value environment.

A Hierarchical Pebble description can be flattened by unfolding both the generic variables and the GENERATE-FOR loops. The block environment  $\Gamma$  and the local variable environment  $\sigma$  support the abstraction mechanisms, and do not affect the underlying evaluation mechanism. Hence we can instantiate generic variables prior to the application of input wires, enabling block definitions to be flattened and incorporated into the main block. Flattened Pebble, itself a subset of Pebble, is used as the output language of this flattening process to facilitate its proof. To avoid parameters and local wire names overwriting each other when instantiating block calls, we rename all such variables beforehand using the function  $\alpha :: \text{block} \rightarrow \text{block}$  (Fig. 7).

To model the static behaviour of the wire environment  $\rho$  in hierarchical descriptions with local bindings, we introduce a local environment  $\mu$  which behaves like a symbol table mapping local parameter names to their original definitions, be they inputs to the circuit or local variable definitions. This leads to the invariant equation below, where  $\rho$  is the environment for modelling wire values in a hierarchical description, while  $\rho_d$  is a dynamic environment for flattened descriptions:

$$\forall id \cdot \rho \text{ id} = \rho_d (\mu \text{ id}) \quad (1)$$

This invariant equation will be used extensively in the correctness proof of the flattening procedure for Hierarchical Pebble in Sect. 4. The flattening procedure itself, defined by  $\Downarrow$ , is given in Fig. 7. Flattening a single statement will result in a pair of lists consisting of primitive gate calls and intermediate wire declarations. The statement list represents those primitive calls required to implement the statement derived from unfolding subsequent parametrised blocks, while the declarations are for the local wire definitions belonging to each unfolded block. The intention is to create the flattened main block from these two lists.

Primitive calls are simply returned with their parameter lists updated with their original variable definitions held in  $\mu$ . Parametrised gate calls create a new instance of the retrieved block using the function  $\alpha$ . Generic variables are bound to their values in  $\sigma_1$ . A static environment  $\mu'$  is created by mapping the parameter names to their original



$$\begin{array}{c}
\rho_1 = \{id_{in_1} \mapsto \text{Defined } v_1, \dots, id_{in_n} \mapsto \text{Defined } v_n\} \\
\rho_2 = \{id_{out_1} \mapsto \text{Undefined}, \dots, id_{out_m} \mapsto \text{Undefined}\} \\
\rho_3 = \{id_1 \mapsto \text{Undefined}, \dots, id_j \mapsto \text{Undefined}\} \\
\Gamma, \delta, \{\} \vdash \langle stmts, \rho_1 \oplus \rho_2 \oplus \rho_3 \rangle \Rightarrow_{stmts} \rho' \\
\hline
\Gamma, \delta \vdash \left\langle \begin{array}{l} \text{BLOCK main} \\ [id_{in_1} : \text{WIRE}, \dots, id_{in_n} : \text{WIRE}] \\ [id_{out_1} : \text{WIRE}, \dots, id_{out_m} : \text{WIRE}] \\ \text{VAR } id_1 : \text{WIRE}; \dots \text{ VAR } id_j : \text{WIRE} \\ \text{BEGIN } stmts \text{ END} \end{array} \right\rangle, [v_1, \dots, v_n] \\
\Rightarrow_{main} [(\rho' id_{out_1}), \dots, (\rho' id_{out_m})]
\end{array}$$

**Fig. 6.** Semantics of the main block in Pebble.

$$\begin{array}{c}
\Gamma, \mu, \sigma \vdash \langle id_{\mathcal{P}} [id_1, \dots, id_n] [id'_1, \dots, id'_m] \rangle \\
\Downarrow_{stmt} ([id_{\mathcal{P}} [\mu id_1, \dots, \mu id_n] [\mu id'_1, \dots, \mu id'_m]], []) \\
(\alpha(\Gamma id)) = \left( \begin{array}{l} \text{BLOCK } id (id_{gen_1}, \dots, id_{gen_q}) [id_{in_1} : \text{WIRE}, \dots, id_{in_n} : \text{WIRE}] \\ [id_{out_1} : \text{WIRE}, \dots, id_{out_m} : \text{WIRE}] \\ \text{VAR } id_{index_1} : \text{NUM}; \dots \text{ VAR } id_{index_j} : \text{NUM} \\ \text{VAR } id_{local_1} : \text{WIRE}; \dots \text{ VAR } id_{local_k} : \text{WIRE} \\ \text{BEGIN } stmts \text{ END} \\ \mu_1 = \{id_{in_1} \mapsto (\mu id_1), \dots, id_{in_n} \mapsto (\mu id_n)\} \\ \mu_2 = \{id_{out_1} \mapsto (\mu id'_1), \dots, id_{out_m} \mapsto (\mu id'_m)\} \\ \mu_3 = \{id_{local_1} \mapsto id_{local_1}, \dots, id_{local_k} \mapsto id_{local_k}\} \\ \sigma_1 = \{id_{gen_1} \mapsto \mathcal{E}_{\sigma}[e_1], \dots, id_{gen_q} \mapsto \mathcal{E}_{\sigma}[e_q]\} \end{array} \right) \\
\Gamma, \mu_1 \oplus \mu_2 \oplus \mu_3, \sigma_1 \vdash \langle stmts \rangle \Downarrow_{stmts} (stmts', locals') \\
\hline
\Gamma, \mu, \sigma \vdash \langle id(e_1, \dots, e_q) [id_1, \dots, id_n] [id'_1, \dots, id'_m] \rangle \\
\Downarrow_{stmt} (stmts', [\text{VAR } id_{local_1} : \text{WIRE}; \dots; \text{VAR } id_{local_k} : \text{WIRE}] \# locals') \\
\hline
\mathcal{E}_{\sigma}[exp_1] > \mathcal{E}_{\sigma}[exp_2] \\
\Gamma, \mu, \sigma \vdash \langle (\text{GENERATE FOR } id_{index}=exp_1 \dots exp_2 \text{ BEGIN } stmts \text{ END}) \rangle \Downarrow_{stmt} ([], []) \\
\hline
\mathcal{E}_{\sigma}[exp_1] \leq \mathcal{E}_{\sigma}[exp_2] \\
\Gamma, \mu, \sigma \oplus \{id_{index} \mapsto \mathcal{E}_{\sigma}[exp_1]\} \vdash \langle stmts \rangle \Downarrow_{stmts} (stmts', locals') \\
\Gamma, \mu, \sigma \vdash \left\langle \begin{array}{l} (\text{GENERATE FOR } id_{index}=(exp_1+1) \dots exp_2) \\ \text{BEGIN } stmts \text{ END} \end{array} \right\rangle \Downarrow_{stmt} (stmts'', locals'') \\
\hline
\Gamma, \mu, \sigma \vdash \left\langle \begin{array}{l} (\text{GENERATE FOR } id_{index}=exp_1 \dots exp_2) \\ \text{BEGIN } stmts \text{ END} \end{array} \right\rangle \\
\Downarrow_{stmt} (stmts' \# stmts'', locals' \# locals'') \\
\hline
\frac{\Gamma, \mu, \sigma \vdash \langle stmt_1 \rangle \Downarrow_{stmt} (stmts_1, locals_1) \dots \Gamma, \mu, \sigma \vdash \langle stmt_n \rangle \Downarrow_{stmt} (stmts_n, locals_n)}{\Gamma, \mu, \sigma \vdash \langle stmt_1; \dots; stmt_n \rangle \Downarrow_{stmts} (stmts_1 \# \dots \# stmts_n, locals_1 \# \dots \# locals_n)}
\end{array}$$

**Fig. 7.** Transition rules for flattening Pebble statements, based on  $\Downarrow$  rules for *stmts* and *stmt*.

names held in  $\mu$ . Local variables are bound to themselves, but are returned as declarations so that they are properly declared at run time. The blocks statements are flattened and returned with the local wire declarations. The two rules for loops apply the unfolding procedure at compile time to their subterms by incrementing the loop bound, creating primitive gate calls to implement the loop at run time. The rule for statements flattens each statement, and collects the intermediate calls and declarations together.

The flattening rule for the main block is shown in Fig. 8. An initial static environment  $\mu$  is created binding input, output and local wire names to themselves as these will be their run-time names. The body of the block is flattened with an empty value environment. The returned list of primitives gate calls forms the body of the flattened main block and the derived intermediate wires are declared local to this block.

$$\begin{array}{c}
 \mu_1 = \{id_{in_1} \mapsto id_{in_1}, \dots, id_{in_n} \mapsto id_{in_n}\} \\
 \mu_2 = \{id_{out_1} \mapsto id_{out_1}, \dots, id_{out_m} \mapsto id_{out_m}\} \\
 \mu_3 = \{id_1 \mapsto id_1, \dots, id_j \mapsto id_j\} \\
 \mu = \mu_1 \oplus \mu_2 \oplus \mu_3 \\
 \Gamma, \mu, \{\} \vdash \langle stms \rangle \Downarrow_{stms} (stms', [\text{VAR } id'_1:\text{WIRE}; \dots \text{VAR } id'_k:\text{WIRE}])
 \end{array}$$


---


$$\Gamma \vdash \left\langle \left( \begin{array}{l}
 \text{BLOCK main } [id_{in_1}:\text{WIRE}, \dots, id_{in_n}:\text{WIRE}] \\
 \qquad [id_{out_1}:\text{WIRE}, \dots, id_{out_m}:\text{WIRE}] \\
 \text{VAR } id_1:\text{WIRE}; \dots \text{VAR } id_j:\text{WIRE} \\
 \text{BEGIN } stms \text{ END}
 \end{array} \right) \right\rangle$$

$$\Downarrow_{main} \left( \begin{array}{l}
 \text{BLOCK main } [id_{in_1}:\text{WIRE}, \dots, id_{in_n}:\text{WIRE}] \\
 \qquad [id_{out_1}:\text{WIRE}, \dots, id_{out_m}:\text{WIRE}] \\
 \text{VAR } id_1:\text{WIRE}; \dots \text{VAR } id_j:\text{WIRE}; \\
 \text{VAR } id'_1:\text{WIRE}; \dots \text{VAR } id'_k:\text{WIRE} \\
 \text{BEGIN } stms' \text{ END}
 \end{array} \right)$$

**Fig. 8.** Transition rules for flattening the main block.

To illustrate how blocks are flattened and local variables are renamed, we shall use the following example which creates a row of two not-gates:

```

BLOCK notrow (n) [vin:WIRE] [vout:WIRE]
  VAR inter:VECTOR (n..0) OF WIRE
  VAR i:NUM
BEGIN
  connect [vin] [inter(0)];
  GENERATE FOR i=0..n-1
  BEGIN
    not [inter(i)] [inter(i+1)]
  END;
  connect [inter(n)] [vout]
END;

```

```
BLOCK main [x:WIRE] [y:WIRE]
BEGIN
  notrow (2) [x] [y]
END;
```

The circuit is defined in terms of two primitive components: `connect` which links two wires together, and the `not` gate. Figure 9 demonstrates how a block call is flattened given previously established environments. The called block is initially renamed; the environment  $\sigma$  is created for the generic values; the environment  $\mu$  maps parameter names to their original names; the block's statements are then flattened to create a list of primitive calls, and returned along with the distinct local wire definitions.

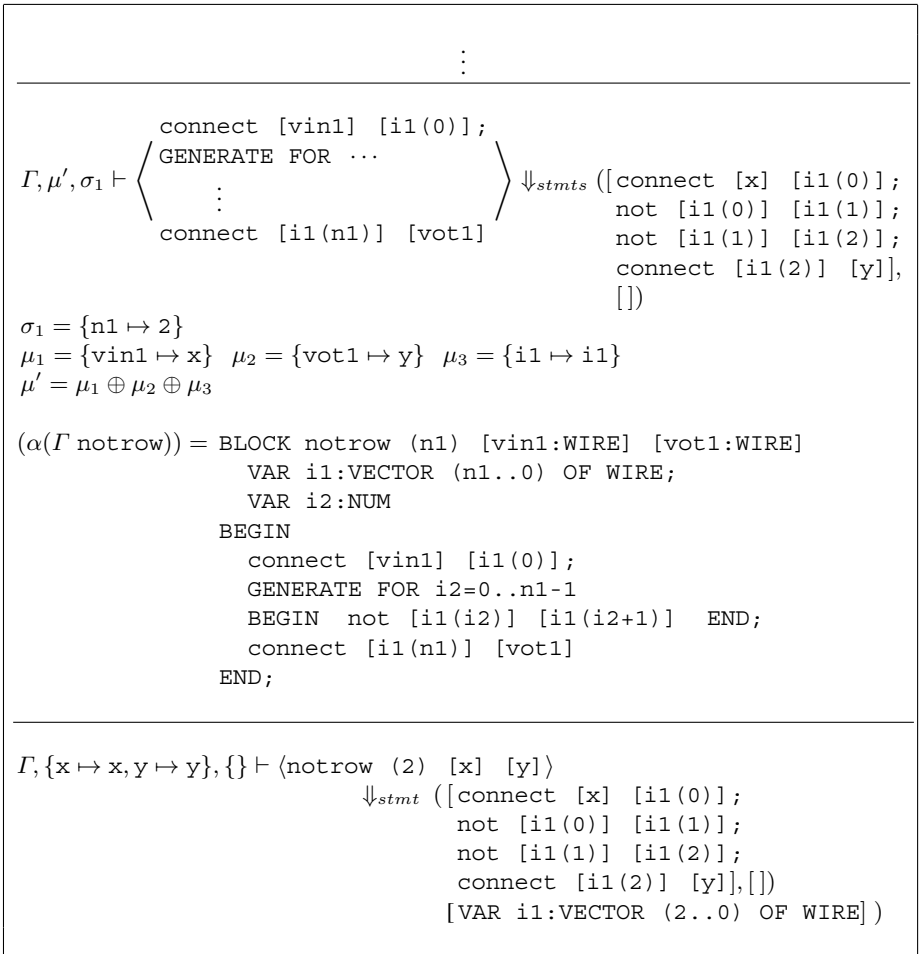


Fig. 9. Fragment of the proof tree for flattening notrow.

## 4 Correctness Proof

We can now present the main correctness theorem for flattening hierarchical blocks and `GENERATE-FOR` loops. This result, given by Equation (2), relies on Lemmas (3), (4) and (5). Each lemma is an instance of the commuting diagram given in Fig. 3, and involves the environment invariant given by Equation (1). At each syntactic level, they show how the sequence of outputs generated by the hierarchical definition can be calculated by first flattening the term and then using the simplified semantics. With the addition of `GENERATE-IF` statements, recursive block definitions can result in non-terminating programs. In these cases the flattening procedure will also fail to terminate.

The main theorem states that, from a given Hierarchical Pebble description consisting of a `main` block, a block environment  $\Gamma$ , a primitive gate environment  $\delta$ , and a sequence of input wires  $[v_1, \dots, v_n]$ , one can calculate the sequence of outputs derived from the circuit's proof trees by first unfolding the description to create one large `main` block, to which the flattening rules can be applied:

$$\begin{aligned} \Gamma, \delta \vdash \langle \text{main}, [v_1, \dots, v_n] \rangle &\Rightarrow_{\text{main}} [v'_1, \dots, v'_m] \\ &\implies \Gamma \vdash \text{main} \Downarrow_{\text{main}} \text{main}' \\ \wedge \quad \delta \vdash \langle \text{main}', [v_1, \dots, v_n] \rangle &\rightarrow_{\text{main}} [v'_1, \dots, v'_m] \end{aligned} \quad (2)$$

This result can be proved by structural induction on the rules for the main block (Figs. 4, 6 and 8). We establish the invariant on wire environments initially, and we use Lemma (3) below to show how it holds on completion of the block's statements so that the same final values are derived.

From a given Hierarchical Pebble statement list, a block environment  $\Gamma$ , a primitive gate environment  $\delta$ , a local value environment  $\sigma$  and a local wire environment  $\rho$ , we can calculate the wire bindings  $\rho'$  derived from a successful completion of the statements, by staging the computation in two. The first stage flattens the statements into a list of primitive calls, where local wire names are mapped to their original definitions using the static environment  $\mu$ , and a distinct list of local wire declarations is created. The second stage applies the Flattened Pebble rules to the primitive gate call list using the dynamic wire environment  $\rho_d$ . An environment  $\rho'_d$  can be derived that will contain the same bindings as those for the hierarchical statements. This implication requires the invariant, given by Equation (1), to hold for wire environments:

$$\begin{aligned} \Gamma, \delta, \sigma \vdash \langle \text{stm}ts, \rho \rangle &\Rightarrow_{\text{stm}ts} \rho' \\ &\implies \Gamma, \sigma, \mu \vdash \text{stm}ts \Downarrow_{\text{stm}ts} (\text{stm}ts', \text{loc}als') \\ \wedge \quad \forall id \cdot \rho \text{ id} &= \rho_d(\mu \text{ id}) \\ \wedge \quad \forall id \cdot \rho' \text{ id} &= \rho'_d(\mu \text{ id}) \\ \wedge \quad \delta \vdash \langle \text{stm}ts', \rho_d \rangle &\rightarrow_{\text{stm}ts} \rho'_d \end{aligned} \quad (3)$$

This lemma can be proved by induction on the length of derivation sequences using Lemma (4) below; it completes the presentation of the main theorem.

The next lemma deals mainly with `GENERATE-FOR` loops (Figs. 4, 5 and 7). From a given Hierarchical Pebble statement, a block environment  $\Gamma$ , a primitive gate environment  $\delta$ , a local value environment  $\sigma$ , and a local wire environment  $\rho$ , one can calculate a set of wire bindings  $\rho'$  derived from the successful completion of the statement, by first

flattening the statement using the static environment  $\mu$ , and then executing the derived statements in the dynamic wire environment  $\rho_d$ :

$$\begin{aligned}
& \Gamma, \delta, \sigma \vdash \langle stmt, \rho \rangle \Rightarrow_{stmt} \rho' \\
& \implies \Gamma, \sigma, \mu \vdash stmt \Downarrow_{stmt} (stmts', locals') \\
& \wedge \forall id \cdot \rho id = \rho_d(\mu id) \\
& \wedge \forall id \cdot \rho' id = \rho'_d(\mu id) \\
& \wedge \delta \vdash \langle stmts', \rho_d \rangle \rightarrow_{stmt} \rho'_d
\end{aligned} \tag{4}$$

This result can be proved by structural induction on statements: primitive gate calls, parametrised gate calls and loops. The first two cases are straightforward, once the invariants of the environments have been established. The third case, however, requires Lemma (5) to show that the appropriate final environment can be derived after staging:

$$\begin{aligned}
& \left( \begin{array}{l} \Gamma, \delta, \sigma_1 \vdash \langle stmts_1, \rho \rangle \Rightarrow_{stmts} \rho' \\ \wedge \Gamma, \delta, \sigma_2 \vdash \langle stmts_2, \rho' \rangle \Rightarrow_{stmts} \rho'' \end{array} \right) \\
& \implies \Gamma, \sigma_1, \mu \vdash stmts_1 \Downarrow_{stmts} (stmts'_1, locals'_1) \\
& \wedge \Gamma, \sigma_2, \mu \vdash stmts_2 \Downarrow_{stmts} (stmts'_2, locals'_2) \\
& \wedge \forall id \cdot \rho id = \rho_d(\mu id) \\
& \wedge \forall id \cdot \rho'' id = \rho''_d(\mu id) \\
& \wedge \delta \vdash \langle stmts_1 \# stmts_2, \rho_d \rangle \rightarrow_{stmts} \rho''_d
\end{aligned} \tag{5}$$

This lemma states that reducing a statement list  $stmts_1$  in the value environment  $\sigma_1$  with wire bindings  $\rho$ , followed by reducing a second statement list  $stmts_2$  in  $\sigma_2$  yielding a final wire environment  $\rho''$ , can be derived by first flattening the two statement lists and then executing the concatenation of the two primitive gate call lists. The lemma can be proved by induction on the length of derivation sequences.

## 5 Compiler Development

This section reflects on the implications of our approach for compiler development. Natural semantic rules, as used in specifying Pebble, rely on notions of pattern matching, inference rules and operational semantics. They can be captured in a theorem prover [2], [17] or translated into Horn clauses via a metalanguage such as Typol [3]. Since the transition rules for flattening Hierarchical Pebble descriptions permit a left-right and top-down construction of the proof tree with no backtracking, we can replace a resolution engine by a functional evaluator based on pattern matching [1] to improve efficiency.

In practice, an implementation in a functional language of the core flattening procedure (Fig. 10) follows naturally from the rules in Fig. 7. For a particular language construct, a function definition is created that pattern matches its goal and obtains the result from the intermediate transitions by means of a `where` clause. This technique offers a way of automatically producing a functional implementation of the compilation tools directly from their specifications.

It is educational to compare the original, hand-developed implementation of the flattening procedure, and the new version in Fig. 10 which results from the specification and proof exercises. The new version is better than the original version in all aspects:

```

data Exp = Number Int | Var String
         | Binop (Exp,Bop,Exp) | Unop (Uop,Exp)
data Type = WIRE
data Dec = VARW (String,Type) | VARN String
data Stmt = PrimCall (String,[String],[String])
           | BlkCall (String,[Exp],[String],[String])
           | Loop (String,Exp,Exp,[Stmt])
data Blk = Block (String,[String],
                 [(String,Type)],[(String,Type)],[Dec],[Stmt])

fetch :: [(String,a)] -> String -> a
...
eval_exp :: Exp -> [(String,Int)] -> Int
...

flatten_stmt :: ((String,Blk)],
               [(String,String)],[(String,Int)]) -> Stmt -> ([Stmt],[Dec])
flatten_stmt (gamma,mu,sigma) (PrimCall (pnm,args1,args2))
  = ([PrimCall (pnm,map (fetch mu) args1,
                map (fetch mu) args2)],[])
flatten_stmt (gamma,mu,sigma) (BlkCall (bnm,gens,args1,args2))
  = (stmts',[(VARW d) | (VARW d) <- decs] ++ locals')
  where
    (Block (nm,genyms,parms1,parms2,decs,stmts))
      = rename (fetch gamma bnm)
    sigma1 = zip genyms
              (map (\ e -> eval_exp e sigma) gens)
    mu1    = zip parms1 (map (fetch mu) args1)
    mu2    = zip parms2 (map (fetch mu) args2)
    mu3    = [(id,id) | (VARW (id,WIRE)) <- decs]
    mu'    = mu1 ++ mu2 ++ mu3
    (stmts',locals')
      = flatten_stmts (gamma,mu',sigma1) stmts
flatten_stmt (gamma,mu,sigma) (Loop (nm,e1,e2,stmts))
  | n1>n2    = ([],[])
  | otherwise = (stmts'++stmts'',locals'++locals'')
  where
    n1 = eval_exp e1 sigma
    n2 = eval_exp e2 sigma
    (stmts',locals')
      = flatten_stmts (gamma,mu,(nm,n1):sigma) stmts
    (stmts'',locals'')
      = flatten_stmt (gamma,mu,sigma)
                    (Loop (nm,Binop (e1,Add,Number 1),e2,stmts))
flatten_stmts :: ((String,Blk)],
                 [(String,String)],[(String,Int)]) -> [Stmt]-> ([Stmt],[Dec])
flatten_stmts (gamma,mu,sigma) = unzip . map flatten_stmt

```

**Fig. 10.** Flattening Pebble in Haskell.

it is clearer, more concise, more robust and more efficient. The main reason is that, based on the formal development, the new version separates variable renaming from the flattening procedure. The original implementation, in contrast, attempted the former without properly considering the latter: the mingling of the two leads to situations where the renaming process is deeply nested within the unfolding procedure, leaving little scope for further optimisations. The new version is amenable to further optimisations, such as the use of de Bruijn indices to avoid the costly rename function [5]. Further refinements would lead to a highly efficient imperative implementation.

Our experience shows that deriving provably-correct compiler implementations can benefit their efficiency, in addition to increasing the confidence in their correctness. Furthermore, it appears possible that the verification of such implementations for domain-specific languages such as Pebble can be mechanised using custom proof engines [2]. Further work, however, is needed to explore this possibility thoroughly.

## 6 Concluding Remarks

While the version of Pebble described in this paper does not include advanced abstraction mechanisms, current work involves extending Pebble with polymorphic variables, records and higher-order functions. These features enable a combinator style of development [10] that tends to simplify the hardware design process.

Our extended compilation strategy infers the types of the polymorphic variables, unfolds the record definitions and instantiates higher-order functions prior to compile time to create a Hierarchical Pebble description. The correctness proof for Polymorphic Pebble is very similar to that for Hierarchical Pebble. An intermediate environment mapping polymorphic variables to types is used to create distinct blocks, and it leads to an invariant equation similar to Equation (1). Higher-Order Pebble enables nested function calls which require lambda lifting before the calls can be unfolded. In this way the ability to generate correct parametrised VHDL will be maintained.

The combinator style of description facilitates the formulation of correctness-preserving algebraic transformations for design development [10]. The proposed extensions of Pebble take us a step closer to providing, for instance, a generic transformation rule [13] which can be used to derive pipelined designs from a non-pipelined design. Further work will generalise our approach to deal with relational descriptions [10].

Our research contributes to insights about abstraction mechanisms and their validated implementations. It also provides a useful foundation on which further work, such as verifying tools for pipeline optimisation [19], can be based. We believe that provably-correct tools will have a profound impact on understanding the scope and effectiveness of hardware synthesis algorithms and their implementation.

**Acknowledgements.** Our thanks to the many Pebble developers, users and advisors, particularly Chris Booth, Ties Bos, Arran Derbyshire, Florent Dupont-De-Dinechin, Mike Gordon, He Jifeng, Tony Hoare, Danny Lee, James Rice, Richard Sandiford, Seng Shay Ping, Nabeel Shirazi, Dimitris Siganos, Henry Styles, Tim Todman and Markus Weinhardt, for their patience, help and encouragement. Thanks also to the comments and suggestions of anonymous reviewers of this paper. We acknowledge the support of

UK Engineering and Physical Sciences Research Council, Celoxica, and Xilinx. This work was carried out as part of Technology Group 10 of UK MOD's Corporate Research Programme.

## References

1. I. Attali, J. Chazarain and S. Gillette, "Incremental evaluation of natural semantics specifications", *Proc. 4th Int. Symp. on Programming Language Implementation and Logic Programming*, LNCS 631, Springer, 1992.
2. L.A. Dennis et. al., "The PROSPER toolkit", *Proc. 6th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, LNCS 1785, Springer, 2000.
3. T. Despeyroux, "Executable specification of static semantics", *Semantics of Data Types*, LNCS 173, Springer, 1984.
4. D. Eisenbiegler, C. Blumenroehr and R. Kumar, "Implementation issues about the embedding of existing high level synthesis algorithms in HOL", *Theorem Proving in Higher Order Logics*, LNCS 1125, Springer, 1996.
5. C. Hankin, *Lambda Calculus, A guide for Computer Scientists*, Oxford University Press, 1994.
6. J. He, G. Brown, W. Luk and J.W. O'Leary, "Deriving two-phase modules for a multi-target hardware compiler", *Designing Correct Circuits*, Springer Electronic Workshop in Computing, 1996.
7. J. He, I. Page and J.P. Bowen, "Towards a provably correct hardware implementation of occam", *Correct Hardware Design and Verification Methods*, LNCS 683, Springer, 1993.
8. C.A.R. Hoare, J. He and A. Sampaio, "Normal form approach to compiler design", *Acta Informatica*, Vol. 30, pp. 701-739, 1993.
9. G. Jones and M. Sheeran, "Timeless truths about sequential circuits", *Concurrent Computations: Algorithms, Architectures and Technology*, S.K. Tewksbury et. al. (eds.), Plenum Press, 1988.
10. G. Jones and M. Sheeran, "Circuit design in Ruby", *Formal Methods for VLSI Design*, J. Staunstrup (ed.), North-Holland, 1990.
11. U. Jørring and W. Scherlis, "Compilers and staging transformations", *Proc. ACM Symp. on Principles of Programming Languages*, ACM Press, 1986.
12. W. Luk and S.W. McKeever, "Pebble: a language for parametrised and reconfigurable hardware design", *Field-Programmable Logic and Applications*, LNCS 1482, Springer, 1998.
13. W. Luk et. al., "A framework for developing parametrised FPGA libraries", *Field-Programmable Logic and Applications*, LNCS 1142, Springer, 1996.
14. W. Luk et. al., "Reconfigurable computing for augmented reality", *Proc. IEEE Symp. on Field-Programmable Custom Computing Machines*, IEEE Computer Society Press, 1999.
15. N. Mansouri and R. Vemuri, "A methodology for completely automated verification of synthesized RTL designs and its integration with a high-level synthesis tool", *Formal Methods in Computer-Aided Design*, LNCS 1522, Springer, 1998.
16. H.R. Nielson and F. Nielson, *Semantics with Applications*, John Wiley and Sons, 1992.
17. F. Pfenning and C. Schrmann, "System description: Twelf – a meta-logical framework for deductive systems", *Proc. Int. Conf. on Automated Deduction*, LNAI 1632, Springer, 1999.
18. M. Sheeran, S. Singh and G. Stalmarck, "Checking safety properties using induction and a SAT-solver", *Proc. Int. Conf. on Formal Methods in CAD*, LNCS 1954, Springer, 2000.
19. M. Weinhardt and W. Luk, "Pipeline vectorization", *IEEE Trans. CAD*, Vol. 20, No. 2, 2001, pp. 234-248.