Equivalence Checking a Floating-point Unit against a High-level C Model

Extended Version

Rajdeep Mukherjee[1], Saurabh Joshi[2], Andreas Griesmayer[3],
Daniel Kroening[1], and Tom Melham[1]

[1] University of Oxford, UK
[2] IIT Hyderabad, India
[3] ARM Limited

{rajdeep.mukherjee,kroening,tom.melham}@cs.ox.ac.uk,
sbjoshi@iith.ac.in, andreas.griesmayer@arm.com

Abstract. Semiconductor companies have increasingly adopted a methodology that starts with a system-level design specification in C/C++/SystemC. This model is extensively simulated to ensure correct functionality and performance. Later, a Register Transfer Level (RTL) implementation is created in Verilog, either manually by a designer or automatically by a high-level synthesis tool. It is essential to check that the C and Verilog programs are consistent. In this paper, we present a two-step approach, embodied in two equivalence checking tools, VERIFOX and HW-CBMC, to validate designs at the software and RTL levels, respectively. VERIFOX is used for equivalence checking of an untimed software model in C against a high-level reference model in C. HW-CBMC verifies the equivalence of a Verilog RTL implementation against an untimed software model in C. To evaluate our tools, we applied them to a commercial floating-point arithmetic unit (FPU) from ARM and an open-source dual-path floating-point adder.

1 Introduction

One of the most important tasks in Electronic Design Automation (EDA) is to check whether the low-level implementation (RTL or gate-level) complies with the system-level specification. Figure 1 illustrates the role of equivalence checking (EC) in the design process. In this paper, we present a new EC tool, VERIFOX, that is used for equivalence checking of an untimed software (SW) model against a high-level reference model. Later, a Register Transfer Level (RTL) model is implemented, either manually by a hardware designer or automatically by a synthesis tool. To guarantee that the RTL is consistent with the SW model, we use an existing tool, HW-CBMC [15], to check the correctness of the synthesized hardware RTL against a SW model.

In this paper, we address the most general and thus most difficult variant of EC: the case where the high-level and the low-level design are substantially different. State-of-the-art tools, such as Hector [14] from Synopsys and SLEC from Calypto,[4] rely on
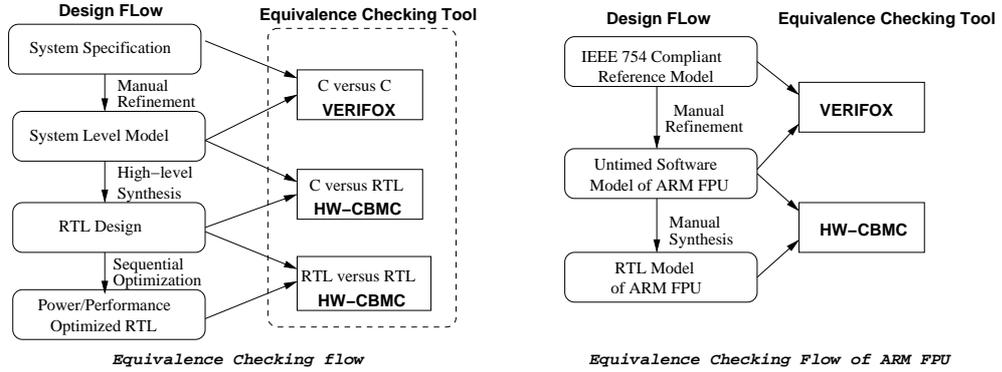
---

[4] http://calypto.com/en/products/slec/

Fig. 1. Electronic Design Automation Flow

equivalence points [18], and hence they are ineffective in this scenario. We present an approach based on bounded analysis, embodied in the tools VERIFOX and HW-CBMC, that can handle arbitrary designs.

VERIFOX is used for equivalence checking of an untimed software model against a high-level reference model and HW-CBMC is used for equivalence checking of the RTL implementation against a software model. EC is broadly classified into two separate categories: combinational equivalence checking (CEC) and sequential equivalence checking (SEC). CEC is used for a pair of models that are cycle accurate and have the same state-holding elements. SEC is used when the high-level model is not cycle accurate or has a substantially different set of state-holding elements [1,11]. It is well-known that EC of floating-point designs is difficult [12, 19]. So there is a need for automatic tools that formally validate floating-point designs at various stages of the synthesis flow, as illustrated by right side flow of Figure 1.

Contributions: In this paper, we sketch two significant equivalence-verification tools:

1. VERIFOX, a tool for equivalence checking of software models given as C programs. We present a path-based symbolic execution tool, VERIFOX, for bounded equivalence checking of floating-point software implementations against a IEEE 754 compliant reference model. VERIFOX supports C89, C99 standards in the front-end. VERIFOX also supports SAT and SMT backends for constraint solving. VERIFOX is available at `http://www.cprover.org/verifox`.

2. HW-CBMC, a tool for C versus RTL equivalence checking. HW-CBMC is used for bounded equivalence checking of Verilog RTL against C/C++ models. HW-CBMC supports IEEE 1364-2005 System Verilog standards and the C89, C99 standards. HW-CBMC generates a monolithic formula from the C and RTL description, which is then checked with SAT/SMT solvers.
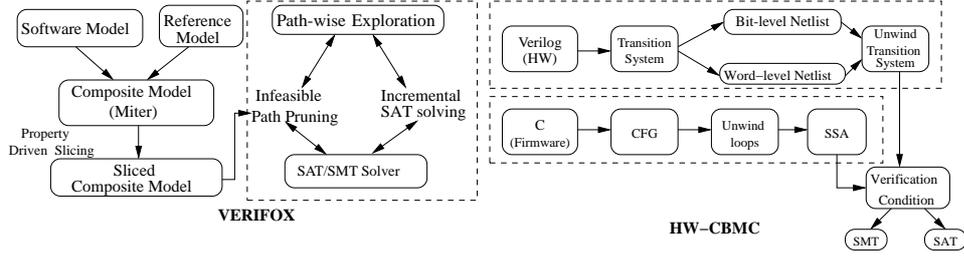
Fig. 2. VERIFOX and HW-CBMC Tool Architecture

2 VERIFOX: A tool for equivalence checking of C programs

VERIFOX is a path-based symbolic execution tool for equivalence checking of C programs. The tool architecture is shown on the left side of Figure 2. VERIFOX supports the C89 and C99 standards. The key feature is symbolic reasoning about equivalence between FP operations. To this end, VERIFOX implements a model of the core IEEE 754 arithmetic operations—single- and double-precision addition, subtraction, multiplication, and division—which can be used as reference designs for equivalence checking. So VERIFOX does not require external reference models for equivalence checking of floating-point designs. This significantly simplifies the users effort to do equivalence checking at software level. The reference model in VERIFOX is equivalent to the Soft-float model.[5] VERIFOX also supports SAT and SMT backends for constraint solving.

Given a reference model, an implementation model in C and a set of partition constraints, VERIFOX performs depth-first exploration of program paths with certain optimizations, such as eager infeasible path pruning and incremental constraint solving. This enables automatic decomposition of the verification state-space into subproblems, by input-space and/or state-space decomposition. The decomposition is done in tandem in both models, exploiting the structure present in the high-level model. The approach generates many but simpler SAT/SMT queries, similar to the technique followed in KLEE [4]. The main focus of our technique is to pass only those verification conditions to the underlying solver for which the corresponding path conditions are feasible with

---

[5] http://www.jhauser.us/arithmetic/SoftFloat.html

Program	Path Constraint 1	Path Constraint 2	Path Constraint 3	Monolithic Path Constraint
```void top(){ if(reset) { x=0; y=0; } else { if(a > b) x=a+b; else y=(a & 3)<<b; }}```	$C_1 \equiv$ $reset_1 \neq 0 \wedge$ $x_2 = 0 \wedge$ $y_2 = 0$	$C_2 \equiv$ $reset_1 = 0 \wedge$ $b_1 \not\geq a_1 \wedge$ $x_3 = a_1 + b_1$	$C_3 \equiv$ $reset_1 = 0 \wedge$ $b_1 \geq a_1 \wedge$ $y_3 = (a_1 \& 3)$ $\ll b_1$	$C \iff ((guard_1 = \neg(reset_1 = 0)) \wedge$ $(x_2 = 0) \wedge (y_2 = 0) \wedge$ $(x_3 = x_1) \wedge (y_3 = y_1) \wedge$ $(guard_2 = \neg(b_1 >= a_1)) \wedge$ $(x_4 = a_1 + b_1) \wedge (x_5 = x_3) \wedge$ $(y_4 = (a_1 \& 3) \ll b_1) \wedge$ $(x_6 = ite(guard_2, x_4, x_5)) \wedge$ $(y_5 = ite(guard_2, y_3, y_4)) \wedge$ $(x_7 = ite(guard_1, 0, x_6)) \wedge$ $(y_6 = ite(guard_1, 0, y_5)))$

Fig. 3. Single-path and Monolithic Symbolic Execution

respect to the property under consideration and the partitioning constraints such as case splitting.

Figure 3 shows three feasible path constraints corresponding to the three paths in the program on the left. In contrast, the last column of Figure 3 shows monolithic path-constraint generated by HW-CBMC.

***Incremental solving in* VERIFOX.** VERIFOX can be run in two different modes: partial incremental and full incremental. In partial incremental mode, only one solver instance is maintained while going down a single path. So when making a feasibility check from one branch b_1 to another branch b_2 along a single path, only the program segment from b_1 to b_2 is encoded as a constraint and added to the existing solver instance. Internal solver states and the information that the solver gathers during the search remain valid as long as all the queries that are posed to the solver in succession are monotonically stronger. If the solver solves a formula ϕ, then posing $\phi \wedge \psi$ as a query to the same solver instance allows one to reuse solver knowledge it has already acquired, because any assignment that falsifies ϕ also falsifies $\phi \wedge \psi$. Thus the solver need not revisit the assignments that it has already ruled out. This results in speeding up the feasibility check of the symbolic state at b_2, as the feasibility check at b_1 was *true*. A new solver instance is used to explore a different path, after the current path is detected as infeasible.

In full incremental mode, only one solver instance is maintained throughout the whole symbolic execution. Let $\phi_{b_1 b_2}$ denote the encoding of the path fragment from b_1 to b_2. It is added in the solver as $B_{b_1 b_2} \Rightarrow \phi_{b_1 b_2}$. Then, $B_{b_1 b_2}$ is added as a *blocking variable*[6] to enforce constraints specified by $\phi_{b_1 b_2}$. Blocking variables are treated specially inside the solvers: unlike regular variables or clauses, the blocking can be removed in subsequent queries without invalidating the solver instance. When one wants to backtrack the symbolic execution, the blocking $B_{b_1 b_2}$ is removed and a unit clause $\neg B_{b_1 b_2}$ is added to the solver, thus effectively removing $\phi_{b_1 b_2}$.

3   HW-CBMC: A tool for equivalence checking of C and RTL

HW-CBMC is used for bounded equivalence checking of C and Verilog RTL. The tool architecture is shown on the right side of Figure 2. HW-CBMC supports IEEE 1364-2005 System Verilog standards and the C89, C99 standards. HW-CBMC maintains two separate flows for hardware and software. The top flow in Figure 2 uses synthesis to obtain either a bit-level or a word-level netlist from Verilog RTL. The bottom flow illustrates the translation of the C program into static single assignment (SSA) form [9]. These two flows meet only at the solver. Thus, HW-CBMC generates a monolithic formula from the C and RTL description, which is then checked with SAT/SMT solvers. HW-CBMC provides specific handshake primitives such as *next_time frame*() and *set_inputs*() that direct the tool to set the inputs to the hardware signals and advance the clock, respectively. The details of HW-CBMC are available online.[7]

---

[6] The SAT community uses the term *assumption variables* or *assumptions*, but we will use the term blocking variable to avoid ambiguity with assumptions in the program.

[7] http://www.cprover.org/hardware/sequential-equivalence/

4 Experimental Results

In this section, we report experimental results for equivalence checking of difficult floating-point designs. All our experiments were performed on an Intel® Xeon® machine with 3.07 GHz clock speed and 48 GB RAM. All times reported are in seconds. MiniSAT-2.2.0 [10] was used as underlying SAT solver with VERIFOX 0.1 and HW-CBMC 5.4. The timeout for all our experiments was set to 2 hours.

Proprietary Floating-point Arithmetic Core: We verified parts of a floating-point arithmetic unit (FPU) of a next generation ARM® GPU. The FP core is primarily composed of single- and double-precision *ADD*, *SUB*, *FMA* and *TBL* functional units, the register files, and interface logic. The pipelined computation unit implements FP operations on a 128-bit data-path. In this paper, we verified the single-precision addition (*FP-ADD*), rounding (*FP-ROUND*), minimum (*FP-MIN*) and maximum (*FP-MAX*) operations. The FP-ADD unit can perform two operations in parallel by using two 64-bit adders over multiple pipeline stages. Each 64-bit unit can also perform operations with smaller bit widths. The FPU decodes the incoming instruction, applies the input modifiers and provides properly modified input data to the respective sub-unit. The implementation is around 38000 LOC, generating tens of thousands of gates. We obtained the SW model (in C) and the Verilog RTL model of the FPU core from ARM. (Due to proprietary nature of the FPU design, we can not share the commercial ARM IP.)

Open-source Dual-path Floating-point Adder: We have developed both a C and a Verilog implementation of an IEEE-754 32-bit single-precision dual-path floating point adder/subtractor. This floating-point design includes various modules for packing, unpacking, normalizing, rounding and handling of infinite, normal, subnormal, zero and NaN (Not-a-Number) cases. We distribute the C and RTL implementation of the dual-path FP adder at `http://www.cprover.org/verifox`.

Reference Model: The IEEE 754 compliant floating-point implementations in VERIFOX are used as the golden reference model for equivalence checking at the software level. For equivalence checking at the RTL phase, we used the untimed software model from ARM as the reference model, as shown on the right side of Figure 1.

Miters for Equivalence Checking: A miter circuit [3] is built from two given circuits *A* and *B* as follows: identical inputs are fed into *A* and *B*, and the outputs of *A* and *B* are compared using a comparator. For equivalence checking at software level, one of the circuits is a SW program and the other is a high-level reference model. For the RTL phase, one of the circuits is a SW program treated as reference model and the other is an RTL implementation.

Case-splitting for Equivalence Checking: Case-splitting is a common practice to scale up formal verification [12, 14, 19] and is often performed by user-specified assumptions. The CPROVER_assume(c) statement instructs HW-CBMC and VERIFOX to restrict the analysis to only those paths satisfying a given condition c. For example, we can limit the analysis to those paths that are exercised by inputs where the rounding mode is nearest-even (RNE) and both input numbers are NaNs by adding the following line:

```
CPROVER_assume(roundingMode==RNE && uf_nan && ug_nan);
```

| Design | Case-splitting | | | | | No-partition |
	INF	ZERO	NaN	SUBNORMAL	NORMAL	Total
Equivalence checking at Software Level (VERIFOX)						
FP-ADD	9.56	11.54	9.95	1124.18	77.74	1566.72
FP-ROUND	1.24	1.36	1.32	3.78	1.63	4.71
FP-MIN	9.76	9.85	9.78	28.67	9.86	48.70
FP-MAX	9.80	9.88	9.97	28.70	9.90	35.81
DUAL-PATH ADDER	3.15	3.11	2.14	88.12	55.28	497.67
Equivalence checking at RTL (HW-CBMC)						
FP-ADD	18.12	18.02	17.87	18.73	39.60	40.72
FP-ROUND	11.87	12.73	13.44	13.67	14.03	14.11
FP-MIN	13.72	13.62	ERROR	14.10	14.08	14.15
FP-MAX	13.70	13.58	ERROR	14.09	14.06	14.05
DUAL-PATH ADDER	0.88	0.87	0.99	169.49	22.42	668.61

Table 1. Equivalence checking of ARM FPU and DUAL-PATH Adder (All time in seconds)

Discussion of Results: Table 1 reports the run times for equivalence checking of the ARM FPU and the dual-path FP adder. Column 1 gives the name of FP design and columns 2–6 show the runtimes for partition modes INF, ZERO, NaN, SUBNORMAL, and NORMAL respectively. For example, the partition constraint 'INF' means addition of two infinite numbers. Column 7 reports the total time for equivalence checking without any partitioning.

VERIFOX successfully proved the equivalence of all FP operations in the SW implementation of ARM FPU against the built-in reference model. However, a bug in FP-MIN and FP-MAX (reported as ERROR in Table 1) is detected by HW-CBMC in the RTL implementation of ARM FPU when checked against the SW model of ARM FPU for the case when both the input numbers are NaN. This happens mostly due to bugs in the high-level synthesis tool or during manual translation of SW model to RTL. VERIFOX and HW-CBMC is able to detect bugs in the SW and RTL models of these designs respectively – thereby emphasizing the need for equivalence checking to validate the synthesis process during the EDA flow. Further, we investigate the reason for higher verification times for subnormal numbers compared to normal, infinity, NaN's and zero's. This is attributed to higher number of paths in subnormal case compared to INF, NaN's and zero's. Closest to our floating-point symbolic execution technique in VERIFOX is the tool KLEE-FP [8]. We could not, however, run KLEE-FP on the software models because the front-end of KLEE-FP failed to parse the ARM models.

5 Related work

The concept of symbolic execution [4, 7, 13] is prevalent in the software domain for automated test generation as well as bug finding. Tools such as Dart [13], Klee [4], EXE [5], Cloud9 [16] employ such a technique for efficient test case generation and bug finding. By contrast, we used path-wise symbolic execution for equivalence checking of software models against a reference model. A user-provided assumption specifies certain testability criteria that render majority of the design logic irrelevant [12, 14, 19], thus giving rise to large number of infeasible paths in the design. Conventional SAT-based bounded model checking [2, 6, 15] can not exploit this infeasibility because these

techniques create a monolithic formula by unrolling the entire transition system up to a given bound, which is then passed to SAT/SMT solver. These tools perform case-splitting at the level of solver through the effect of constant propagation. Optimizations such as eager path pruning combined with incremental encoding enable VERIFOX to address this limitation.

6   Concluding Remarks

In this paper we presented VERIFOX, our path-based symbolic execution tool, which is used for equivalence checking of arbitrary software models in C. The key feature of VERIFOX is symbolic reasoning on the equivalence between floating-point operations. To this end, VERIFOX implements a model of the core IEEE 754 arithmetic operations, which can be used for reference models. Further, to validate the synthesis of RTL from software model, we used our existing tool, HW-CBMC, for equivalence checking of RTL designs against the software model used as reference. We successfully demonstrated the utility of our equivalence checking tool chain, VERIFOX and HW-CBMC, on a large commercial FPU core from ARM and a dual-path FP adder. Experience suggests that the synthesis of software models to RTL is often error prone—this emphasizes the need for automated equivalence checking tools at various stages of EDA flow. In the future, we plan to investigate various path exploration strategies and path-merging techniques in VERIFOX to further scale equivalence checking to complex data and control intensive designs.

Acknowledgements

References

1. Baumgartner, J., Mony, H., Paruthi, V., Kanzelman, R., Janssen, G.: Scalable sequential equivalence checking across arbitrary design transformations. In: ICCD. pp. 259–266. IEEE (2006)
2. Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded model checking. Advances in Computers 58, 117–148 (2003)
3. Brand, D.: Verification of large synthesized designs. In: ICCAD. pp. 534–537 (1993)
4. Cadar, C., Dunbar, D., Engler, D.R.: KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs. In: OSDI. pp. 209–224. USENIX (2008)
5. Cadar, C., Ganesh, V., Pawlowski, P.M., Dill, D.L., Engler, D.R.: EXE: automatically generating inputs of death. ACM Trans. Inf. Syst. Secur. 12(2) (2008)
6. Clarke, E., Kroening, D.: Hardware verification using ANSI-C programs as a reference. In: Proceedings of the 2003 Asia and South Pacific Design Automation Conference. pp. 308–311. ASP-DAC, ACM (2003)
7. Clarke, L.A.: A system to generate test data and symbolically execute programs. IEEE Trans. Software Eng. 2(3), 215–222 (1976)

8. Collingbourne, P., Cadar, C., Kelly, P.H.J.: Symbolic crosschecking of floating-point and SIMD code. In: EuroSys. pp. 315–328 (2011)
9. Cytron, R., Ferrante, J., Rosen, B.K., Wegman, M.N., Zadeck, F.K.: An efficient method of computing static single assignment form. In: POPL. pp. 25–35. ACM (1989)
10. Eén, N., Biere, A.: Effective preprocessing in SAT through variable and clause elimination. In: SAT. pp. 61–75 (2005)
11. van Eijk, C.A.J.: Sequential equivalence checking without state space traversal. In: DATE. pp. 618–623. IEEE (1998)
12. Fujita, M.: Verification of arithmetic circuits by comparing two similar circuits. In: CAV. vol. 1102, pp. 159–168. Springer (1996)
13. Godefroid, P., Klarlund, N., Sen, K.: DART: directed automated random testing. In: PLDI. pp. 213–223 (2005)
14. Kölbl, A., Jacoby, R., Jain, H., Pixley, C.: Solver technology for system-level to RTL equivalence checking. In: DATE. pp. 196–201. IEEE (2009)
15. Kroening, D., Clarke, E., Yorav, K.: Behavioral consistency of C and Verilog programs using bounded model checking. In: DAC. pp. 368–371 (2003)
16. Kuznetsov, V., Kinder, J., Bucur, S., Candea, G.: Efficient state merging in symbolic execution. In: PLDI. pp. 193–204 (2012)
17. Weiser, M.: Program slicing. In: ICSE. pp. 439–449. IEEE (1981)
18. Wu, W., Hsiao, M.S.: Mining global constraints for improving bounded sequential equivalence checking. In: DAC. pp. 743–748. ACM (2006)
19. Xue, B., Chatterjee, P., Shukla, S.K.: Simplification of C-RTL equivalent checking for fused multiply add unit using intermediate models. In: ASP-DAC. pp. 723–728. IEEE (2013)

A Appendix

This appendix provides simple, illustrative examples of the use of VERIFOX and HW-CBMC, as well as further technical details.

Worked Example of VERIFOX

Figure 4 demonstrates the working of VERIFOX as a property verifier in the absence of a reference model. Note that equivalence checking is a special case of property verification where the property is replaced by a reference model. Hence, VERIFOX can be configured as a property verifier or as an equivalence checker.

Let us consider a software model as shown in column 1 in Figure 4. The program implements a high-level power management strategy to orchestrate various modules, such as, *core*, *memory* etc. Depending on the interrupt status (*env*), power modes (*mode*) and power gated logic (*power_gated*), the call to *core* or *memory* is made. These units are complex implementations of a processor core or a memory unit.

State-of-the-art verification tools may not be able verify the whole system due to resource limitations. Therefore, it is a common practice to write additional constraints, also known as *assumptions*, that exercise only a fragment of the entire state-space. Verification engine can use these assumptions to partition the state-space, thus decomposing a hard proof into simpler sub-proofs.

Software Model (in C)	Sliced program (wrt. property)	Pruned program (wrt. assumptions)
| ```
#define threshold 15
if(reset) {
  mode=TURN_OFF;
  feedback=0;
} else {
// Trigger IP if env is set
if(env) {
// check the voltage level
if(voltage_level < threshold)
 power_gated = 1;
else power_gated = 0;
// check the low-power modes
if(mode == STAND_BY ||
   mode == TURN_OFF) {
// power gated logic,
// call to core
if(power_gated) {
 core(reset,mode,power_gated
 ser_in,&buf_out);
 feedback = LOW;
}
else { // normal logic
  core(reset,mode,power_gated
  ser_in,&buf_out);
  feedback = buf_out; }}
}
else
 // call to memory
 memory(size);
}
``` | ```
#define threshold 15
if(reset) {
 mode=TURN_OFF;
 feedback=0;
} else {
// Trigger IP if env is set
if(env) {
// check the voltage level
if(voltage_level < threshold)
 power_gated = 1;
else power_gated = 0;
// check the low-power modes
if(mode == STAND_BY ||
mode == TURN_OFF) {
// power gated logic,
// call to core
if(power_gated) {
 core(reset,mode,power_gated
 ser_in,&buf_out);
  feedback = LOW;
}
else { // normal logic
  core(reset,mode,power_gated
  ser_in,&buf_out);
  feedback = buf_out;
  }
}}
}
```$Assertion : (!(rst != 0) || (feedback > 0))$ | $Assumption : (reset == 0) \wedge (env == 1)$<br>$\wedge (mode == STAND\_BY)$<br>$\wedge (voltage\_level == 10)$<br><br>```
#define threshold 15
// non-reset logic
else {
if(env) {
 // check the voltage level
 if(voltage_level < threshold)
  power_gated = 1;

 // check the low-power modes
 if(mode == STAND_BY ||
   mode == TURN_OFF) {
  // power gated logic,
  // call to core
  if(power_gated) {
   core(reset,mode,power_gated
   ser_in,&buf_out);
   feedback = LOW;
  }
 }
}
}
```$Assertion : (!(rst != 0) || (feedback > 0))$ |

Fig. 4. Example demonstrating automated slicing and path pruning in VERIFOX

Column 2 presents the result of property-driven slicing on the input program. This step is purely syntactic, meaning that we perform a backward dependency analysis [17] starting from the property which only preserve those program fragments that are relevant to the given property. The sliced program is then passed to the symbolic execution engine that performs eager infeasibility based path-pruning. The result of infeasible path pruning based on assumption is shown in column 3. This step is semantic because VERIFOX determines the feasibility of paths in the sliced program in an eager fashion with respect to the user-provided assumptions using satisfiability queries.

An important point to note here is that the number of path constraints after slicing and infeasible path pruning are significantly less compared to the initial program. Additionally, these per-path constraints are much easier to solve compared to a monolithic formula generated from a BMC-style symbolic execution tool.

***Command to run* VERIFOX.** Below are the commands to run VERIFOX in partial or full incremental mode. When VERIFOX is used as an equivalence checker, the input file is usually a miter in C which must include both the reference model and the implementation model. However, in the absence of a reference model, one can write assertions inside the software model to configure VERIFOX as a property verifier. The command line switch `--unwind` is used to specify the unwind depth for the software model. To use the SMT backend with VERIFOX, the command line switch is `--smt2`, followed by the name of the SMT solver, for example `--z3`. Note that the SMT solver must be installed in the system. The switch `--help` shows the available command line options for using VERIFOX.

```
// partial incremental mode with SAT
verifox-pi filename.c --unwind N
// full incremental mode with SAT
verifox-fi filename.c --unwind N
// partial incremental mode with SMT
verifox-pi filename.c --smt2 --z3
```

Worked Example of HW-CBMC

Figure 5 demonstrates the working of HW-CBMC as a C-RTL equivalence checker. Columns 1–3 present a C model of an up-counter, an RTL model of the same device, and a miter that feeds the same input to the C and RTL model and asserts equivalence of their outputs. HW-CBMC can be configured in *bit-level* or *word-level* mode. In bit-level mode, the input models are synthesized to And Inverter Graphs (AIG)[8] and then passed to the SAT solver. In word-level mode, the input models are synthesized into an intermediate word-level format, which are then despatched to a word-level SMT solver.

***Command to run* HW-CBMC.** Shown below are the commands to configure HW-CBMC in bit-level or word-level mode. The first command using `--gen-interface` is used to generate the interface for the hardware modules automatically. These interface signals are required to construct the miter as shown in column 3 of Figure 5. Note that

---

[8] http://fmv.jku.at/aiger/

C Program	Verilog RTL	Miter
| ```c
struct st_up_counter{
  unsigned char out;
};
struct st_up_counter
sup_counter;

void upcounter(unsigned char *out,
  _Bool enable, _Bool clk,
  _Bool reset)
{
 unsigned char out_old;
 out_old = sup_counter.out;
 if(reset)
 {
  sup_counter.out = 0;
 }
 else if(enable)
 {
  sup_counter.out =
   out_old + 1;
 }
}
``` | ```verilog
module up_counter(out,
   enable,clk, reset);
output [7:0] out;
input enable, clk, reset;
reg [7:0] out;
always @(posedge clk)
if(reset)
begin
 out<=8'b0;
end
else if(enable)
begin
out<=out+1;
end
endmodule
``` | ```c
typedef unsigned char _u8;
struct module_up_counter {
 _u8 out;
 _Bool enable;
 _Bool clk;
 _Bool reset;
};
extern struct
module_up_counter up_counter;
int main()
{
 // Inputs of C program
 _Bool enable;
 _Bool clk;
 _Bool reset;
 unsigned char out;

      // reset the design
 // call to C function
 upcounter(&out, 0, clk, 1);
 // set Verilog inputs
 up_counter.enable = 0;
 up_counter.reset = 1;
 set_inputs();
 next_timeframe();
 assert(up_counter.out
   == sup_counter.out);

 while(1) {
  // Start counting, set
  // enable = 1 and reset = 0
  up_counter.reset = 0;
  up_counter.enable = 1;
  set_inputs();
  next_timeframe();
  upcounter(&out, 1, clk, 0);
  assert(up_counter.out
   == sup_counter.out);
 }
}
``` |

Fig. 5. Example of equivalence checking using HW-CBMC

the `<VERILOG-FILE-NAME>` can be specified as (*.v*) or (*.sv*) file, where (*.v*) is an extension for Verilog files and (*.sv*) is an extension for SystemVerilog files. We assume that the `<MITER-FILE-NAME>` includes the reference model in C and implements the miter. Note that HW-CBMC expects the reference model and the miter implementation to be C programs. The command line switch `--aig` instructs the tool to operate in bit-level mode. Without this option, the default operating mode in HW-CBMC is word-level mode. The switch `--bound` and `--unwind` is used to specify the unwind depth for the hardware and software transition system respectively. The switch `--module` specifies the name of the top level module in the Verilog design file. HW-CBMC also provides an option, `--vcd` to dump counterexamples in Value Change Dump (*vcd*) format in case of assertion failure, which can be analyzed for debugging purpose using waveform viewer, such as `gtkwave`.[9] The switch `--help` shows the available command line options for using HW-CBMC.

---

[9] http://gtkwave.sourceforge.net

```
// generate interface
hw-cbmc <VERILOG-FILE-NAME> --module <TOP-MODULE> --gen-interface
// bit-level mode
hw-cbmc <VERILOG-FILE-NAME> <MITER-FILE-NAME> --module <TOP-MODULE>
--bound N --unwind M --aig --vcd <VCD-FILE-NAME>
// word-level mode
hw-cbmc <VERILOG-FILE-NAME> <MITER-FILE-NAME> --module <TOP-MODULE>
--bound N --unwind M --vcd <VCD-FILE-NAME>
```

Monolithic and Path-wise Approach to Equivalence Checking

We investigated the structure of the ARM FPU and dual-path adder examples discussed the paper to analyze the effect on runtimes of the monolithic and path-based equivalence checking approaches followed by HW-CBMC and VERIFOX respectively.

　　We observe that the pipelined implementation of ARM FPU forces VERIFOX to traverse deep into a particular path and then backtrack to a much higher level in the symbolic tree due to infeasibility of the current path. This causes VERIFOX to throw away several path fragments that were earlier considered feasible while going deep in the path only to be discovered as infeasible much later. This results in the wastage of significant computation time in VERIFOX. On the other hand, the dual-path adder contains a state-machine that implements separate cases for the addition of different types of numbers. This allows VERIFOX to perform an early infeasibility check and prune most of the irrelevant logic upfront in the symbolic execution phase using assumptions. On the other hand, the monolithic constraint generated by HW-CBMC for the dual-path FP adder was extremely difficult to solve. In this way, our experiments give some insight into how the path-based symbolic execution in VERIFOX and the monolithic BMC-based approach in HW-CBMC are sensitive to the structure of the original floating-point design.

Synthesizable Constructs in HW-CBMC

Our Verilog front-end in HW-CBMC support IEEE 1364.1 2005 Verilog standards. This includes the entire synthesizable fragment of Verilog. The detailed list of synthesizable Verilog constructs supported by our Verilog front-end is available in our website www.cprover.org/ebmc/manual/verilog_language_features.shtml.

Miter Construction for Equivalence Checking

Figure 6 shows an example miter for checking equivalence of a 64-bit floating-point adder at the software level and RTL phase using VERIFOX and HW-CBMC respectively.

　　For the miter in VERIFOX, we provide the same floating-point numbers as inputs to the reference design (built inside VERIFOX) and an externally provided untimed SW implementation (in C). We then set the `rounding mode` of the reference model and the SW implementation accordingly. Subsequently, the results of addition from the reference model (*sum_ref*) and the SW implementation (*sum_impl*) are checked for equivalence using the function, `assert(compareFloat(sum_ref, sum_impl));`.

Miter for VERIFOX	Miter for HW-CBMC

```
int miter(float f, float g) {
 roundmode        rmode;
 softfloat_uint64_t nan_payload;
 float            sum_ref,sum_impl;
 int ROUNDMODE;
 switch(ROUNDMODE) {
   case 0: { // ROUND TO NEAREST EVEN
     fesetround(FE_TONEAREST);
     rmode = 3;
     break;
   }
   case 1: { // ROUND UP
     fesetround(FE_UPWARD);
     rmode = 0;
     break;
   }
   case 2: { // ROUND DOWN
     fesetround(FE_DOWNWARD);
     rmode = 1;
     break;
   }
   case 3: { // TOWARDZERO
     fesetround(FE_TOWARDZERO);
     rmode = 2;
     break;
   }
 }

 // Invoke the reference model
 sum_ref = f + g;
 // Invoke the ARM FPU ADD
 nan_payload = 0x00080000;
 sum_impl = sfadd64(f,g,rmode,nan_payload);
 // check the output
 assert(compareFloat(sum_ref,sum_impl));
}
```

```
int miter(float f, float g) {
 float C_result,Verilog_result;
 roundmode         rmode;
 int ROUNDMODE;
 // reset the design
 add64.reset_n = 0;
 set_inputs();next_timeframe();
 // pass the inputs to the RTL
 add64.reset_n = 1;
 add64.src0 = *(unsigned*)&f;add64.src1 = *(unsigned*)&g;
 set_inputs();next_timeframe();
 // settings for RTL floating-point addition
 add64.pipe_ready = 1;add64.valid_in = 1;add64.lane_mask = 3;
 switch(ROUNDMODE) {
   case0 : { // ROUND TO NEAREST EVEN
     add64.round_mode = 0;
     rmode = 3;break; }
   case1 : { // ROUND UP
     add64.round_mode = 1;
     rmode = 0;break; }
   case2 : { // ROUND DOWN
     add64.round_mode = 2;
     rmode = 1;break; }
   case3 : { // TOWARDZERO
     add64.round_mode = 3;
     rmode = 2;break; }
 }
 set_inputs();next_timeframe();
 next_timeframe();next_timeframe();
 next_timeframe();
 // Invoke the ARM FPU ADD
 nan_payload = 0x00080000;
 C_result = sfadd64(f,g,rmode,nan_payload);
 // RTL result must be ready here
 Verilog_result = *(float*)&add64.res;
 // check the output
 assert(compareFloat(C_result,Verilog_result));
}
```

Fig. 6. Miter for equivalence checking of a double precision floating-point adder from ARM

In a similar way, the miter in HW-CBMC is constructed by providing the same floating-point numbers as input to the SW and HW RTL implementations. Note that the inputs are set to the HW signals in HW-CBMC using a function set_inputs(). Since the ARM FPU is a pipeline implementation with pipeline depth 4, we unwind the HW transition system up to a bound of 4 using the function next_timeframe(). Subsequently, the results computed by the HW design and the C reference model are compared using the compareFloat() function.

Miter for Combinational Equivalence Checking in HW-CBMC

Figure 7 shows an example miter for checking combinational equivalence of a 32-bit floating-point adder/subtractor circuit. We provide the same floating-point numbers as inputs to the reference design (in C) and the hardware implementation (in RTL Verilog) using set_inputs(). Subsequently, we indicate that we want to perform a floating-point addition by setting isAdd=1. The results computed by the hardware design and

```
void miter(float f, float g) {
    // setting up the inputs to hardware FPU
    fp_add_sub.f = *(unsigned*)&f;
    fp_add_sub.g = *(unsigned*)&g;
    fp_add_sub.isAdd = 1;
    // propagates inputs of the hardware circuit
    set_inputs();
    // get result from hardware circuit
    float Verilog_result = *(float*)&fp_add_sub.result;
    // compute fp-add in Software with rounding mode RNE
    float C_result = add(RNE, f, g);
    // compare the outputs
    assert(compareFloat(C_result, Verilog_result));
}
```

Fig. 7. Miter for combinational equivalence checking for a 32-bit floating-point adder/subtractor for the case of addition in HW-CBMC

the C reference model are compared using the `compareFloat()` function. Note that this is a combinational circuit, so there is no call to `next_timeframe()`.