

Formal Verification of High-Level Synthesis

Yann Herklotz, James D. Pollard, Nadesh Ramanathan, John Wickerson

Imperial College London



Outline

Example

Verification

The Need to Design Hardware Accelerators

Field-programmable gate arrays (FPGAs) becoming more popular as flexible hardware acceleration.

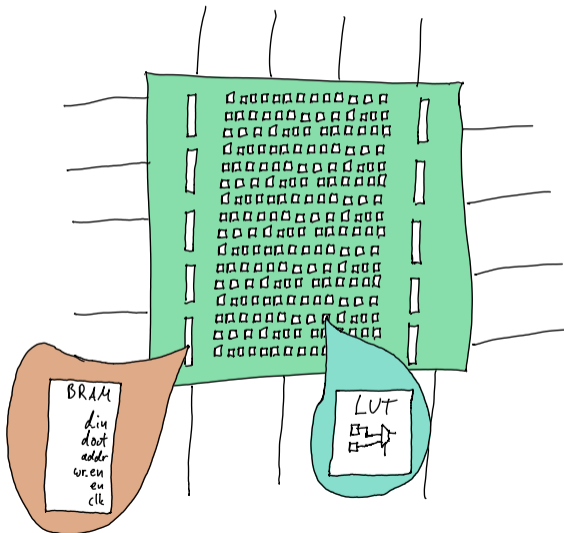
Compared to microcontrollers:

- Can greatly **reduce latency**.
- Lower **power**.
- **Higher performance**.

But:


- Needs knowledge about hardware design.
- **Less flexible**.

FPGA Layout



So How do we Program an FPGA?

- FPGAs contain **LUTs** and programmable interconnects.



So How do we Program an FPGA?


- FPGAs contain **LUTs** and programmable interconnects.
- Programmed using **hardware description languages**.



Fine control



Long to design



So How do we Program an FPGA?


- FPGAs contain **LUTs** and programmable interconnects.
- Programmed using **hardware description languages**.
- Simulation quite slow.



Fine control



Long to design



So How do we Program an FPGA?


- FPGAs contain **LUTs** and programmable interconnects.
- Programmed using **hardware description languages**.
- Simulation quite slow.
- High-Level Synthesis is an alternative.



Quick to design



Less control



So How do we Program an FPGA?


- FPGAs contain **LUTs** and programmable interconnects.
- Programmed using **hardware description languages**.
- Simulation quite slow.
- High-Level Synthesis is an alternative.
- Faster testing through execution.



Quick to design



Less control



Motivation for Formal Verification

Difficult to debug HLS tools:

- Simulation can take a long time.
- Correctness is important in hardware, testing is done at every level.

Motivation for Formal Verification


Difficult to debug HLS tools:

- Simulation can take a long time.
- Correctness is important in hardware, testing is done at every level.

High-level synthesis is often quite unreliable:


- Intel's OpenCL could not be fuzzed because of too many issues (Lidbury et al. [2015]).
- We fuzzed HLS tools and found they failed on **2.5%** of simple random test cases.

Solution



Use CompCert, a fully verified C compiler, and add an HLS backend.

Solution



Current progress: fully verified HLS tool for a subset of C.

Support for: all **control flow**, **fixedpoint**, **non-recursive functions** and **local arrays/structs/unions**.

Outline

Example

Verification

Example: RTL

```
int main() {  
    int x[2] = {3, 6};  
    int i = 1;  
    return x[i];  
}
```

Example of a very simple program performing loads and stores.

Example: RTL

- **three address code (RTL)**
instructions are represented as a control-flow graph (CFG).
- Each instruction links to the next one.

```
main() {  
    x5 = 3  
    int32[stack(0)] = x5  
    x4 = 6  
    int32[stack(4)] = x4  
    x1 = 1  
    x3 = stack(0) (int)  
    x2 = int32[x3 + x1 * 4 + 0]  
    return x2  
}
```


Example: HTL Overview

The representation of the **finite state-machine with datapath (FSMD)** is abstract and called **HTL**.

Definition datapath := PTree.t Verilog.stmnt.

Definition controllogic := PTree.t Verilog.stmnt.

Example: HTL Overview

The representation of the **finite state-machine with datapath (FSMD)** is abstract and called **HTL**.

Definition datapath := PTree.t Verilog.stmnt.

Definition controllogic := PTree.t Verilog.stmnt.

```
Record module: Type := mkmodule {  
  mod_datapath: datapath; mod_controllogic: controllogic;  
  mod_wf: map_well_formed mod_controllogic  
    /\ map_well_formed mod_datapath;  
  mod_reset: reg;  
  mod_ram: ram_spec;  
  ...  
}.
```

Example: Translation (RTL \rightarrow HTL)

Translation from **control-flow graph (CFG)** into a **finite state-machine with datapath (FSMD)**.

Example: Translation (RTL \rightarrow HTL)

Translation from **control-flow graph (CFG)** into a **finite state-machine with datapath (FSMD)**.

- **Control-flow** is translated into a **finite state-machine**.

Example: Translation (RTL \rightarrow HTL)

Translation from **control-flow graph (CFG)** into a **finite state-machine with datapath (FSMD)**.

- **Control-flow** is translated into a **finite state-machine**.
- Each **RTL instructions** translated into equivalent **Verilog statements**.

$x3 = x3 + x5 + 0 \quad \longrightarrow \quad \text{reg_3} \leq \{\text{reg_3} + \{\text{reg_5} + 32'd0\}\}$

Example: Translation (RTL \rightarrow HTL)

Translation from **control-flow graph (CFG)** into a **finite state-machine with datapath (FSMD)**.

- **Control-flow** is translated into a **finite state-machine**.
- Each **RTL instructions** translated into equivalent **Verilog statements**.
- Function **stack** implemented as **RAM**.

Example: Translation (RTL \rightarrow HTL)

Translation from **control-flow graph (CFG)** into a **finite state-machine with datapath (FSMD)**.

- **Control-flow** is translated into a **finite state-machine**.
- Each **RTL instructions** translated into equivalent **Verilog statements**.
- Function **stack** implemented as **RAM**.
- Pointers for loads and stores translated to RAM addresses.

$$\begin{aligned} &x5 + x1 * 4 + 0 \\ \longrightarrow & \{ \{ \{ \text{reg}_5 + 32'd0 \} + \{ \text{reg}_1 * 32'd4 \} \} / 32'd4 \} \end{aligned}$$

Example: Translation (RTL \rightarrow HTL)

Translation from **control-flow graph (CFG)** into a **finite state-machine with datapath (FSMD)**.

- **Control-flow** is translated into a **finite state-machine**.
- Each **RTL instructions** translated into equivalent **Verilog statements**.
- Function **stack** implemented as **RAM**.
- Pointers for loads and stores translated to RAM addresses.
 - **Byte** addressed to **word** addressed.

$$\begin{aligned} & x5 + x1 * 4 + 0 \\ \longrightarrow & \{ \{ \{ \text{reg}_5 + 32'd0 \} + \{ \text{reg}_1 * 32'd4 \} \} / 32'd4 \} \end{aligned}$$

Example: Translation (HTL \rightarrow Verilog)

```
module main(reset, clk, finish, return_val);
  input [0:0] reset, clk;
  output reg [0:0] finish = 0;
  output reg [31:0] return_val = 0;
  reg [31:0] reg_3 = 0, addr = 0, d_in = 0,
            reg_5 = 0, wr_en = 0,
            state = 0, reg_2 = 0,
            reg_4 = 0, d_out = 0, reg_1 = 0;
  reg [0:0] en = 0, u_en = 0;
  reg [31:0] stack [1:0];
  // RAM interface
  always @(negedge clk)
    if ({u_en != en}) begin
      if (wr_en) stack[addr] <= d_in;
      else d_out <= stack[addr];
      en <= u_en;
    end
end
```

- Finally, translate the FSM into Verilog.

Example: Translation (HTL \rightarrow Verilog)

```
module main(reset, clk, finish, return_val);
  input [0:0] reset, clk;
  output reg [0:0] finish = 0;
  output reg [31:0] return_val = 0;
  reg [31:0] reg_3 = 0, addr = 0, d_in = 0,
            reg_5 = 0, wr_en = 0,
            state = 0, reg_2 = 0,
            reg_4 = 0, d_out = 0, reg_1 = 0;
  reg [0:0] en = 0, u_en = 0;
  reg [31:0] stack [1:0];
  // RAM interface
  always @(negedge clk)
    if ({u_en != en}) begin
      if (wr_en) stack[addr] <= d_in;
      else d_out <= stack[addr];
      en <= u_en;
    end
end
```

- Finally, translate the FSM into Verilog.
- This includes a RAM interface.

Example: Translation (HTL → Verilog)

```
// Data-path
always @(posedge clk)
  case (state)
    32'd11: reg_2 <= d_out;
    32'd8: reg_5 <= 32'd3;
    32'd7: begin
      u_en <= (~ u_en); wr_en <= 32'd1;
      d_in <= reg_5; addr <= 32'd0;
    end
    32'd6: reg_4 <= 32'd6;
    32'd5: begin
      u_en <= (~ u_en); wr_en <= 32'd1;
      d_in <= reg_4; addr <= 32'd1;
    end
    32'd4: reg_1 <= 32'd1;
    32'd3: reg_3 <= 32'd0;
    32'd2: begin
      u_en <= (~ u_en); wr_en <= 32'd0;
      addr <= {{{reg_3 + 32'd0} + {reg_1 * 32'd4}} / 32'd4};
    end
    32'd1: begin finish = 32'd1; return_val = reg_2; end
    default: ;
  endcase
```

- Finally, translate the FSM into Verilog.
- This includes a RAM interface.
- Data path is translated into a case statement.

Example: Translation (HTL → Verilog)

```
// Data-path
always @(posedge clk)
  case (state)
    32'd11: reg_2 <= d_out;
    32'd8: reg_5 <= 32'd3;
    32'd7: begin
      u_en <= ( ~ u_en); wr_en <= 32'd1;
      d_in <= reg_5; addr <= 32'd0;
    end
    32'd6: reg_4 <= 32'd6;
    32'd5: begin
      u_en <= ( ~ u_en); wr_en <= 32'd1;
      d_in <= reg_4; addr <= 32'd1;
    end
    32'd4: reg_1 <= 32'd1;
    32'd3: reg_3 <= 32'd0;
    32'd2: begin
      u_en <= ( ~ u_en); wr_en <= 32'd0;
      addr <= {{{reg_3 + 32'd0} + {reg_1 * 32'd4}} / 32'd4};
    end
    32'd1: begin finish = 32'd1; return_val = reg_2; end
    default: ;
  endcase
```

- Finally, translate the FSM into Verilog.
- This includes a RAM interface.
- Data path is translated into a case statement.
- Ram loads and stores automatically turn off RAM.

Example: Translation (HTL → Verilog)

```
// Control logic
always @(posedge clk)
  if ({reset == 32'd1}) state <= 32'd8;
  else case (state)
    32'd11: state <= 32'd1;      32'd4: state <= 32'd3;
    32'd8: state <= 32'd7;      32'd3: state <= 32'd2;
    32'd7: state <= 32'd6;      32'd2: state <= 32'd11;
    32'd6: state <= 32'd5;      32'd1: ;
    32'd5: state <= 32'd4;      default: ;
  endcase
endmodule
```

- Finally, translate the FSM into Verilog.
- This includes a RAM interface.
- Data path is translated into a case statement.
- Ram loads and stores automatically turn off RAM.
- Control logic is translated into another case statement with a reset.

Outline

Example

Verification

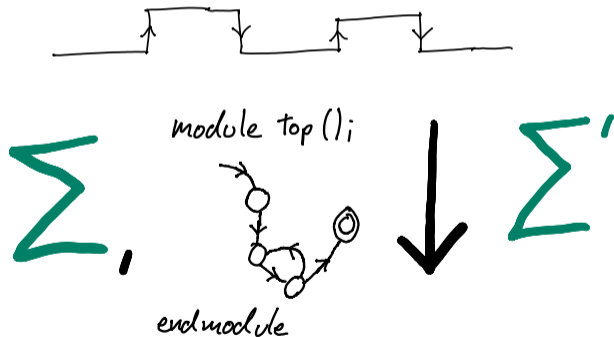
Verilog Semantics (Adapted from Lööw et al. (2019))

- Top-level semantics are **small-step operational semantics**.



Verilog Semantics (Adapted from Löow et al. (2019))

- Top-level semantics are **small-step operational semantics**.
- Within each small step, a Verilog module is executed in one big step.



How do we prove the HLS tool correct?

- We have an **algorithm** describing the **translation**.
- Have to **prove** that it does not change **behaviour** with respect to our language semantics.

How do we prove the HLS tool correct?

- We have an **algorithm** describing the **translation**.
- Have to **prove** that it does not change **behaviour** with respect to our language semantics.

Behaviour	Guarantee
Converging	Means a result is obtained, Verilog and C results must be equal.
Diverging	C is in an infinite loop, Verilog must execute indefinitely.
Wrong	Such as undefined behaviour, no guarantees need to be shown.

How do we prove the HLS tool correct?

- We have an **algorithm** describing the **translation**.
- Have to **prove** that it does not change **behaviour** with respect to our language semantics.


Theorem (Main Backward Simulation)

$$\forall C, V, B, \quad \mathbf{HLS}(C) = \mathbf{OK}(V) \wedge \mathbf{Safe}(C) \implies (V \Downarrow B \implies C \Downarrow B).$$


where

$$\mathbf{Safe}(C) : \forall B, C \Downarrow B \implies B \in \mathbf{Safe}$$

With Division approximately $27\times$ slower



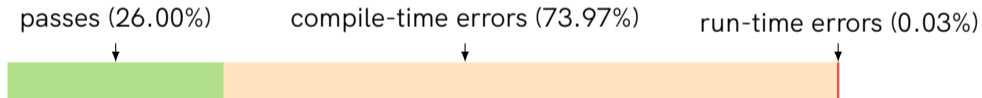
Without Division about $2\times$ slower



Fuzzing Vericert with Csmith

Fuzzed Vericert with Csmith to check correctness theorem.

- One bug was found in the pretty printing.
- Many compile-time errors are expected.
- Mainly rejected because of wrong size.



Conclusion

Written a formally verified high-level synthesis tool in **Coq** based on **CompCert**.

- Base translation **proven correct** by proving translation of CFG into FSMD.

Conclusion

Written a formally verified high-level synthesis tool in **Coq** based on **CompCert**.

- Base translation **proven correct** by proving translation of CFG into FSM.
- Small optimisations implemented such as **Ram Inference**.

Conclusion

Written a formally verified high-level synthesis tool in **Coq** based on **CompCert**.

- Base translation **proven correct** by proving translation of CFG into FSMD.
- Small optimisations implemented such as **Ram Inference**.
- Performance without divisions comparable to LegUp without optimisations.

Thank you

Documentation



<https://vericert.ymhg.org>

GitHub



<https://github.com/ymherklotz/vericert>

OOPSLA'21 Preprint



https://ymhg.org/papers/fvhls_oopsla21.pdf

References

Christopher Lidbury, Andrei Lascu, Nathan Chong, and Alastair F. Donaldson. Many-core compiler fuzzing. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '15*, pages 65–76, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450334686. doi: 10.1145/2737924.2737986. URL <https://doi.org/10.1145/2737924.2737986>.