Porting and Optimizing CompCert for an interlocked VLIW processor

Cyril Six (CIFRE PhD student)
Supervised by:
Sylvain Boulmé (Verimag PACSS)
Benoît Dupont de Dinechin (Kalray)
David Monniaux (Verimag PACSS)

cyril.six@univ-grenoble-alpes.fr
sylvain.boulme@univ-grenoble-alpes.fr
benoit.dinechin@kalray.eu
david.monniaux@univ-grenoble-alpes.fr

November 25th, 2019

About us

- PACSS team - Proofs and Code analysis for Safety and Security
  - David Monniaux, CNRS senior researcher
  - Sylvain Boulmé, researcher
- Kalray - Fabless semiconductor company based in Grenoble
  - Benoît Dupont de Dinechin, CTO

Presentation plan

1. Introduction
   - VLIW in-order processors
   - Coq and CompCert architecture
2. Our work
   - Formal blockstep semantics for VLIW in CompCert
   - Certified intrablock postpass scheduling
3. Results
   - Experimentations
   - Future and ongoing work

Outline

1. Introduction
   - VLIW in-order processors
   - Coq and CompCert architecture
2. Our work
3. Results

Kalray processor

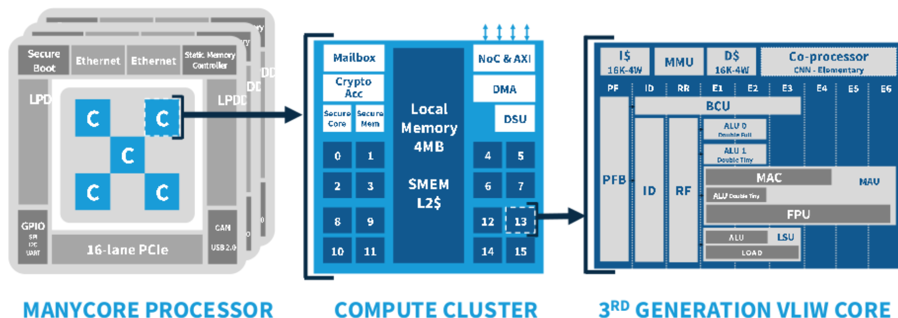

- 1 processor = 5 compute clusters
- 1 compute clusters = 16 cores @ (600MHz - 1.2GHz)
- Network on Chip allowing point-to-point communication
- Main DDR memory of 4 GB

Kalray VLIW k1c core

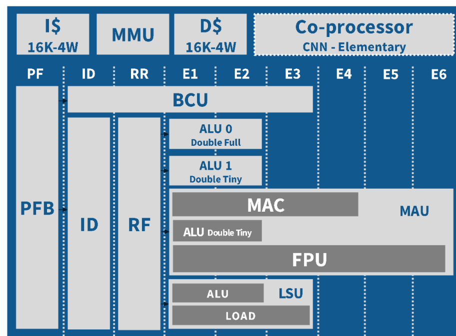

- ALU: Arithmetic-Logic Unit (x2)
- LSU: Load-Store Unit
- MAU: Multiply-Accumulate Unit
- BCU: Branch Control Unit

- 64x64-bit user registers per core
- Very Large Instruction Word (VLIW): explicit Instruction Level Parallelism
- 5 execution units: ALU0, ALU1, LSU, MAU, BCU
- In-order, pipelined execution

Example of k1c code

```
addw $r2 = $r1, $r0      /* ALU */
;;                       /* bundle delimitor */
mulu $r2 = $r2, 2
addw $r0 = $r2, $r1      /* MAU + ALU */
;;
addw $r0 = $r1, 0
addw $r1 = $r0, 0     /* ALU + ALU */
;;
mulu $r1 = $r1, 2
addw $r3 = $r2, 42
j toto                   /* BCU + ALU + MAU */
```

- Bundles are explicitly delimited by the programmer/compiler
- In-order execution
- Parallelism inside each bundle

In-order vs out-of-order

- More predictible, more precise computation of Worst Case Execution Time (WCET)
- Simpler control structure
  - Uses less CPU die space and energy
  - May be more reliable (less complex design)
- => Good for safety-critical applications

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

13 cycles

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				

13 cycles

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				
2	lwz_{r4}^{r0}	lwz_{r1}^{r0}			

13 cycles

K1c pipeline

```c
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				
2	lwz_{r4}^{r0}	lwz_{r1}^{r0}			
3	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}		

13 cycles

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				
2	lwz_{r4}^{r0}	lwz_{r1}^{r0}			
3	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}		
4	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}	

13 cycles

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				
2	lwz_{r4}^{r0}	lwz_{r1}^{r0}			
3	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}		
4	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}	
5	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	*STALL*	lwz_{r4}^{r0}	lwz_{r1}^{r0}

13 cycles

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				
2	lwz_{r4}^{r0}	lwz_{r1}^{r0}			
3	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}		
4	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}	
5	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	STALL	lwz_{r4}^{r0}	lwz_{r1}^{r0}
6	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	STALL	STALL	lwz_{r4}^{r0}

13 cycles

K1c pipeline

```c
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

Cycle	Issue	Read Regs	E1	E2	E3
1	lwz_{r1}^{r0}				
2	lwz_{r4}^{r0}	lwz_{r1}^{r0}			
3	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}		
4	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	lwz_{r4}^{r0}	lwz_{r1}^{r0}	
5	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	STALL	lwz_{r4}^{r0}	lwz_{r1}^{r0}
6	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	STALL	STALL	lwz_{r4}^{r0}
7	$add_{r0}^{r1,r3}$	lwz_{r3}^{r0}	$add_{r1}^{r1,r4}$	STALL	STALL

13 cycles

K1c pipeline

```
int add4(int *t){
  return (t[0] + t[1] + t[2] + t[3]);
}
```

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    /* 2 cycles stall ($r4) */
    addw $r1 = $r1, $r4
    ;;
    lwz $r3 = 8[$r0]
    ;;
    /* 2 cycles stall ($r3) */
    addw $r0 = $r1, $r3
    ;;
    lwz $r2 = 12[$r0]
    ;;
    /* 2 cycles stall ($r2) */
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

13 cycles

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    lwz $r3 = 8[$r0]
    ;;
    lwz $r2 = 12[$r0]
    ;;
    addw $r1 = $r1, $r4
    ;;
    addw $r0 = $r1, $r3
    ;;
    addw $r0 = $r0, $r2
    ;;
    ret
    ;;
```

8 cycles

Compilers and VLIW processors

- VLIW compilers should be able to generate bundles to have a good performance
- Instructions should also be reordered to minimize the latencies
- => This is usually done by a scheduling pass, after register allocation

Code to schedule..

```
add4:                          # Sched time
    lwz $r1 = 0[$r0]           #    0
    lwz $r4 = 4[$r0]           #    1
    addw $r1 = $r1, $r4        #    3
    lwz $r3 = 8[$r0]           #    2
    addw $r0 = $r1, $r3        #    5
    lwz $r2 = 12[$r0]          #    3
    addw $r0 = $r0, $r2        #    6
    ret                        #    6
```

After scheduling

```
add4:
    lwz $r1 = 0[$r0]
    ;;
    lwz $r4 = 4[$r0]
    ;;
    lwz $r3 = 8[$r0]
    ;;
    lwz $r2 = 12[$r0]
    addw $r1 = $r1, $r4
    ;;
    addw $r0 = $r1, $r3
    ;;
    addw $r0 = $r0, $r2
    ret
    ;;
```

Register allocation and scheduling

- Register allocation
  - Allocates physical registers (bounded) to virtual registers (unbounded)
  - Performs "spilling" if not able to
- Instruction scheduling
  - After register allocation (postpass): more precise informations, can make bundles, but extra dependencies on registers
  - To deal with the register dependencies: instruction scheduling before register allocation (prepass)
- We present here the postpass scheduling optimization

Outline

1. Introduction
   - VLIW in-order processors
   - Coq and CompCert architecture
2. Our work
3. Results

Coq in a few words

- Programming language with a proof assistant
- Idea: you write programs in the Coq functional language..

```
Fixpoint sum_list (l: nat list) : nat :=
  match l with
  | nil -> 0
  | e::l -> e + sum_list l
  end.
```

- .. and then you prove them!

```
Theorem sum_list_distributive:
  forall l l', sum_list (l++l') = sum_list l + sum_list l'.
Proof.
  (* ... *)
Qed.
```

- Advantage: the proof is *machine checked*, so less prone to human errors

CompCert in a few words

- Machine checked, formally verified compiler: proof of specification preservation
- Written in Coq and OCaml
- Targets: PPC, ARM, RISC-V, x86
- Performance: close to GCC -O1

CompCert architecture

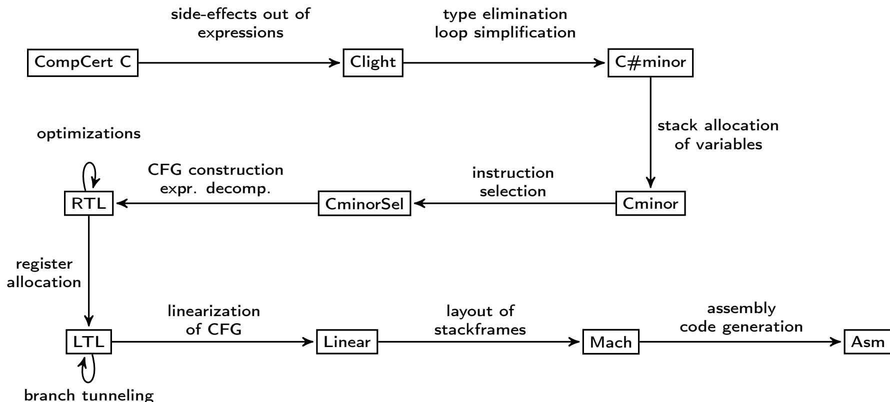

- What we want: intrablock postpass scheduling at Asm level
- Drawing on board

Our modifications to the CompCert architecture

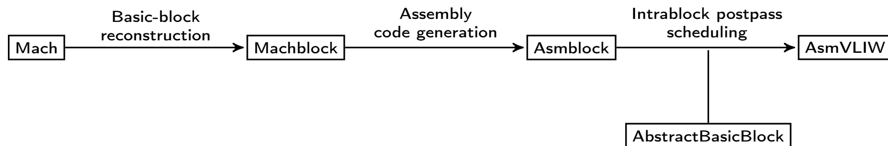

- Machblock and Asmblock: one block = one basic block, sequential semantics
- AsmVLIW: one block = one bundle, parallel semantics within a bundle

Outline

1 Introduction

2 Our work
- Formal blockstep semantics for VLIW in CompCert
- Certified intrablock postpass scheduling

3 Results

Semantics of Asm

State: (rs, m)

- Registers state *rs*: mapping from registers (`PC`, `RA`, r1, r2, ..) to values
- Memory state *m*: mapping from addresses to values

- exec_instr: function −> instruction −> regset −> mem −> outcome
- outcome: either Stuck, or Next rs' m'
- Instructions reside in memory, and are pointed by `PC` register
- Exemples of execution:
  - (exec_instr f (Pcall s) rs m) returns $(rs[\texttt{RA} \leftarrow rs[\texttt{PC}];\ \texttt{PC} \leftarrow \texttt{@}s], m)$
  - (exec_instr f (Paddd r0 r1 r2) rs m) returns $(rs[r_0 \leftarrow rs[r_1] + rs[r_2]], m)$

Formal definition of a basic block

```
Inductive basic: Type :=    (* basic instructions *)
Inductive control: Type := (* control-flow instructions *)
Record bblock := {
  header: list label;  body: list basic;  exit: option control;
  correct: wf_bblock body exit  (* must contain at least 1 instr. *)
}
```

- Exemples: Pcall is a *control*, Paddd is a *basic*.
- We use a bblock for basic blocks (sequential) and bundles (parallel) alike
- Executing the bblock ($R_0 := R_1; R_1 := R_0;$ `jump @toto`) in parallel should lead to $rs[R_0 \leftarrow rs[R_1]; R_1 \leftarrow rs[R_0]; PC \leftarrow @toto]$

Parallel deterministic in-order semantics

- Idea: have a memorized state for the reads, and a running state for the writes
- Instead of (rs, m), we use four components in a bundle execution:
  - rsr, mr: the regset and the memory state, prior to executing the bundle
  - rsw, mw: the running state, where all the writes will occur (in order)

Parallel deterministic in-order semantics

- Idea: have a memorized state for the reads, and a running state for the writes
- Instead of (rs, m), we use four components in a bundle execution:
  - rsr, mr: the regset and the memory state, prior to executing the bundle
  - rsw, mw: the running state, where all the writes will occur (in order)

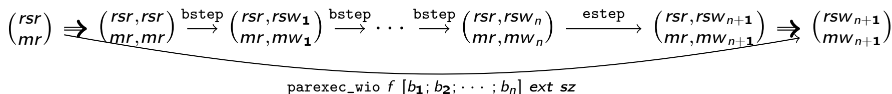

Parallel deterministic in-order semantics

- Idea: have a memorized state for the reads, and a running state for the writes
- Instead of (rs, m), we use four components in a bundle execution:
  - rsr, mr: the regset and the memory state, prior to executing the bundle
  - rsw, mw: the running state, where all the writes will occur (in order)

$$\binom{rsr}{mr} \Rightarrow \binom{rsr, rsr}{mr, mr} \xrightarrow{\texttt{bstep}} \binom{rsr, rsw_1}{mr, mw_1} \xrightarrow{\texttt{bstep}} \cdots \xrightarrow{\texttt{bstep}} \binom{rsr, rsw_n}{mr, mw_n} \xrightarrow{\texttt{estep}} \binom{rsr, rsw_{n+1}}{mr, mw_{n+1}} \Rightarrow \binom{rsw_{n+1}}{mw_{n+1}}$$

$\texttt{parexec_wio}\ f\ [b_1; b_2; \cdots; b_n]\ ext\ sz$

- This first version models an atomic parallel execution, where:
  - All reads are done in parallel, prior to any write
  - The writes are done sequentially in the same order
- Example: $(R_0 := R_1; R_1 := R_0; \texttt{jump @toto})$
  - $rsw_1 = rsr[R_0 \leftarrow rsr[R_1]] = rs[R_0 \leftarrow r_1]$
  - $rsw_2 = rsw_1[R_1 \leftarrow rsr[R_0]] = rs[R_0 \leftarrow r_1; R_1 \leftarrow r_0]$
  - $rs' = rsw_3 = rsw_2[\mathrm{PC} \leftarrow @toto] = rs[R_0 \leftarrow r_1; R_1 \leftarrow r_0; \mathrm{PC} \leftarrow @toto]$

Extending the semantic

- Issue with the semantic: does not model the possibility of concurrent writes

Extending the semantic

- Issue with the semantic: does not model the possibility of concurrent writes
- Solution: say in Coq "Executing bundle b with initial states (rs, m) gives outcome o iff there exists a permutation of writes of b that gives o by the previous semantic (deterministic in-order)"

Extending the semantic

- Issue with the semantic: does not model the possibility of concurrent writes
- Solution: say in Coq "Executing bundle b with initial states (rs, m) gives outcome o iff there exists a permutation of writes of b that gives o by the previous semantic (deterministic in-order)"
- To determinize it: only accept outcomes if they are unique

So far..

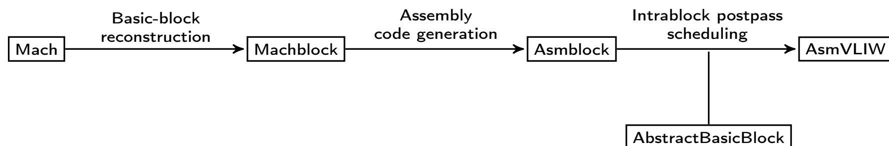

- How to reorder a block from Asmblock into several bundles of AsmVLIW, and prove it correct?

Outline

1 Introduction

2 Our work
- Formal blockstep semantics for VLIW in CompCert
- Certified intrablock postpass scheduling

3 Results

Forward simulations in CompCert

- Simulation diagrams are used to prove semantic preservation
- For this transformation, we use the *Lock-step* and the *Plus* simulations

Lock-step simulation

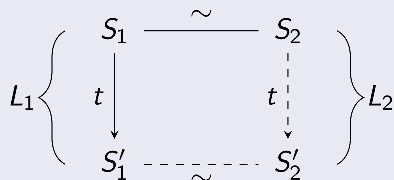

"Plus" simulation

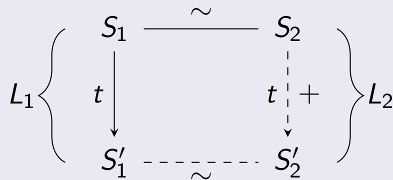

- Lock: If $S_1 \xrightarrow[step_{L_1}]{t} S'_1$ and $S_1 \sim S_2$, then $\exists S'_2, S_2 \xrightarrow[step_{L_2}]{t} S'_2$ and $S'_1 \sim S'_2$.

Previous work of J.B. Tristan on scheduling

- Gallium PhD 2009 by J.B. Tristan, *Formal verification of translation validators*

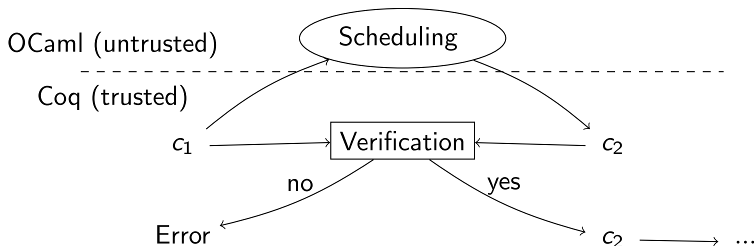

- Verification proof: $V(c_1, c_2) = true \implies c_1 \sim c_2$
- Advantages: easier to prove and modular
- Implemented it at the Mach level, with symbolic evaluation
- Drawbacks: scalability issue + Mach level

Architecture of the pass Asmblock -> AsmVLIW

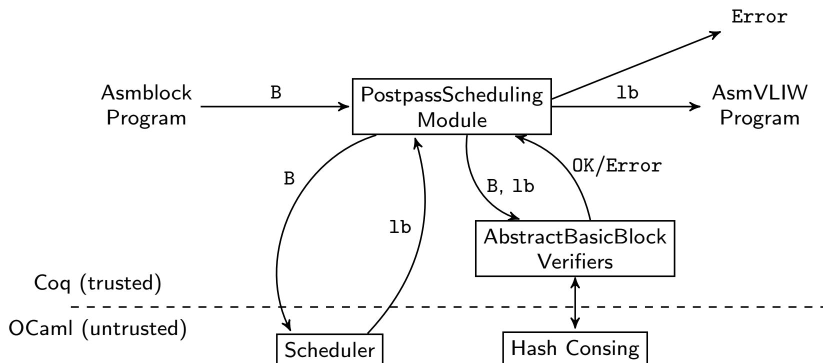

- Axiom schedule: bblock −> list bblock.
- Forward simulation in two parts:
  - Proving reordering, in sequential semantics: *plus simulation*
  - Proving for each bundle that parallel execution = sequential execution: *lockstep simulation*
- (Drawing on board)

AbstractBasicBlock

- New language called AbstractBasicBlock to abstract the details of instructions
- Put “This instruction writes in that register, and reads in this and that registers” in a canonical form

```
Inductive exp :=
  | Read (x:R.t) | Op (o:op) (le: list_exp) | Old (e: exp)
with list_exp :=
  | Enil | Econs (e:exp) (le:list_exp) | LOld (le: list_exp).

Definition inst := list (R.t * exp). (* list of assignments *)
```

- The assignments can be executed sequentially or in parallel
- Examples of traductions:
  - Picall r => $[\#RA \leftarrow Read(\#PC);\ \texttt{PC} \leftarrow Read(\#r)]$
  - Paddw r0 r1 r2 => $[\#r_0 \leftarrow (OpAddw[\#r_1;\ \#r_2])]$
- We use it to execute symbolically the instructions

Intrablock Reordering verifier of AbstractBasicBlock

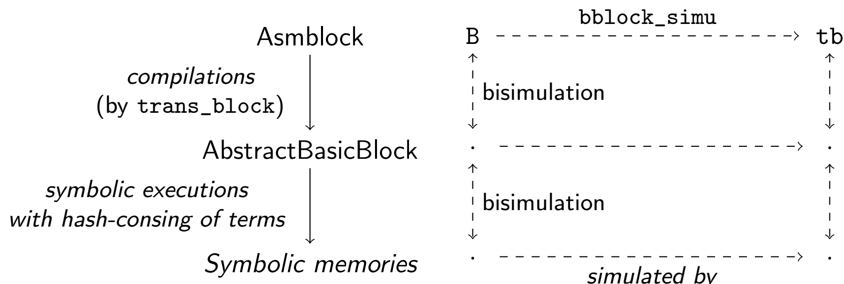

- Symbolic execution: computing symbolic memories, final value of the pseudo register in function of the initial values.
- Example:
  - $B_1 = [r_1 := r_1 + r_2;\ r_3 := load[r_2, m];\ r_1 := r_1 + r_3]$
  - $B_2 = [r_3 := load[r_2, m];\ r_1 := r_1 + r_2;\ r_1 := r_1 + r_3]$
- These two blocks are equivalent to this assignment:
  $[r_1 \leftarrow (r_1 + r_2) + load[r_2, m] \mid\mid r_3 \leftarrow load[r_2, m]]$

The untrusted scheduler in a few words

- We want to assign a scheduling date to each instruction
- This schedule should satisfy two constraints:
  - Latency constraints: we must respect the dependencies of each instruction
  - Resource constraints: a bundle must not consume more resources than available
  - (drawing on board)
- We implemented several schedulers:
  - Naive greedy one, just packs instructions together (linear time)
  - Variant of Coffman-Graham list scheduler, with critical path heuristic (linear time)
  - Optimal list scheduler by Integer Linear Programming + branch and bound (not linear)
- Experimentally, we barely get better results with ILP than with list scheduling

Outline

1 Introduction

2 Our work

3 Results
- Experimentations
- Future and ongoing work

Experimental time of the oracle and verifier

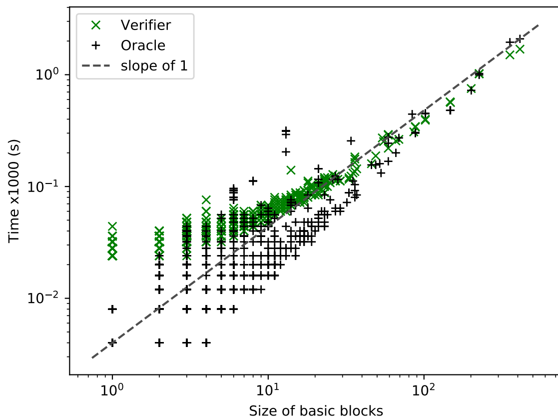

- Timings obtained by instrumenting the OCaml code

Impact of the optimizations on our backend

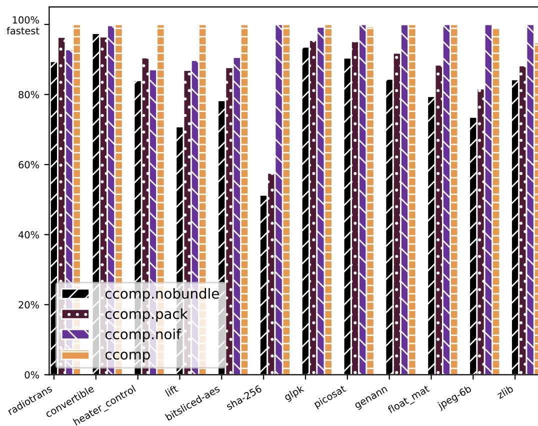

- Impact of scheduling: average gain of 25%, up to 100%

Comparison with GCC

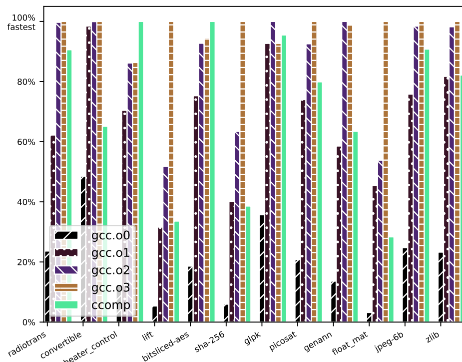

- Results to take with a pinch of salt: GCC backend still in development
- Always better than -O0: 2 to 17 times better
- For most benchmarks, faster than -O1 by 20%. Sometimes better than -O2 and -O3.
- For most others, between 20% and 30% slower than GCC -O3

Optimizations that GCC do compared to us

- Certain strength reductions (replacing multiplication in loop by addition)
- Code motion across basic blocks (e.g. loop invariant code motion)
- Loop unrolling and other loop optimizations
- Prepass scheduling
- Memory alias analysis

Outline

1 Introduction

2 Our work

3 Results
- Experimentations
- Future and ongoing work

Towards prepass scheduling

- WIP: Prepass superblock scheduling in RTL
  - Blocks with one entry point, several exits
- Pass of instruction duplication (tail duplication)
  - Verifier done and proved!
  - Heuristics to implement
- Integration of non-trapping loads
  - Almost done

Other stuff

- Integrate memory alias analysis in the checker
    - Instead of viewing memory as a single pseudo register, have something more elaborate
- Other optimizations..
    - Loop invariant code motion?
    - Loop unrolling?

Thanks for your attention

Do you have any questions?

Parallel in-order semantics (2)

```
Fixpoint parexec_wio_body bdy rsr rsw mr mw : outcome :=
  match bdy with
  | nil => Next rsw mw
  | bi::bdy' => NEXT rsw', mw' <- bstep bi rsr rsw mr mw
                IN parexec_wio_body bdy' rsr rsw' mr mw'
  end.

Definition parexec_wio f bdy ext sz rs m :=
  NEXT rsw', mw' <- parexec_wio_body bdy rs rs m m
  IN estep f ext sz rs rsw' mw'.
```

Non deterministic parallel semantics

- (parexec_bblock f b rs m o) holds if there exists a permutation of writes that gives o by a deterministic in-order execution
- We can reason on permutations of instructions instead of permutations of writes

```
Definition parexec_bblock f b rs m o: Prop :=
exists bdy1 bdy2, Permutation (bdy1 ++ bdy2) b.(body)
 /\ o=(NEXT rsw', mw' <- parexec_wio f bdy1 b.(exit)
                                    (Ptrofs.repr (size b)) rs m
      IN parexec_wio_body bdy2 rs rsw' m mw').
```

- We would like to determinize it, to use in CompCert

Deterministic out-of-order parallel semantics (2)

- (det_parexec f b rs m rs' m') holds if: (rs', m') is the unique outcome of the non-deterministic parallel execution

```
Definition det_parexec f b rs m rs' m': Prop :=
  forall o, parexec_bblock f b rs m o -> o = Next rs' m'.
```

- Remark: in this semantic, Stuck executions cannot happen

AbstractBasicBlock (2)

- Translation from AsmVLIW to AbstractBasicBlock:
  trans_block: bblock -> (list inst)
- We prove a bisimulation for sequential, and a bisimulation for parallel semantics

Bisimulation for sequential:

```
  match_states (State rs m) s ->
  match_outcome (exec_bblock ge fn b rs m)
                (exec Ge (trans_block b) s).
```

Bisimulation for parallel:

```
  match_states (State rs1 m1) s1' ->
  parexec_bblock ge fn b rs1 m1 o2 ->
exists o2', prun Ge (trans_block b) s1' o2'
        /\ match_outcome o2 o2'.
```

Parallelizability checker through AbstractBasicBlock

- We translate the bundle to a block of AbstractBasicBlock
- We prove the following theorem with the sequential bisimulation + parallel bisimulation + correctness of `is_parallelizable` + other minor lemmas:

```
bblock_para_check bundle = true ->
exec_bblock ge f bundle rs m = Next rs' m' ->
det_parexec ge f bundle rs m rs' m'.
```

Towards proving reordering

- Definition of bblock simulation:

```
Definition bblock_simu ge f b b' :=
  forall rs m, exec_bblock ge f b rs m <> Stuck ->
  exec_bblock ge f b rs m = exec_bblock ge f b' rs m.
```

- Definition of a concatenation function, and a predicate (is_concat b lb), where lb is a list bblock
- Definition of a function (verified_schedule b) that:
  - Calls the oracle, retrieving a list of bundles
  - Concatenates together the bundles to form a bblock B
  - Calls the reordering verifier from AbstractBasicBlock (detailed later)
- We then prove the following property:

```
Theorem verified_schedule_correct: forall ge f B lb,
  (verified_schedule B) = (OK lb) ->
  exists tb, is_concat tb lb /\ bblock_simu ge f B tb.
```

Verified hash consing

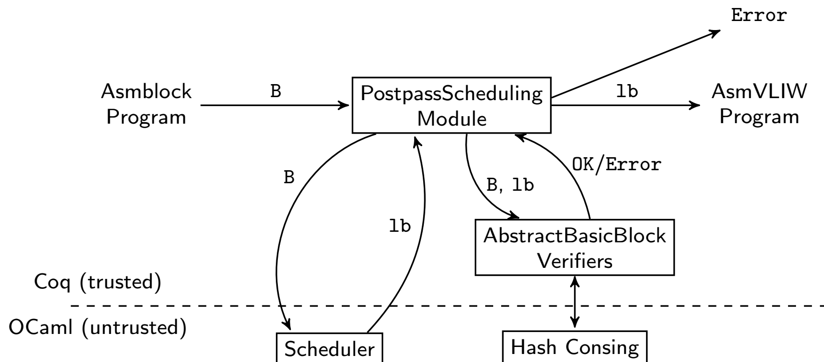

- Memoization involves calling an untrusted OCaml oracle to give memoized terms out of terms
- Dynamic checker in Coq that ensures the memoized term and the term have the same evaluation function
- Check done with OCaml pointer equality
  - Axiom: if pointer equality returns true, then the two values are structurally equals