

# Porting IsaSAT to LLVM (and Inprocessing)

Mathias Fleury and Peter Lammich

CADE'29



# Introduction



# How do We Make SAT Solvers Correct?

## DRAT Proofs

- requires to check the proof for each file
- not all techniques can be represented by current proof formats

# How do We Make SAT Solvers Correct?

## DRAT Proofs

- requires to check the proof for each file
- not all techniques can be represented by current proof formats

## Program Verification This talk!

- works for every input, so *no overhead*
- does not crash even if program runs for a year

## Why do We Care?

- We have abstract CDCL, we have inprocessing, but can we combine them?
- Corner cases easy to miss known papers have flaws!
- Metatheory, study variants See session 1

# Why do We Care?

**Definition 4.2.2** (Clause Redundancy). *A witness labelled clause  $(\omega : C)$  is redundant with respect to a formula  $F$  if  $\omega(C) = \top$  and  $F|_{\alpha} \models F|_{\omega}$  for  $\alpha = \neg C$ . This is also denoted as  $F \wedge C \equiv_{sat}^{\omega} F$ .*

We formalize that part of the proof and extend it to *partial* truth assignments,

CaDiCaL [37]. Rule **WEAKEN**<sup>+</sup> is defined in our calculus based on the most general redundancy property and so it allows to employ every clause elimination procedure implemented in CaDiCaL including variable elimination [86], vivifica-

# Why do We Care?

**Definition 4.2.2** (Clause Redundancy). *A witness labelled clause  $(\omega : C)$  is redundant with respect to a formula  $F$  if  $\omega(C) = \top$  and  $F|_{\alpha} \models F|_{\omega}$  for  $\alpha = \neg C$ . This is also denoted as  $F \wedge C \equiv_{sat}^{\omega} F$ .*

We formalize that part of the proof and extend it to *partial* truth assignments,

CaDiCaL [37]. Rule **WEAKEN**<sup>+</sup> is defined in our calculus based on the most general redundancy property and so it allows to employ every clause elimination procedure implemented in CaDiCaL including variable elimination [86], vivifica-

Either partial models or VE (Master thesis by Katharina Wagner)

I did not realize that either before Isabelle refused a proof  
Implementation heavily tested... on total modals

# Why do We Care?

**Definition 4.2.2** (Clause Redundancy). *A witness labelled clause  $(\omega : C)$  is redundant with respect to a formula  $F$  if  $\omega(C) = \top$  and  $F|_{\alpha} \models F|_{\omega}$  for  $\alpha = \neg C$ . This is also denoted as  $F \wedge C \equiv_{sat}^{\omega} F$ .*

We formalize that part of the proof and extend it to *partial* truth assignments,

CaDiCaL [37]. Rule **WEAKEN**<sup>+</sup> is defined in our calculus based on the most general redundancy property and so it allows to employ every clause elimination procedure implemented in CaDiCaL including variable elimination [86], vivifica-

Either partial models or VE (Master thesis by Katharina Wagner)

I did not realize that either before Isabelle refused a proof

Implementation heavily tested... on total modals

NB: this is not (yet) in IsaSAT



# Why do We Care?

**Definition 4.2.2** (Clause Redundancy). *A witness labelled clause  $(\omega : C)$  is redundant with respect to a formula  $F$  if  $\omega(C) = \top$  and  $F|_{\alpha} \models F|_{\omega}$  for  $\alpha = \neg C$ . This is also denoted as  $F \wedge C \equiv_{sat}^{\omega} F$ .*

We formalize that part of the proof and extend it to *partial* truth assignments,

CaDiCaL [37]. Rule **WEAKEN**<sup>+</sup> is defined in our calculus based on the most general redundancy property and so it allows to employ every clause elimination procedure implemented in CaDiCaL including variable elimination [86], vivifica-

Either partial models or VE (Master thesis by Katharina Wagner)

I did not realize that either before Isabelle refused a proof

Implementation heavily tested... on total modals

NB: this is not (yet) in IsaSAT

See session 12, talk by Nikolaj Bjorner


## Why do We Care?

- We have abstract CDCL, we have inprocessing, but can we combine them?
- Corner cases easy to miss known papers have flaws!
- Metatheory, study variants See session 1
- No gap between specification and implementation  
most papers: “we can combine”


# Porting to LLVM



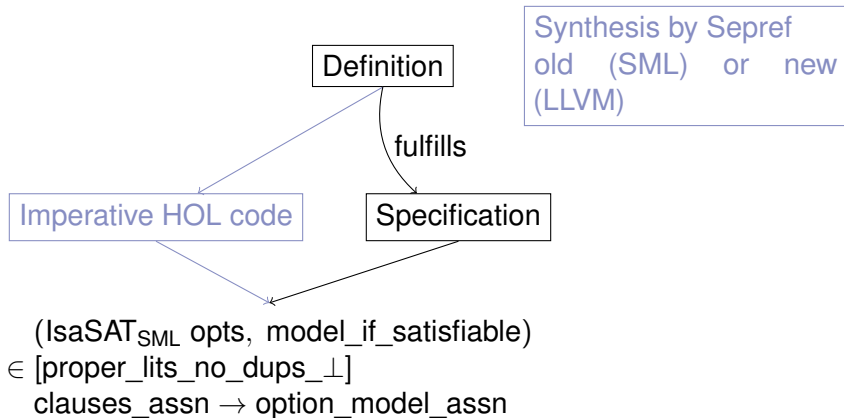
# Correctness Theorems



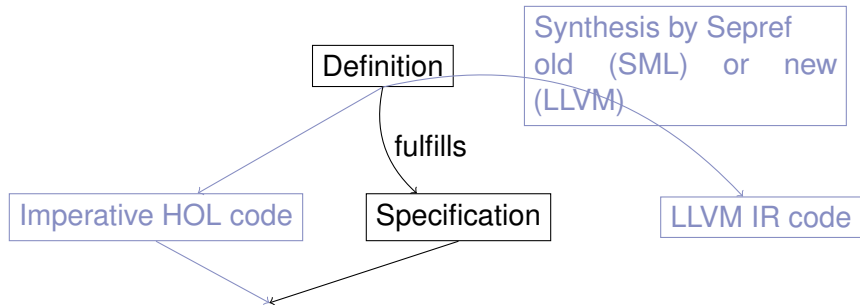
# Correctness Theorems



# Correctness Theorems

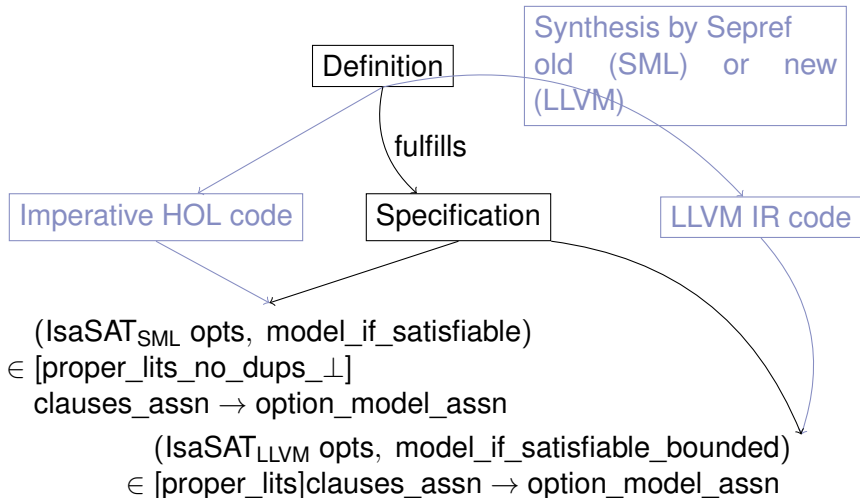


## Correctness Theorems



(IsaSAT<sub>SML</sub> opts, model\_if\_satisfiable)  
∈ [proper\_lits\_no\_dups\_⊥]  
clauses\_assn → option\_model\_assn

# Correctness Theorems





## Correctness Theorems

In LLVM: Stop if memory allocation  $> 2^{63}$   
In SML: switch to GMP (but: no compiler supports that)  
And: no machines support that

by Sepref  
(LL) or new

Imperative HOL code

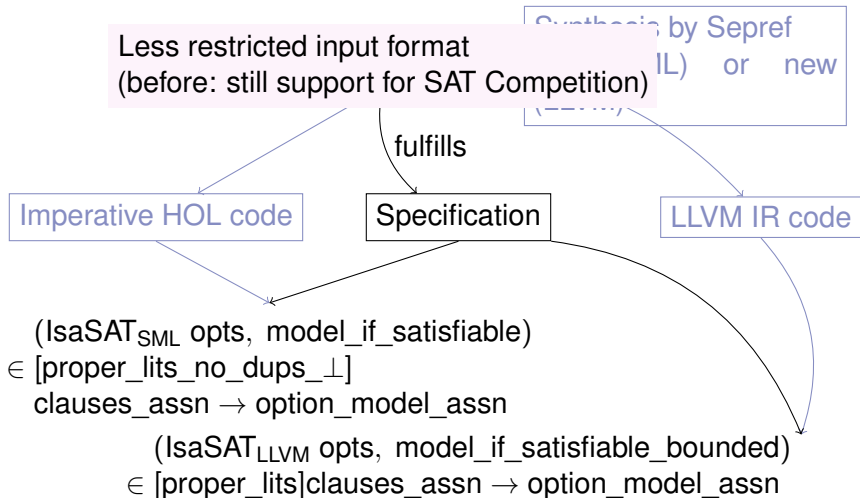
Specification

LLVM IR code

(IsaSAT<sub>SML</sub> opts, model\_if\_satisfiable)  
∈ [proper\_lits\_no\_dups\_⊥]  
clauses\_assn → option\_model\_assn

(IsaSAT<sub>LLVM</sub> opts, model\_if\_satisfiable\_bounded)  
∈ [proper\_lits] clauses\_assn → option\_model\_assn

## Correctness Theorems



# Correctness Theorems

Performance: LLVM twice as fast!

Synthesis by Sepref  
(IL) or new  
(LLVM)

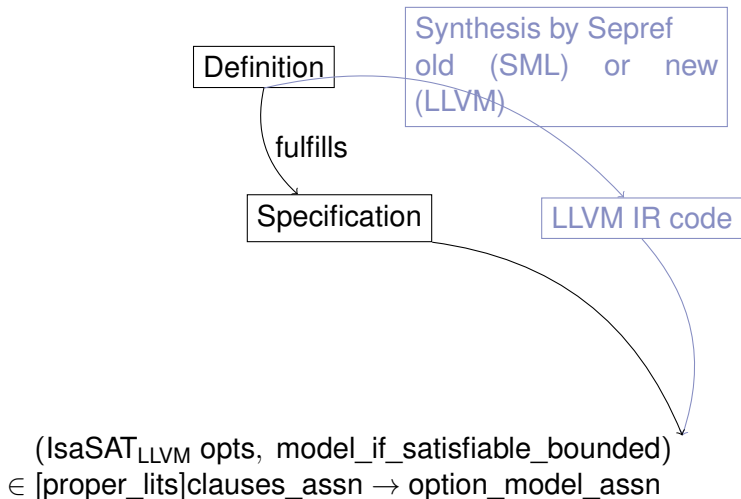
fulfills

Specification

LLVM IR code

$(\text{IsaSAT}_{\text{LLVM}} \text{ opts}, \text{model\_if\_satisfiable\_bounded})$   
 $\in [\text{proper\_lits}] \text{clauses\_assn} \rightarrow \text{option\_model\_assn}$

# Correctness Theorems



## Lessons Learned


### Aggressive reorganization of theories... but not enough

- IsaSAT is the largest development based on Isabelle/LLVM (94kloc)
- Sepref enforces a non-parallel style (and is not very fast)
- Currently 1h 20 for the non-synthesis part, 2h with it  
same for (smaller) SML version
- Synthesis rather slow and (for technical reasons) single threaded


## Lessons Learned

Local reasoning is better

Before



After



This is ongoing refactoring: the new things are ported, the olds whenever I change them

## Lessons Learned

### Performance bugs happen


A normal run:

c propagate	: 83.48% (581.66 s)
c reduce	: 0.12% (0.82 s)
c subsumption	: 0.06% (0.39 s)
c pure_lits	: 0.05% (0.33 s)
c binary_simp	: 0.02% (0.15 s)
c GC	: 0.16% (1.10 s)

(Hacked)Implementation by calling C function (via sed/grep)  
optional, but very useful

But you have to trust the (unverified) statistics!

# Performance Results





# CDF on the SAT Competition 2022




Figure: Higher is better (best: Kissat, unverified)

## Conclusion




## Summary

- IsaSAT is now much better
  - Bug: after 2 years: GC and reduction swapped!
  - I still suspect a heuristic bug
- We have now inprocessing with general rules specialized for our needs
  - See paper for Subsumption-Resolution
- We know how to improve the formalization but it won't solve all problems
- For Main track SAT Competition, DRAT proof (via sed to interface LLVM to C code)
  - "officially", I am not aware of any bug during proof checking

# Summary


- IsaSAT is now much better



- We have now inprocessing
- We know how to improve the formalization but it won't solve all problems

Questions?

# A Personal History of Solver Verification



# A Personal History of Solver Verification

