TU Delft XAIT AITL TLAB

How to discover short, shorter, and the shortest proofs of UNSAT formulas

Presenter: Konstantin Sidorov

Konstantin Sidorov, Koos van der Linden,
Gonçalo Homem de Almeida Correia, Mathijs de Weerdt, Emir Demirović
Delft University of Technology, The Netherlands

Motivation: what is the deal with *short* proofs?

- Any UNSAT claim of the solver can be *justified with a sequence of derived clauses*
- This gives a sound approach to *analyze the "thought process"* post-hoc!
- A *short* sequence of clauses = *fast* execution of a solver
- What is *the shortest* derivation of an UNSAT verdict?

Key challenge: symmetry breaking

- *Changing the clause order does not change the proof*... most of the time
- We only want to enumerate each set of inferences *once*

Prior research: DAG encoding

- Vertices are clauses, edges encode resolution steps
- Encode the proof structure as a SAT problem
- Break symmetries by *topological ordering*

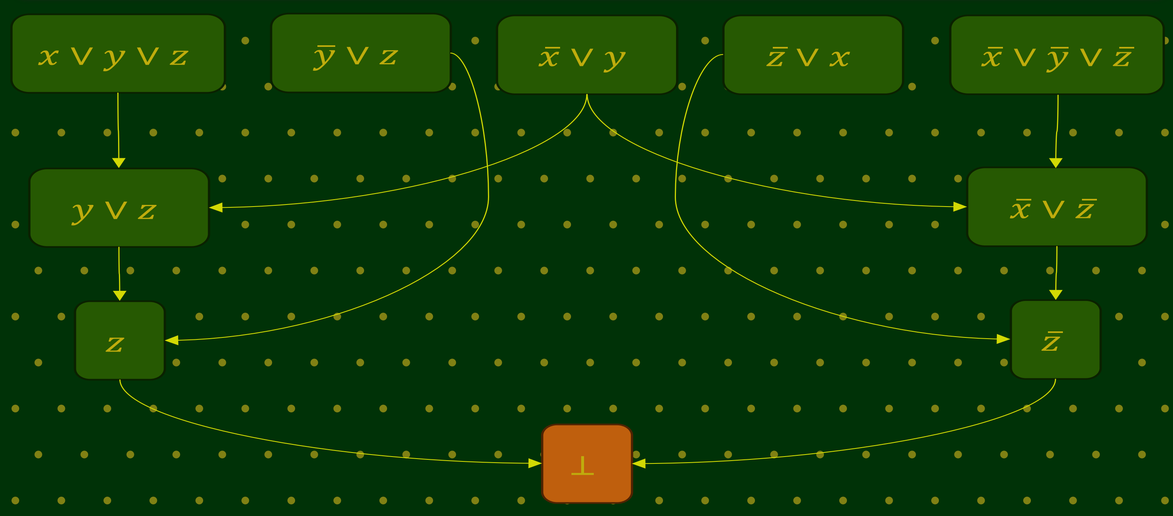

Does it break all symmetries?

No.

Some clauses can be derived in one way...

...or another, but with a *different* DAG

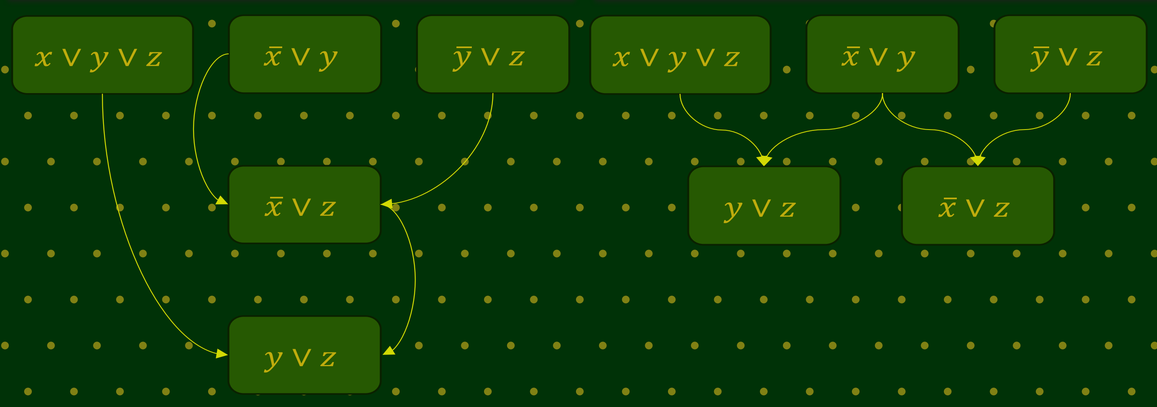

But we do break all symmetries, right?

- We design a *custom branch-and-bound* procedure
- Track the clauses that *have not been derived* at the first opportunity
- Branching *skips* those clauses

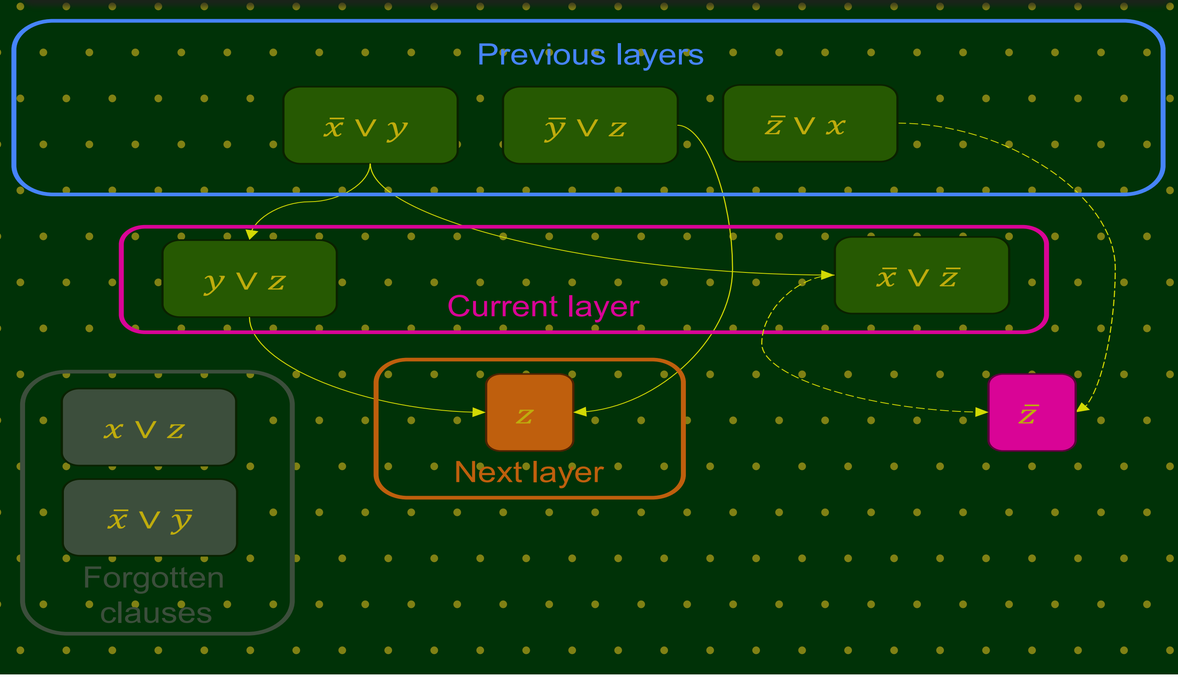

Experimental results

CaDiCaL produces *50% extra clauses* against the optimal proofs for 3-CNFs

Our approach improves time to the shortest proofs *by orders of magnitude*