

Modern CDCL SAT Solvers

SAT / SMT Summer School

12. June 2012

Fondazione Bruno Kessler

Trento, Italy


Armin Biere

Institute for Formal Models and Verification


Johannes Kepler University, Linz, Austria

<http://fmv.jku.at/biere/talks/Biere-SATSMT12.pdf>


<http://fmv.jku.at/cleaneling/cleaneling00a.zip>



Results of the SAT competition/race winners on the SAT 2009 application benchmarks, 20mn timeout



Results of the SAT competition/race winners on the SAT 2009 application benchmarks, 20mn timeout



- dates back to the 50'ies:

1st version DP is *resolution based* \Rightarrow SatELite preprocessor [EénBiere05]

2st version D(P)LL splits space for time \Rightarrow **CDCL**

- **ideas:**

- 1st version: eliminate the two cases of assigning a variable in space or
- 2nd version: case analysis in time, e.g. try $x = 0, 1$ in turn and recurse

- most successful SAT solvers are based on variant (CDCL) of the second version works for very large instances
- recent (≤ 15 years) optimizations:
backjumping, learning, UIPs, dynamic splitting heuristics, fast data structures
(we will have a look at each of them)

forever

if $F = \top$ **return** *satisfiable*

if $\perp \in F$ **return** *unsatisfiable*

pick remaining variable x

add all resolvents on x

remove all clauses with x and $\neg x$

\Rightarrow SatELite preprocessor [EénBiere05]

$DPLL(F)$

$F := BCP(F)$

boolean constraint propagation

if $F = \top$ **return** *satisfiable*

if $\perp \in F$ **return** *unsatisfiable*


pick remaining variable x and literal $l \in \{x, \neg x\}$

if $DPLL(F \wedge \{l\})$ returns *satisfiable* **return** *satisfiable*

return $DPLL(F \wedge \{\neg l\})$

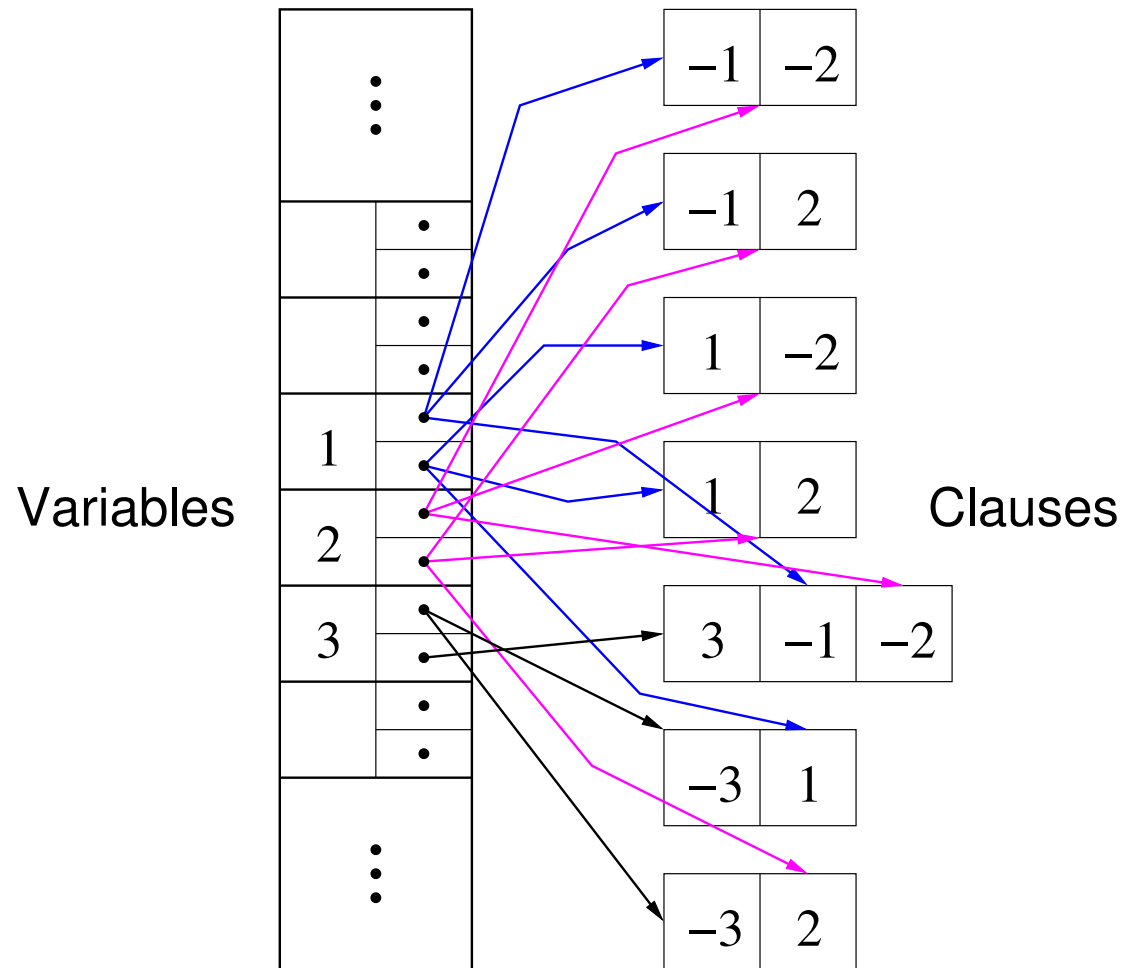
\Rightarrow



CDCL





clauses

- $\neg a \vee \neg b \vee \neg c$
- $\neg a \vee \neg b \vee c$
- $\neg a \vee b \vee \neg c$
- $\neg a \vee b \vee c$
- $a \vee \neg b \vee \neg c$
- $a \vee \neg b \vee c$
- $a \vee b \vee \neg c$
- $a \vee b \vee c$








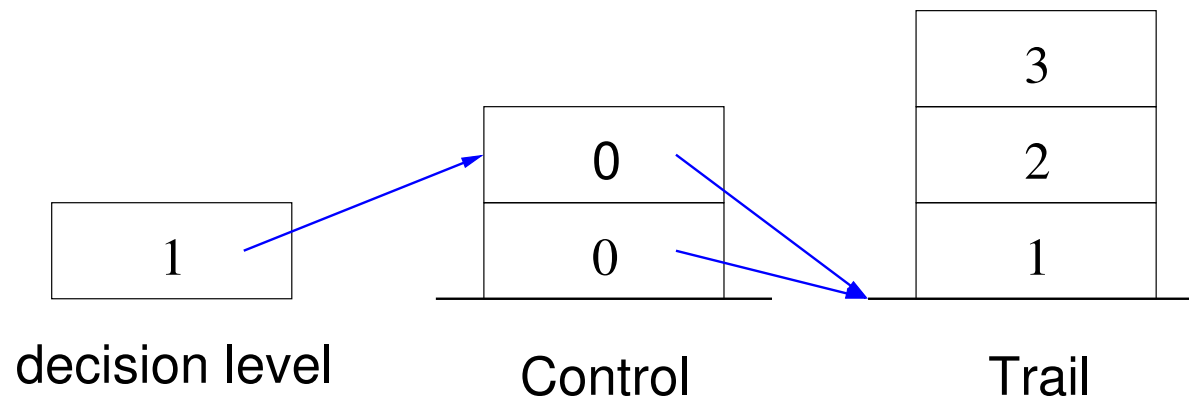
Decide





Assign







BCP

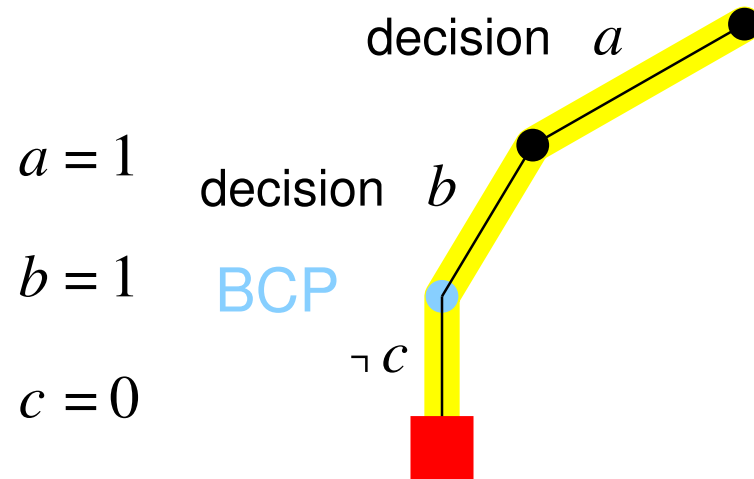


Decide









clauses

$\neg a \vee \neg b \vee \neg c$

$\neg a \vee \neg b \vee c$

$\neg a \vee b \vee \neg c$

$\neg a \vee b \vee c$

$a \vee \neg b \vee \neg c$

$a \vee \neg b \vee c$

$a \vee b \vee \neg c$

$a \vee b \vee c$


learn $\neg a \vee \neg b$

$a = 1$

$b = 0$

$c = 0$

decision a



clauses

$\neg a \vee \neg b \vee \neg c$

$\neg a \vee \neg b \vee c$

$\neg a \vee b \vee \neg c$

$\neg a \vee b \vee c$

$a \vee \neg b \vee \neg c$

$a \vee \neg b \vee c$

$a \vee b \vee \neg c$

$a \vee b \vee c$

$\neg a \vee \neg b$


learn

$\neg a$

$a = 1$

$b = 0$

$c = 0$



clauses

$\neg a \vee \neg b \vee \neg c$

$\neg a \vee \neg b \vee c$

$\neg a \vee b \vee \neg c$

$\neg a \vee b \vee c$

$a \vee \neg b \vee \neg c$

$a \vee \neg b \vee c$

$a \vee b \vee \neg c$

$a \vee b \vee c$

$\neg a \vee \neg b$

$\neg a$


learn

c

$a = 1$

$b = 0$

$c = 0$



clauses

$\neg a \vee \neg b \vee \neg c$

$\neg a \vee \neg b \vee c$

$\neg a \vee b \vee \neg c$

$\neg a \vee b \vee c$

$a \vee \neg b \vee \neg c$

$a \vee \neg b \vee c$

$a \vee b \vee \neg c$

$a \vee b \vee c$

$\neg a \vee \neg b$

$\neg a$

c

learn

\perp

empty clause

- **static heuristics:**

- one *linear* order determined before solver is started
- usually quite fast to compute, since only calculated once
- and thus can also use more expensive algorithms

- **dynamic heuristics**

- typically calculated from number of occurrences of literals (in unsatisfied clauses)
- could be rather expensive, since it requires traversal of all clauses (or more expensive updates in BCP)
- effective *second order* dynamic heuristics (e.g. VSIDS in Chaff)

- Dynamic Largest Individual Sum (DLIS)
 - fastest dynamic *first order* heuristic (e.g. GRASP solver)
 - choose literal (variable + phase) which occurs most often (ignore satisfied clauses)
 - requires explicit traversal of CNF (or more expensive BCP)
- look-ahead heuristics (e.g. SATZ or MARCH solver) **failed literals, probing**
 - trial assignments and BCP for all/some unassigned variables (both phases)
 - if BCP leads to conflict, enforce toggled assignment of current trial decision
 - optionally learn binary clauses and perform equivalent literal substitution
 - decision: most balanced w.r.t. prop. assignments / sat. clauses / reduced clauses
 - related to our recent Cube & Conquer paper [HeuleKullmanWieringaBiere-HVC'11]

Chaff [MoskewiczMadiganZhaoZhangMalik'01]

- increment score of involved variables by 1
- decay score of all variables every 256'th conflict by halving the score
- sort priority queue after decay and not at every conflict

MiniSAT uses EVSIDS [EénSörensson'03/'06]

- update score of involved variables as actually LIS would also do
- dynamically adjust increment: $\delta' = \delta \cdot \frac{1}{f}$ typically increment δ by 5%
- use floating point representation of score
- “rescore” to avoid overflow in regular intervals
- EVSIDS linearly related to NVSIDS

(consider only one variable)

$$\delta_k = \begin{cases} 1 & \text{if involved in } k\text{-th conflict} \\ 0 & \text{otherwise} \end{cases}$$

$$i_k = (1 - f) \cdot \delta_k$$

$$s_n = (\dots (i_1 \cdot f + i_2) \cdot f + i_3) \cdot f \dots) \cdot f + i_n = \sum_{k=1}^n i_k \cdot f^{n-k} = (1 - f) \cdot \sum_{k=1}^n \delta_k \cdot f^{n-k} \quad (\text{NVSIDS})$$

$$S_n = \frac{f^{-n}}{(1 - f)} \cdot s_n = \frac{f^{-n}}{(1 - f)} \cdot (1 - f) \cdot \sum_{k=1}^n \delta_k \cdot f^{n-k} = \sum_{k=1}^n \delta_k \cdot f^{-k} \quad (\text{EVSIDS})$$

[GoldbergNovikov-DATE'02]


- observation:
 - recently added conflict clauses contain all the good variables of VSIDS
 - the order of those clauses is not used in VSIDS
- basic idea:
 - simply try to satisfy recently learned clauses first
 - use VSIDS to choose the decision variable for one clause
 - if all learned clauses are satisfied use other heuristics
 - intuitively obtains another order of localization (no proofs yet)
- mixed results as other variants VMTF, CMTF (var/clause move to front)

- keeping all learned clauses slows down BCP kind of quadratically
 - so SATO and ReSAT just kept only “short” clauses
- better periodically delete “useless” learned clauses
 - keep a certain number of learned clauses “search cache”
 - if this number is reached MiniSAT reduces (deletes) half of the clauses
 - keep *most active*, then *shortest*, then *youngest* (LIFO) clauses
 - after reduction maximum number kept learned clauses is increased geometrically
- LBD (Glue) based (apriori!) prediction for usefulness [AudemardLaurent’09]
 - LBD (Glue) = number of decision-levels in the learned clause
 - allows arithmetic increase of number of kept learned clauses

- for satisfiable instances the solver may get stuck in the unsatisfiable part
 - even if the search space contains a large satisfiable part
- often it is a good strategy to abandon the current search and restart
 - restart after the number of decisions reached a *restart limit*
- avoid to run into the same dead end
 - by randomization (either on the decision variable or its phase)
 - and/or just keep all the learned clauses
- for completeness dynamically increase restart limit

Inner/Outer Restart Intervals

378 restarts in 104408 conflicts



```
int inner = 100, outer = 100;
int restarts = 0, conflicts = 0;


for (;;)
{
    ... // run SAT core loop for 'inner' conflicts

    restarts++;
    conflicts += inner;

    if (inner >= outer)
    {
        outer *= 1.1;
        inner = 100;
    }
    else
        inner *= 1.1;
}
```

Luby's Restart Intervals

70 restarts in 104448 conflicts




```
unsigned
luby (unsigned i)
{
    unsigned k;

    for (k = 1; k < 32; k++)
        if (i == (1 << k) - 1)
            return 1 << (k - 1);

    for (k = 1;; k++)
        if ((1 << (k - 1)) <= i && i < (1 << k) - 1)
            return luby (i - (1 << (k-1)) + 1);
}


limit = 512 * luby (++restarts);
... // run SAT core loop for 'limit' conflicts
```


- phase assignment:
 - assign decision variable to 0 or 1?
 - only thing that matters in *satisfiable* instances
- “phase saving” as in RSat:
 - pick phase of last assignment (if not forced to, do not toggle assignment)
 - initially use statically computed phase (typically LIS)
 - so can be seen to maintain a *global full assignment*
- rapid restarts: varying restart interval with bursts of restarts
 - not only theoretically avoids local minima
 - works nicely together with phase saving




If y has never been used to derive a conflict, then skip \bar{y} case.

Immediately *jump back* to the \bar{x} case – assuming x was used.









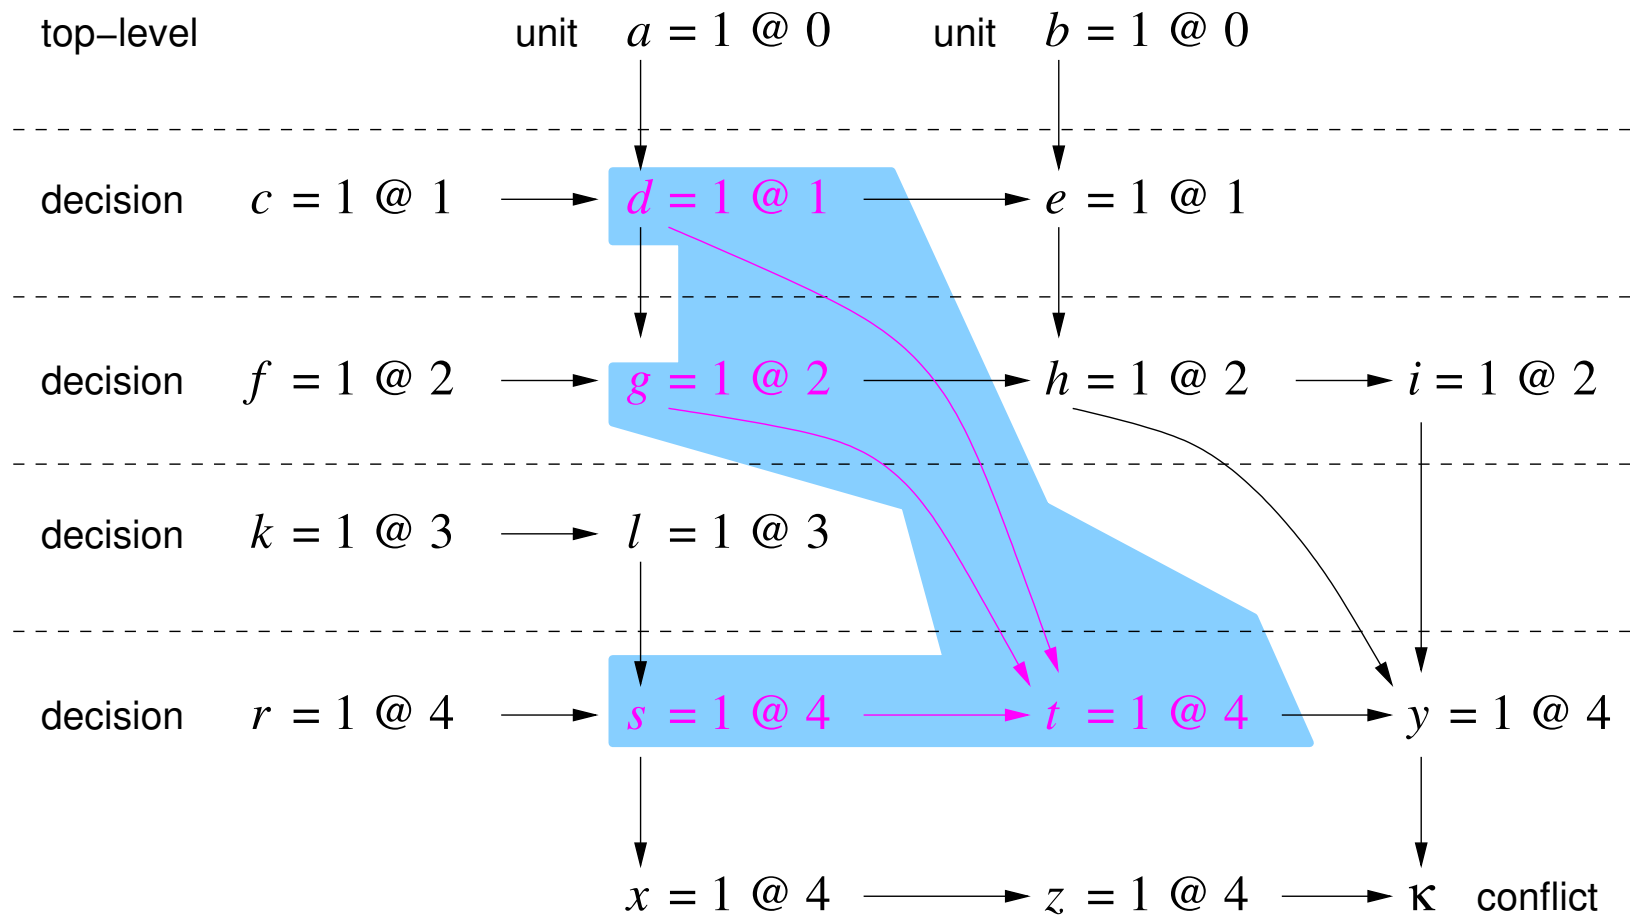
a simple cut always exists: set of roots (decisions) contributing to the conflict

```
Status Solver::search (long limit) {
    long conflicts = 0; Clause * conflict; Status res = UNKNOWN;
    while (!res)
        if (empty) res = UNSATISFIABLE;
        else if ((conflict = bcp ())) analyze (conflict), conflicts++;
        else if (conflicts >= limit) break;
        else if (reducing ()) reduce ();
        else if (restarting ()) restart ();
        else if (!decide ()) res = SATISFIABLE;
    return res;
}
```


```
Status Solver::solve () {
    long conflicts = 0, steps = 1e6;
    Status res;
    for (;;)
        if ((res = search (conflicts)) break;
        else if ((res = simplify (steps)) break;
        else conflicts += 1e4, steps += 1e6;
    return res;
}
```








$$d \wedge g \wedge s \rightarrow t \quad \equiv \quad (\bar{d} \vee \bar{g} \vee \bar{s} \vee t)$$




$$\neg(y \wedge z) \equiv (\bar{y} \vee \bar{z})$$




$$(\bar{h} \vee \bar{i} \vee \bar{t} \vee y)$$


$$(\bar{y} \vee \bar{z})$$




$$\frac{(\bar{h} \vee \bar{i} \vee \bar{t} \vee y) \quad (\bar{y} \vee \bar{z})}{(\bar{h} \vee \bar{i} \vee \bar{t} \vee \bar{z})}$$




$$\frac{(\bar{h} \vee \bar{i} \vee \bar{t} \vee y) \quad (\bar{y} \vee \bar{z})}{(\bar{h} \vee \bar{i} \vee \bar{t} \vee \bar{z})}$$




$$(\bar{h} \vee \bar{i} \vee \bar{t} \vee \bar{z})$$



$$\frac{(\bar{d} \vee \bar{g} \vee \bar{s} \vee t) \quad (\bar{h} \vee \bar{i} \vee \bar{t} \vee \bar{z})}{(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i} \vee \bar{z})}$$




$$\frac{(\bar{x} \vee z) \quad (\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i} \vee \bar{z})}{(\bar{x} \vee \bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})}$$



$$\frac{(\bar{s} \vee x) \quad (\bar{x} \vee \bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})}{(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})}$$


self subsuming resolution




$$(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})$$

UIP = *unique implication point* dominates conflict on the last level


- can be found by graph traversal in the order of made assignments
 - *trail* respects this order
 - traverse reasons of variables on trail starting with conflict
- count “open paths”
 - initially size of clause with only false literals
 - decrease counter if new reason / antecedent clause resolved
 - if all paths converged, i.e. counter = 1, then this node is a UIP
 - decision of current decision level is a UIP and thus a *sentinel*




$$\frac{(\bar{l} \vee \bar{r} \vee s) \quad (\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})}{(\bar{l} \vee \bar{r} \vee \bar{d} \vee \bar{g} \vee \bar{h} \vee \bar{i})}$$



$$(\bar{d} \vee \bar{g} \vee \bar{l} \vee \bar{r} \vee \bar{h} \vee \bar{i})$$




$$(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})$$



$$\frac{(\bar{h} \vee i) \quad (\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h} \vee \bar{i})}{(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h})}$$

self subsuming resolution



$$(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h})$$


Two step algorithm:

1. mark all variables in 1st UIP clause
2. remove literals with all antecedent literals also marked


Correctness:

- removal of literals in step 2 are self subsuming resolution steps.
- implication graph is acyclic.


Confluence: produces a unique result.



$$(\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h})$$



$$\begin{array}{c}
 \frac{(\bar{e} \vee \bar{g} \vee h) \quad (\bar{d} \vee \bar{g} \vee \bar{s} \vee \bar{h})}{(\bar{d} \vee \bar{b} \vee e) \quad (\bar{e} \vee \bar{d} \vee \bar{g} \vee \bar{s})} \\
 \frac{(b)}{(\bar{d} \vee \bar{g} \vee \bar{s})}
 \end{array}$$



$$(\bar{d} \vee \bar{g} \vee \bar{s})$$

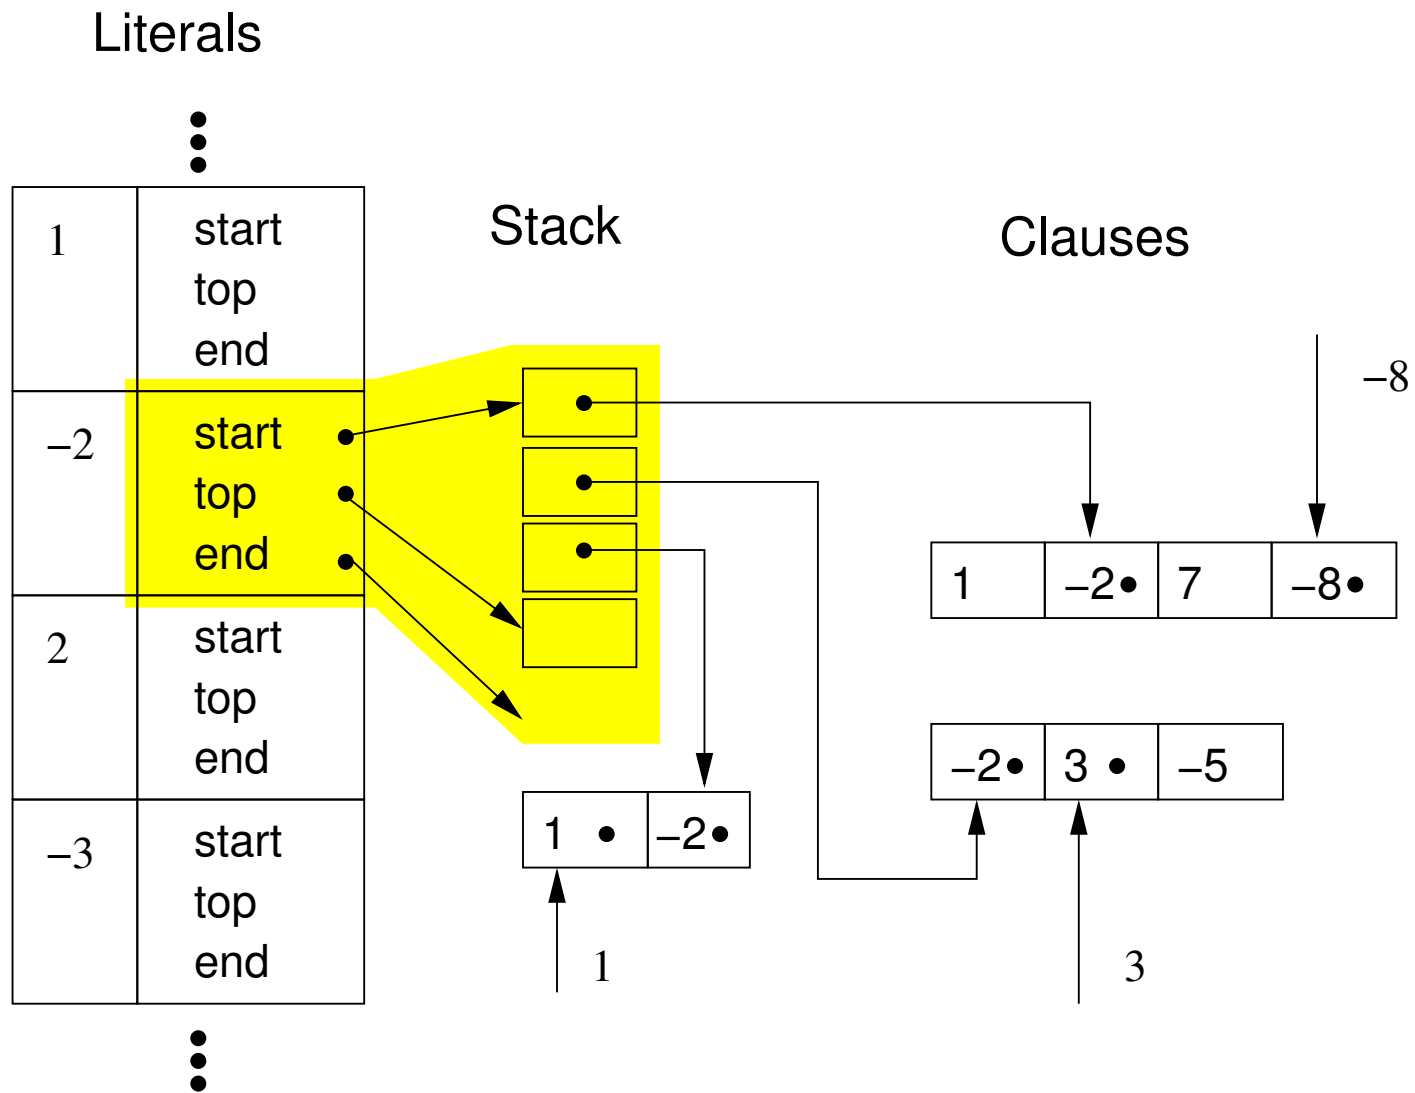
Four step algorithm:


1. mark all variables in 1st UIP clause
2. for each candidate literal: search implication graph
3. start at antecedents of candidate literals
4. if search always terminates at marked literals remove candidate


Correctness and **Confluence** as in local version!!!


Optimization: terminate early with failure if new decision level is “pulled in”


- original idea from SATO [ZhangStickel'00]
 - maintain the invariant: always watch two non-false literals
 - if a watched literal becomes *false* replace it
 - if no replacement can be found clause is either unit or empty
 - original version used *head* and *tail* pointers on Tries
- improved variant from Chaff [MoskewiczMadiganZhaoZhangMalik'01]
 - watch pointers can move arbitrarily SATO: *head* forward, *trail* backward
 - no update needed during backtracking
- *one* watch is enough to ensure correctness but loses *arc consistency*
- reduces *visiting* clauses by 10x, particularly useful for large and many learned clauses



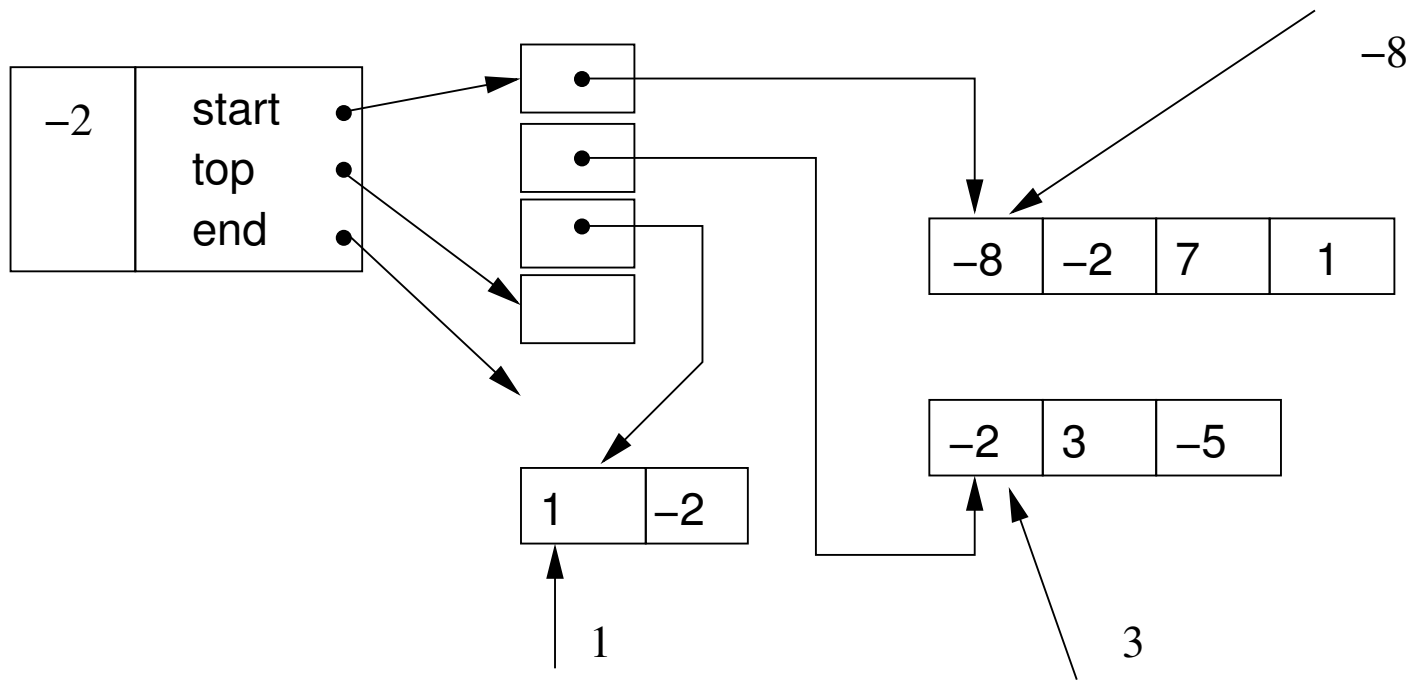










still seems to be best way for *real* sharing of clauses in multi-threaded solvers





invariant: first two literals are watched

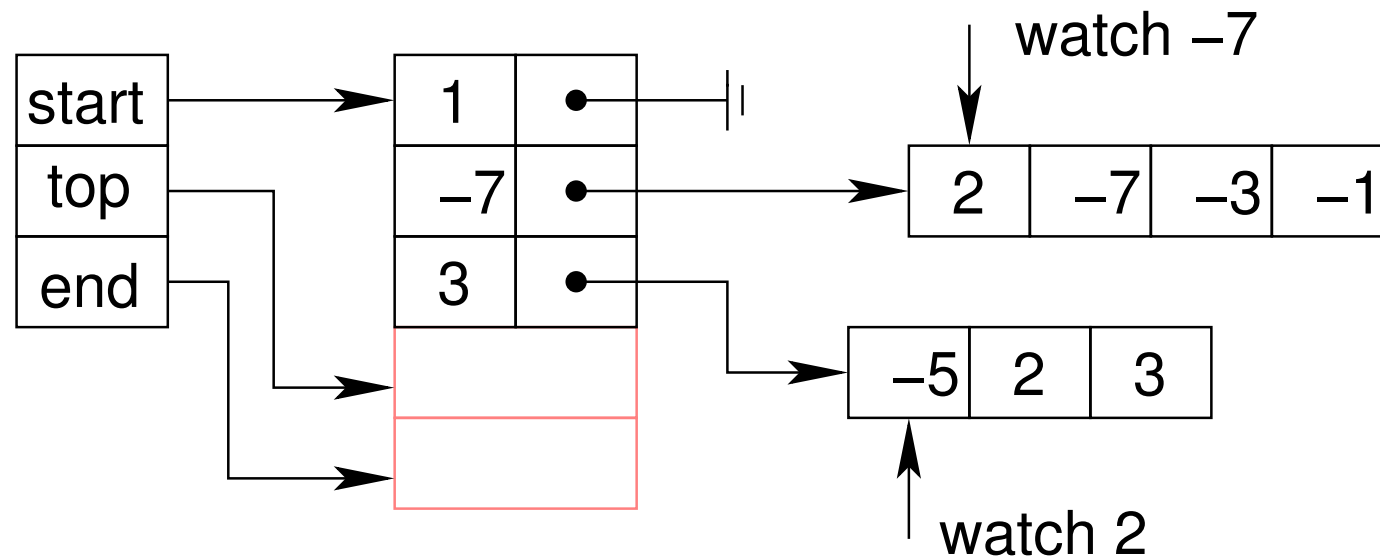




invariant: first two literals are watched

Additional Binary Clause Watcher Stack





observation: often the *other* watched literal satisfies the clause

so cache this literals in watch list to avoid pointer dereference

for binary clause no need to store clause at all

can easily be adjusted for ternary clauses (with full occurrence lists)

LINGELING uses more compact pointer-less variant

we are still working on tracking down the origin before [Freeman'95] [LeBerre'01]

- key technique in look-ahead solvers such as Satz, OKSolver, March
 - failed literal probing at all search nodes
 - used to find the best decision variable and phase
- simple algorithm
 1. assume literal l , propagate (BCP), if this results in conflict, add unit clause $\neg l$
 2. continue with all literals l until *saturation* (nothing changes)
- quadratic to cubic complexity
 - BCP linear in the size of the formula 1st linear factor
 - each variable needs to be tried 2nd linear factor
 - and tried again if some unit has been derived 3rd linear factor

- lifting
 - complete case split: literals implied in all cases become units
 - similar to Stålmarm's method and Recursive Learning [PradhamKunz'94]
- asymmetric branching
 - assume all but one literal of a clause to be false
 - if BCP leads to conflict remove originally remaining unassigned literal
 - implemented for a long time in MiniSAT but switched off by default
- generalizations:
 - vivification [PietteHamadiSais ECAI'08]
 - distillation [JinSomenzi'05][HanSomenzi DAC'07] probably most general (+ tries)

- similar to look-ahead heuristics: polynomially bounded search
 - may be recursively applied (however, is often too expensive)
- Stålmarck's Method
 - works on triplets (intermediate form of the Tseitin transformation):
 $x = (a \wedge b), y = (c \vee d), z = (e \oplus f)$ etc.
 - generalization of BCP to (in)equalities between variables
 - **test rule** splits on the two values of a variable
- Recursive Learning (Kunz & Pradhan)
 - (originally) works on circuit structure (derives implications)
 - splits on different ways to *justify* a certain variable value

[DavisPutnam60][Biere SAT'04] [SubbarayanPradhan SAT'04] [EénBiere SAT'05]

- use DP to existentially quantify out variables as in [DavisPutnam60]
- only remove a variable if this does not add (too many) clauses
 - do not count tautological resolvents
 - detect units on-the-fly
- schedule removal attempts with a priority queue [Biere SAT'04] [EénBiere SAT'05]
 - variables ordered by the number of occurrences
- strengthen and remove subsumed clauses (on-the-fly)
(SATeLite [EénBiere SAT'05] and Quantor [Biere SAT'04])

- for each (new or strengthened) clause
 - traverse list of clauses of the least occurring literal in the clause
 - check whether traversed clauses are subsumed or
 - strengthen traversed clauses by self-subsumption [EénBiere SAT'05]
 - use Bloom Filters (as in “bit-state hashing”), aka signatures
- check old clauses being subsumed by new clause: **backward (self) subsumption**
 - new clause (self) subsumes existing clause
 - new clause smaller or equal in size
- check new clause to be subsumed by existing clauses **forward (self) subsumption**
 - can be made more efficient by one-watcher scheme [Zhang-SAT'05]

one clause $C \in F$ with l

all clauses in F with \bar{l}

fix a CNF F

$$\bar{l} \vee \bar{a} \vee c$$

$$a \vee b \vee l$$

$$\bar{l} \vee \bar{b} \vee d$$

all resolvents of C on l are tautological

\Rightarrow

C can be removed

Proof assume assignment σ satisfies $F \setminus C$ but not C

can be extended to a satisfying assignment of F by flipping value of l

Definition A literal l in a clause C of a CNF F **blocks** C w.r.t. F if for every clause $C' \in F$ with $\bar{l} \in C'$, the resolvent $(C \setminus \{l\}) \cup (C' \setminus \{\bar{l}\})$ obtained from resolving C and C' on l is a tautology.

Definition [Blocked Clause] A clause is **blocked** if has a literal that blocks it.

Definition [Blocked Literal] A literal is **blocked** if it blocks a clause.

Example

$$(a \vee b) \wedge (a \vee \bar{b} \vee \bar{c}) \wedge (\bar{a} \vee c)$$

only **first clause** is not blocked.

second clause contains two blocked literals: a and \bar{c} .

literal c in the **last clause** is blocked.

after removing either $(a \vee \bar{b} \vee \bar{c})$ or $(\bar{a} \vee c)$, the clause $(a \vee b)$ becomes blocked
actually all clauses can be removed

COI Cone-of-Influence reduction

MIR Monotone-Input-Reduction

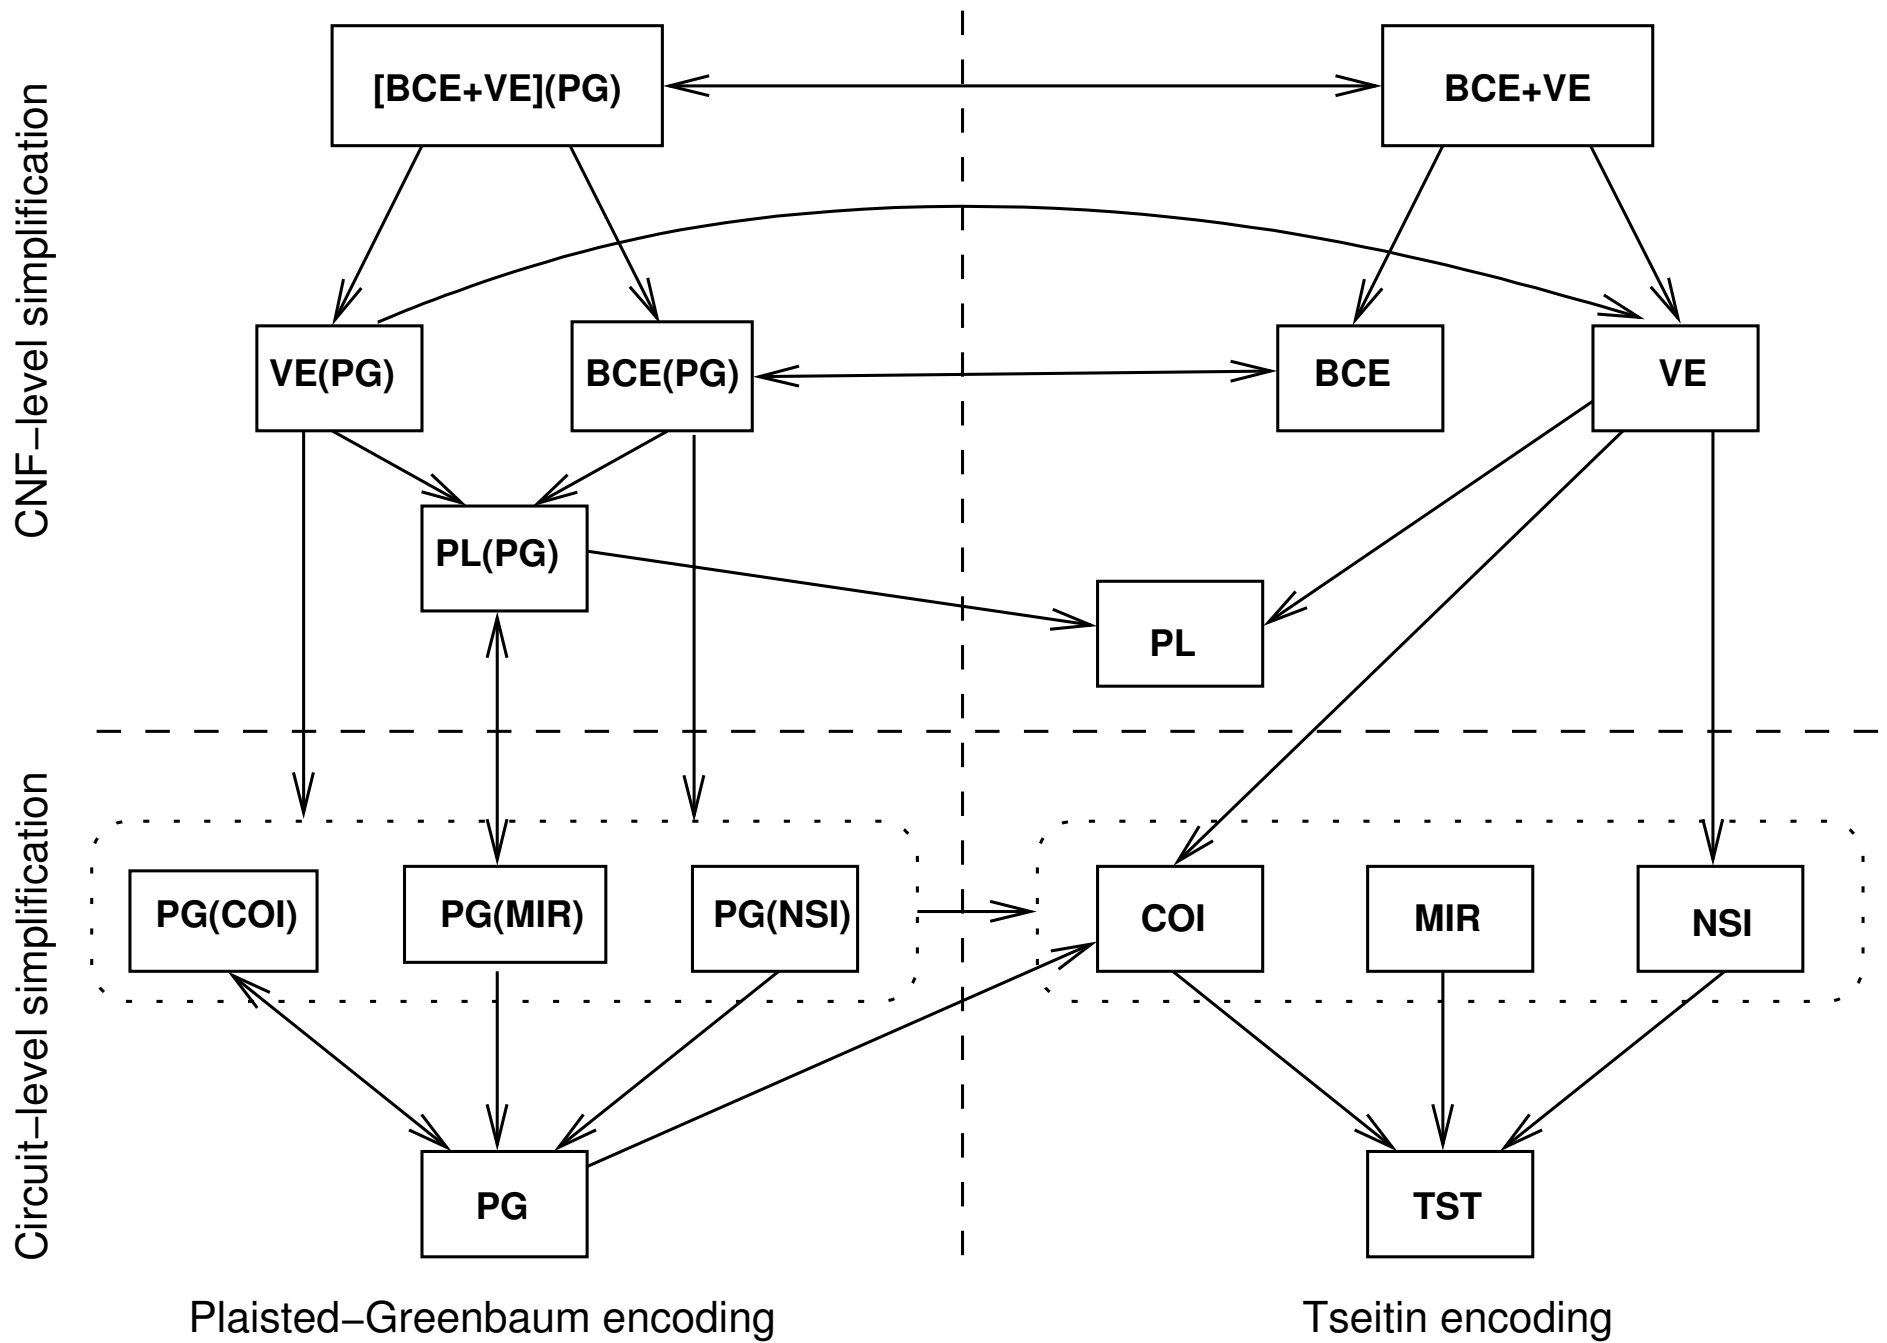
NSI Non-Shared Inputs reduction

PG Plaisted-Greenbaum polarity based encoding

TST standard Tseitin encoding

VE Variable-Elimination as in DP / Quantor / SATeLite

BCE Blocked-Clause-Elimination



[BCE+VE](PG)

BCE+VE

VE(PG)

BCE(PG)

BCE

VE

PL(PG)

PL

PG(COI)

PG(MIR)

PG(NSI)

COI

MIR

NSI

PG

TST

PrecoSAT [Biere'09], Lingeling [Biere'10], now also in CryptoMiniSAT (Mate Soos)

- preprocessing can be extremely beneficial
 - most SAT competition solvers use variable elimination (VE) [EénBiere SAT'05]
 - equivalence / XOR reasoning
 - probing / failed literal preprocessing / hyper binary resolution
 - however, even though polynomial, **can not be run until completion**
- simple idea to benefit from full preprocessing without penalty
 - **“preempt” preprocessors** after some time
 - **resume preprocessing** between restarts
 - limit preprocessing time in relation to search time

equivalent literal substitution find strongly connected components in binary implication graph, replace equivalent literals by representatives

boolean ring reasoning extract XORs, then Gaussian elimination etc.

hyper-binary resolution focus on producing binary resolvents

hidden/asymmetric tautology elimination discover redundant clauses through probing

covered clause elimination use covered literals in probing for redundant clauses

unhiding randomized algorithm (one phase linear) for clause removal and strengthening

- allows to use *costly* preprocessors
 - without increasing run-time “much” in the worst-case
 - still useful for benchmarks where these costly techniques help
 - good examples: probing and distillation even VE can be costly
- additional benefit:
 - makes units / equivalences learned in search available to preprocessing
 - particularly interesting if preprocessing simulates encoding optimizations
- danger of hiding “bad” implementation though ...
- ... and hard(er) to get right! “Inprocessing Rules” JärvisaloHeuleBiere’12 at IJCAR