

Understanding Modern SAT Solvers

VTSA'12

Summer School 2012

Verification Technology, Systems & Applications

September 2012

Max Planck Institut Informatik

Saarbrücken, Germany

Armin Biere

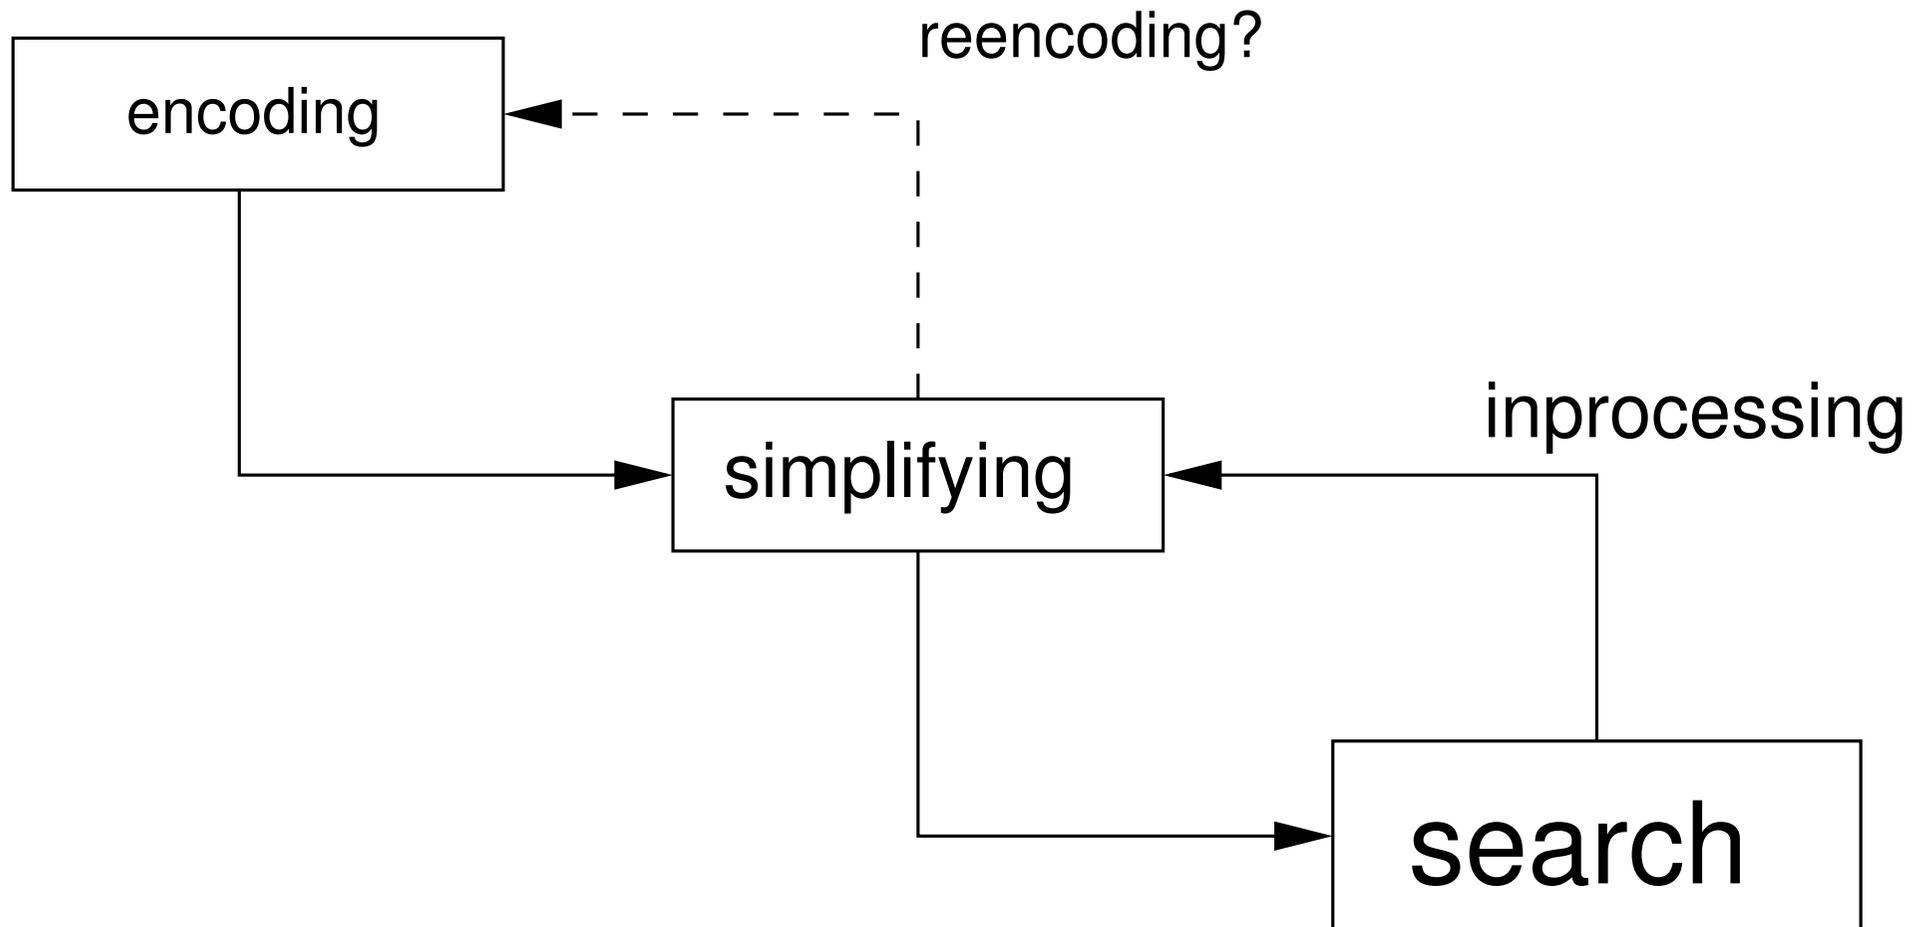
Institute for Formal Models and Verification

Johannes Kepler University, Linz, Austria

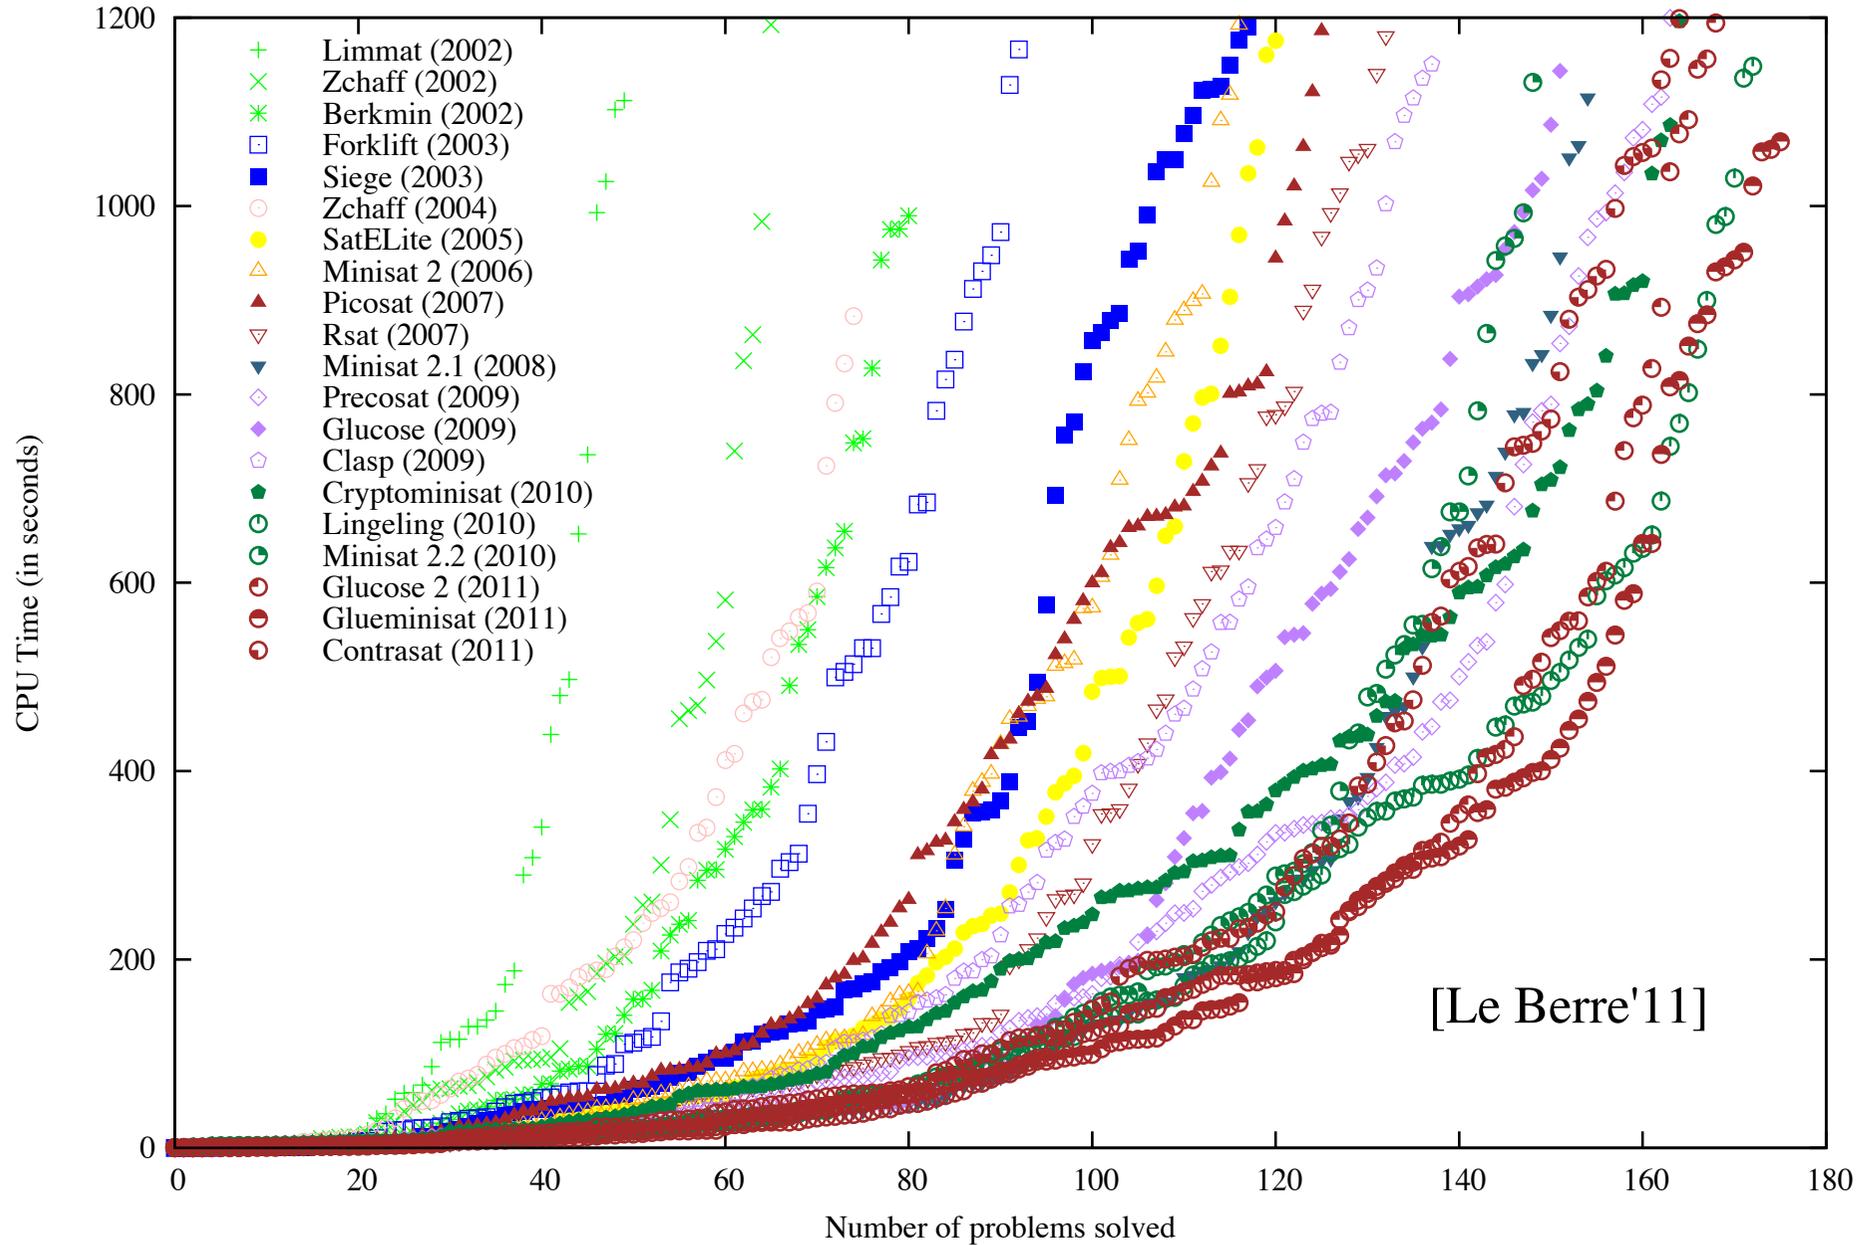
<http://fmv.jku.at/picosat>

<http://fmv.jku.at/biere/talks/Biere-VTSA12-talk.pdf>

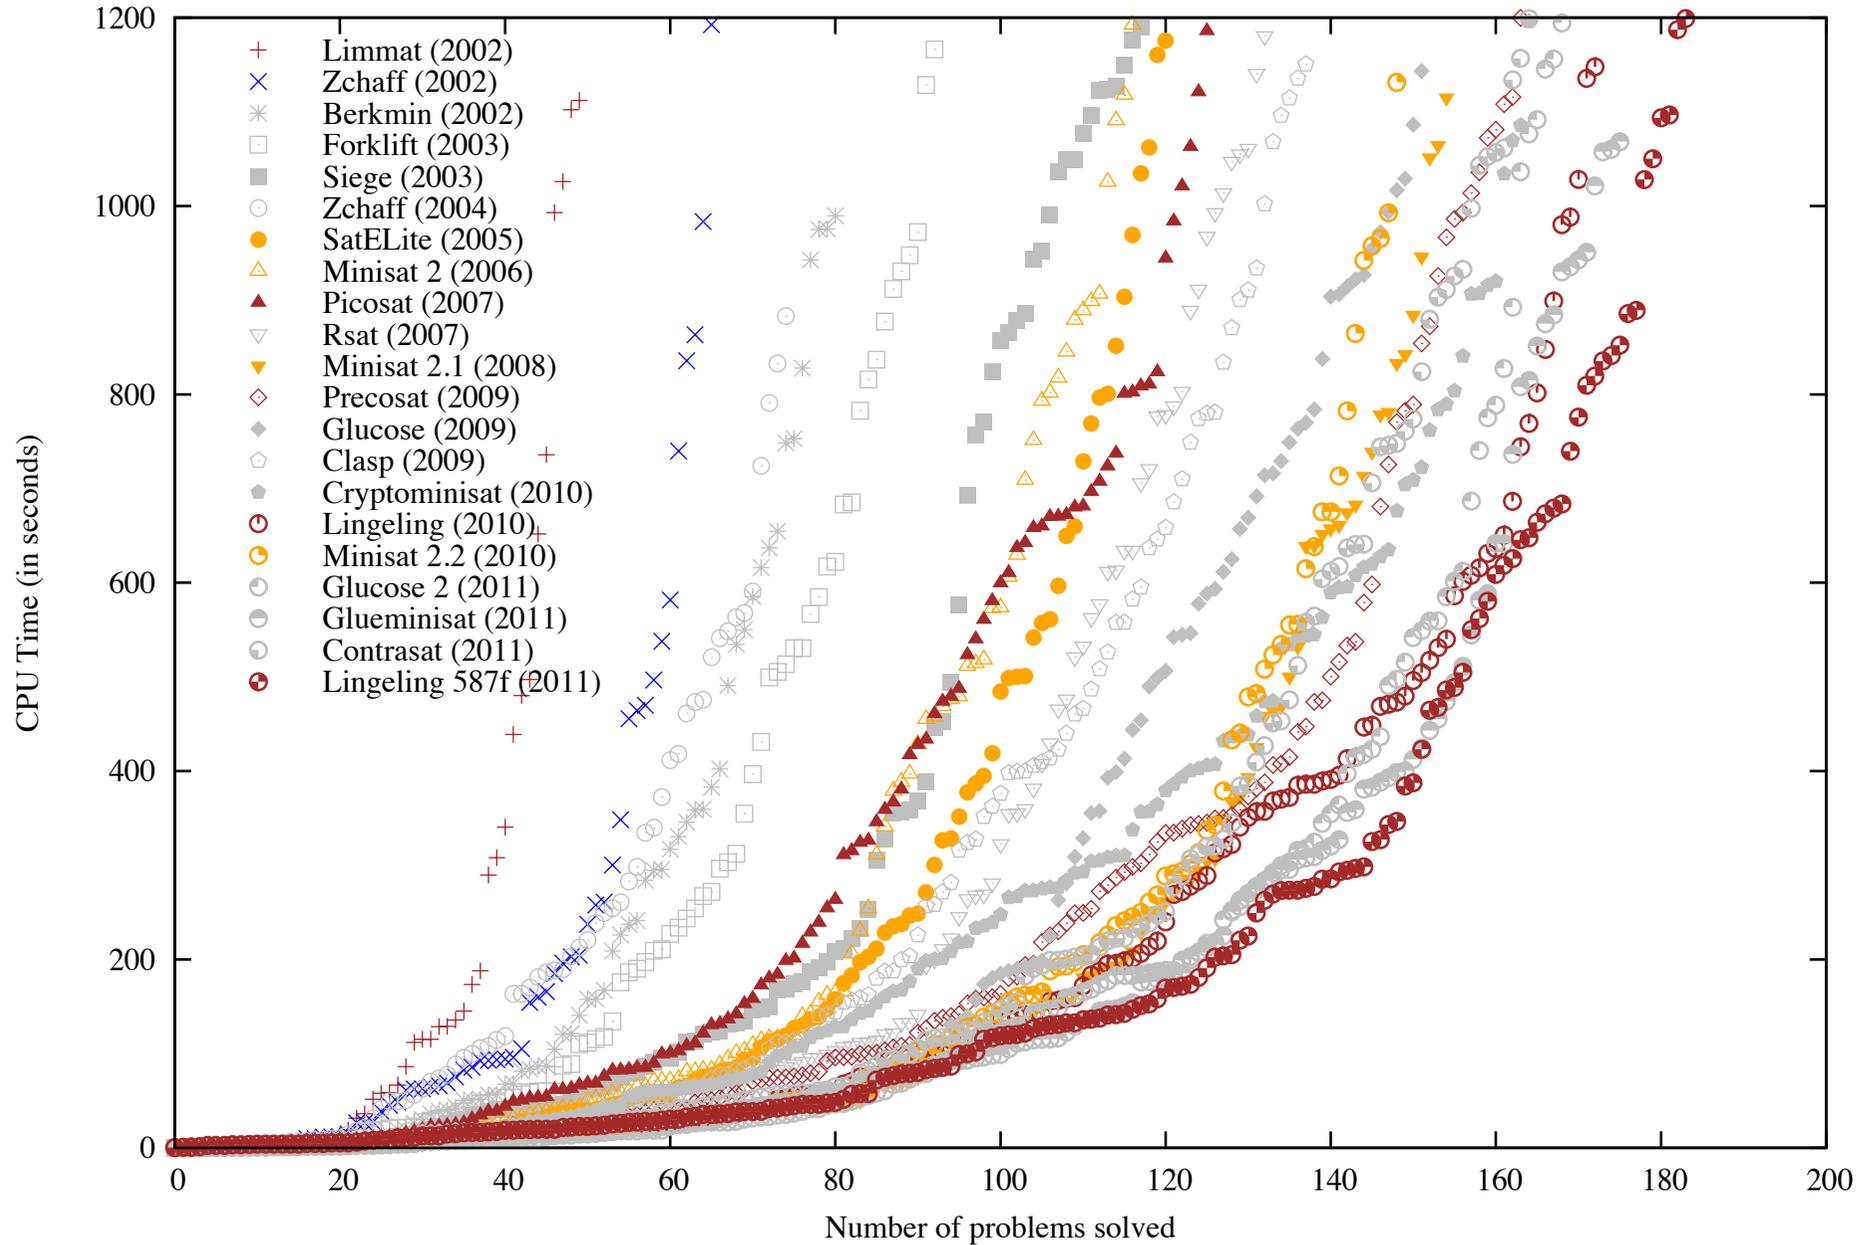
<http://fmv.jku.at/cleaneling/cleaneling00f.zip>





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







- propositional logic:
 - variables **tie** **shirt**
 - negation \neg (not)
 - disjunction \vee disjunction (or)
 - conjunction \wedge conjunction (and)
- three conditions / clauses:
 - clearly one should not wear a **tie** without a **shirt** $\neg\mathbf{tie} \vee \mathbf{shirt}$
 - not wearing a **tie** nor a **shirt** is impolite $\mathbf{tie} \vee \mathbf{shirt}$
 - wearing a **tie** and a **shirt** is overkill $\neg(\mathbf{tie} \wedge \mathbf{shirt}) \equiv \neg\mathbf{tie} \vee \neg\mathbf{shirt}$
- is the formula $(\neg\mathbf{tie} \vee \mathbf{shirt}) \wedge (\mathbf{tie} \vee \mathbf{shirt}) \wedge (\neg\mathbf{tie} \vee \neg\mathbf{shirt})$ satisfiable?

- a class of rather low-level kind of problems:
 - propositional variables only, e.g. either hold (true) or not (false)
 - logic operators \neg , \vee , \wedge , actually restricted to conjunctive normal form (CNF)
 - but no quantifiers such as “for all such things”, or “there is one such thing”
 - can we find an assignment of the variables to true or false, such that a set of clauses is satisfied simultaneously
- theory: it is **the** standard NP complete problem [Cook’70]
- encoding: how to get your problem into CNF
- simplifying: how can the problem or the CNF be simplified (preprocessing)
- solving: how to implement fast solvers

- Davis and Putnam procedure
 - **DP**: elimination procedure [DavisPutnam'60]
 - **DPLL**: splitting [DavisLogemannLoveland'62]
- modern SAT solvers are mostly based on **DPLL**, actually **CDCL**
CDCL = Conflict Driven Clause Learning
 - learning: GRASP [MarquesSilvaSakallah'96], ReISAT [BayardoSchrag'97]
 - watched literals, VSIDS: [mz]Chaff [MoskewiczMadiganZhaoZhangMalik-DAC'01]
 - improved heuristics: MiniSAT [EénSörensson-SAT'03] actually version from 2005
- preprocessing is still a hot topic:
 - most practical solvers use SatELite style preprocessing [EénBiere'05] **DP**
 - *Inprocessing* in PrecoSAT, Lingeling, CryptoMiniSAT, ... [JärvisaloHeuleBiere'12]

- satisfiability solving for first order formulae
 - extension of SAT but interpreted over fixed theories
 - originally without quantifiers *but* quantifiers are important
 - fully automatic decision procedures which also can provide models
- theories of interest
 - equality, uninterpreted functions
 - real / integer arithmetic
 - bit-vectors, arrays
- particularly important are bit-vectors with and without arrays for HW/SW verification
 - our SMT solver Boolector ranked #1 in these categories (SMT 2008/2009/2012)

- *bounded model checking* in electronic design automation (EDA)
 - routinely used for falsification in all major design houses
 - unbounded extensions also use SAT, e.g. *sequential* equivalence checking
- SAT as working horse in *static software verification*
- static device driver verification at Microsoft (SLAM, SDV)
 - predicate abstraction with SMT solvers
 - spurious counter example checking
- *software configuration*, e.g. Eclipse IDE ships with SAT4J
- cryptanalysis and other *combinatorial problems* (bio-informatics)

- QBF can be seen as extension to SAT:
 - existentially quantified variables as in SAT
 - but some variables can be universally quantified
- QBF is the *the classical* PSPACE complete problem
 - as SAT is *the* NP-complete problem
 - two other important PSPACE complete problems:
 - * (Propositional) Linear Temporal Logic (LTL) satisfiability
 - * symbolic model checking / symbolic reachability

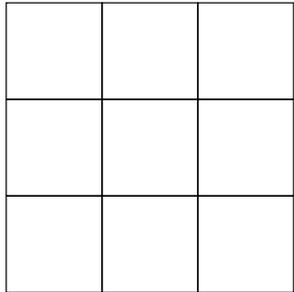









$$\forall \text{tie} [\exists \text{shirt} [\overbrace{(\text{tie} \vee \text{shirt}) \wedge (\neg \text{tie} \vee \neg \text{shirt})}^{\text{tie} \neq \text{shirt}}]] \quad \neq \quad \exists \text{tie} [\forall \text{shirt} [\overbrace{(\text{tie} \vee \text{shirt}) \wedge (\neg \text{tie} \vee \neg \text{shirt})}^{\text{tie} \neq \text{shirt}}]]$$

satisfiable
unsatisfiable

- semantics given as **expansion** of quantifiers

$$\exists x[f] \equiv f[0/x] \vee f[1/x] \qquad \forall x[f] \equiv f[0/x] \wedge f[1/x]$$

- expansion as translation from SAT to QBF is exponential
 - SAT problems have only existential quantifiers
 - expansion of universal quantifier can double formula size
- large number of different approaches to solve QBF versus “mono-culture” in SAT
 - scalability for practically interesting problem still an issue
 - nevertheless first real applications appear, e.g. black-box equivalence checking
 - steady progress: currently fastest solvers DepQBF and Qube

s_0  s_1  s_2  s_3  s_4  s_5  s_6  s_7  s_8  s_9 

$$\begin{aligned}
 \not\models & \quad \forall s_0[\text{empty}(s_0) \rightarrow \\
 & \quad \exists x_1[\text{circle}(s_0, x_1, s_1) \wedge \quad x_i, y_i \text{ plays (4 bits each)} \\
 & \quad \forall y_2[\text{cross}(s_1, y_2, s_2) \rightarrow \\
 & \quad \quad \exists x_3[\text{circle}(s_2, x_3, s_3) \wedge \\
 & \quad \quad \forall y_4[\text{cross}(s_3, y_4, s_4) \rightarrow \\
 & \quad \quad \quad \exists x_5[\text{circle}(s_4, x_5, s_5) \wedge \\
 & \quad \quad \quad \forall y_6[\text{cross}(s_5, y_6, s_6) \rightarrow \\
 & \quad \quad \quad \quad \exists x_7[\text{circle}(s_6, x_7, s_7) \wedge \\
 & \quad \quad \quad \quad \forall y_8[\text{cross}(s_7, y_8, s_8) \rightarrow \\
 & \quad \quad \quad \quad \quad \exists x_9[\text{circle}(s_8, x_9, s_9) \wedge \text{win}_{\text{circle}}(s_9)]]]]]]]]]]]]]]
 \end{aligned}$$

s_i configurations

(9 × 3 bits each)

original code

```
if(!a && !b) h();  
else if(!a) g();  
else f();
```



```
if(!a) {  
    if(!b) h();  
    else g();  
} else f();
```



optimized code

```
if(a) f();  
else if(b) g();  
else h();
```



```
if(a) f();  
else {  
    if(!b) h();  
    else g();  
}
```

How to check that these two versions are equivalent?

1. represent procedures as *independent* boolean variables

original :=

if $\neg a \wedge \neg b$ **then** h
else if $\neg a$ **then** g
else f

optimized :=

if a **then** f
else if b **then** g
else h

2. compile if-then-else chains into boolean formulae

$$\text{compile}(\mathbf{if\ } x \mathbf{\ then\ } y \mathbf{\ else\ } z) \equiv (x \wedge y) \vee (\neg x \wedge z)$$

3. check equivalence of boolean formulae

$$\text{compile}(\mathit{original}) \Leftrightarrow \text{compile}(\mathit{optimized})$$

$$\begin{aligned} \textit{original} &\equiv \mathbf{if} \neg a \wedge \neg b \mathbf{ then } h \mathbf{ else if } \neg a \mathbf{ then } g \mathbf{ else } f \\ &\equiv (\neg a \wedge \neg b) \wedge h \vee \neg(\neg a \wedge \neg b) \wedge \mathbf{if} \neg a \mathbf{ then } g \mathbf{ else } f \\ &\equiv (\neg a \wedge \neg b) \wedge h \vee \neg(\neg a \wedge \neg b) \wedge (\neg a \wedge g \vee a \wedge f) \end{aligned}$$

$$\begin{aligned} \textit{optimized} &\equiv \mathbf{if } a \mathbf{ then } f \mathbf{ else if } b \mathbf{ then } g \mathbf{ else } h \\ &\equiv a \wedge f \vee \neg a \wedge \mathbf{if } b \mathbf{ then } g \mathbf{ else } h \\ &\equiv a \wedge f \vee \neg a \wedge (b \wedge g \vee \neg b \wedge h) \end{aligned}$$

$$(\neg a \wedge \neg b) \wedge h \vee \neg(\neg a \wedge \neg b) \wedge (\neg a \wedge g \vee a \wedge f) \Leftrightarrow a \wedge f \vee \neg a \wedge (b \wedge g \vee \neg b \wedge h)$$


Reformulate it as a satisfiability (SAT) problem:

Is there an assignment to a, b, f, g, h ,
which results in different evaluations of *original* and *optimized*?


or equivalently:

Is the boolean formula $\text{compile}(\textit{original}) \not\leftrightarrow \text{compile}(\textit{optimized})$ satisfiable?

such an assignment would provide an easy to understand counterexample



$$b \vee a \wedge c$$



$$(a \vee b) \wedge (b \vee c)$$

equivalent?

$$b \vee a \wedge c$$

\Leftrightarrow

$$(a \vee b) \wedge (b \vee c)$$

Definition formula in *Conjunctive Normal Form* (CNF) is a conjunction of clauses

$$C_1 \wedge C_2 \wedge \dots \wedge C_n$$

each *clause* C is a disjunction of literals

$$C = L_1 \vee \dots \vee L_m$$

and each *literal* is either a plain variable x or a negated variable \bar{x} .

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

Note 1: two notions for negation: in \bar{x} and \neg as in $\neg x$ for denoting negation.

Note 2: original SAT problem is actually formulated for CNF

Note 3: solvers (mostly) expect CNF as input

- common ASCII file format of SAT solvers, used by SAT competitions
- variables are represented as natural numbers, literals as integers
- header “p cnf <vars> <clauses>”, comment lines start with “c”

In order to show the *validity* of

$$b \vee a \wedge c \Leftrightarrow (a \vee b) \wedge (b \vee c)$$

negate,

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

simplify and show *unsatisfiability* of

$$\neg b \wedge (\neg a \vee \neg c) \wedge (a \vee b) \wedge (b \vee c)$$

c the first two lines are comments

c ex1.cnf: a=1, b=2, c=3

p cnf 3 4

-2 0

-1 -3 0

1 2 0

2 3 0

```
// compile with: gcc -o ex1 ex1.c picosat.o
#include "picosat.h"
#include <stdio.h>
int main () {
    int res;
    picosat_init ();

    picosat_add (-2); picosat_add (0);
    picosat_add (-1); picosat_add (-3); picosat_add (0);
    picosat_add (1); picosat_add (2); picosat_add (0);
    picosat_add (2); picosat_add (3); picosat_add (0);

    res = picosat_sat (-1);

    if (res == 10) printf ("s SATISFIABLE\n");
    else if (res == 20) printf ("s UNSATISFIABLE\n");
    else printf ("s UNKNOWN\n");

    picosat_reset ();
    return res;
}
```

assume invalid equivalence resp. implication: $(a \vee b) \Rightarrow (a \text{ xor } b)$
 its negation $(a \vee b) \wedge (a = b)$
 as CNF $(a \vee b) \wedge (\neg a \vee b) \wedge (\neg b \vee a)$

```
c ex2.cnf: a=1,b=2
p cnf 2 3
1 2 0
-1 2 0
-2 1 0
```


SAT solver then allows to extract one satisfying assignment:

```
$ picosat ex2.cnf
s SATISFIABLE
v 1 2 0
```

this is the *only one* since “assuming” the opposite values individually is UNSAT

```
$ picosat ex2.cnf -a -1; picosat ex2.cnf -a -2
s UNSATISFIABLE
s UNSATISFIABLE
```

CNF



$$\begin{aligned}
 & o \wedge \\
 & (x \leftrightarrow a \wedge c) \wedge \\
 & (y \leftrightarrow b \vee x) \wedge \\
 & (u \leftrightarrow a \vee b) \wedge \\
 & (v \leftrightarrow b \vee c) \wedge \\
 & (w \leftrightarrow u \wedge v) \wedge \\
 & (o \leftrightarrow y \oplus w)
 \end{aligned}$$

$$o \wedge (x \rightarrow a) \wedge (x \rightarrow c) \wedge (x \leftarrow a \wedge c) \wedge \dots$$

$$o \wedge (\bar{x} \vee a) \wedge (\bar{x} \vee c) \wedge (x \vee \bar{a} \vee \bar{c}) \wedge \dots$$

1. generate a new variable x_s for each non input circuit signal s
2. for each gate produce complete input / output constraints as clauses
3. collect all constraints in a big conjunction

the transformation is *satisfiability equivalent*:

the result is satisfiable iff and only the original formula is satisfiable

not equivalent to the original formula: it has new variables

just *project* satisfying assignment onto the original variables

Negation: $x \leftrightarrow \bar{y} \Leftrightarrow (x \rightarrow \bar{y}) \wedge (\bar{y} \rightarrow x)$
 $\Leftrightarrow (\bar{x} \vee \bar{y}) \wedge (y \vee x)$


Disjunction: $x \leftrightarrow (y \vee z) \Leftrightarrow (y \rightarrow x) \wedge (z \rightarrow x) \wedge (x \rightarrow (y \vee z))$
 $\Leftrightarrow (\bar{y} \vee x) \wedge (\bar{z} \vee x) \wedge (\bar{x} \vee y \vee z)$

Conjunction: $x \leftrightarrow (y \wedge z) \Leftrightarrow (x \rightarrow y) \wedge (x \rightarrow z) \wedge ((y \wedge z) \rightarrow x)$
 $\Leftrightarrow (\bar{x} \vee y) \wedge (\bar{x} \vee z) \wedge ((y \wedge z) \vee x)$
 $\Leftrightarrow (\bar{x} \vee y) \wedge (\bar{x} \vee z) \wedge (\bar{y} \vee \bar{z} \vee x)$

Equivalence: $x \leftrightarrow (y \leftrightarrow z) \Leftrightarrow (x \rightarrow (y \leftrightarrow z)) \wedge ((y \leftrightarrow z) \rightarrow x)$
 $\Leftrightarrow (x \rightarrow ((y \rightarrow z) \wedge (z \rightarrow y))) \wedge ((y \leftrightarrow z) \rightarrow x)$
 $\Leftrightarrow (x \rightarrow (y \rightarrow z)) \wedge (x \rightarrow (z \rightarrow y)) \wedge ((y \leftrightarrow z) \rightarrow x)$
 $\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge ((y \leftrightarrow z) \rightarrow x)$
 $\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge (((y \wedge z) \vee (\bar{y} \wedge \bar{z})) \rightarrow x)$
 $\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge ((y \wedge z) \rightarrow x) \wedge ((\bar{y} \wedge \bar{z}) \rightarrow x)$
 $\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge (\bar{y} \vee \bar{z} \vee x) \wedge (y \vee z \vee x)$

- goal is smaller CNF less variables, less clauses, so easier to solve (?!)
- extract multi argument operands to remove variables for intermediate nodes
- half of AND, OR node constraints/clauses can be removed for *unnegated* nodes [PlaistedGreenbaum'86]
 - node occurs negated if it has an ancestor which is a negation
 - half of the constraints determine parent assignment from child assignment
 - those are unnecessary if node is not used negated
 - those have to be carefully applied to DAG structure
- further structural circuit optimizations ...

CNF Blocked Clause Elimination simulates many encoding / circuit optimizations



- encoding directly into CNF is hard, so we use intermediate levels:
 1. application level
 2. bit-precise semantics world-level operations: bit-vector theory
 3. bit-level representations such as AIGs or vectors of AIGs
 4. CNF
- encoding application level formulas into word-level: as generating machine code
- word-level to bit-level: bit-blasting similar to hardware synthesis
- encoding “logical” constraints is another story

addition of 4-bit numbers x, y with result s also 4-bit: $s = x + y$

$$[s_3, s_2, s_1, s_0]_4 = [x_3, x_2, x_1, x_0]_4 + [y_3, y_2, y_1, y_0]_4$$

$$[s_3, \cdot]_2 = \text{FullAdder}(x_3, y_3, c_2)$$

$$[s_2, c_2]_2 = \text{FullAdder}(x_2, y_2, c_1)$$

$$[s_1, c_1]_2 = \text{FullAdder}(x_1, y_1, c_0)$$

$$[s_0, c_0]_2 = \text{FullAdder}(x_0, y_0, \text{false})$$


where

$$[s, o]_2 = \text{FullAdder}(x, y, i) \quad \text{with}$$

$$s = x \text{ xor } y \text{ xor } i$$

$$o = (x \wedge y) \vee (x \wedge i) \vee (y \wedge i) = ((x + y + i) \geq 2)$$

- widely adopted bit-level intermediate representation
 - see for instance our AIGER format <http://fmv.jku.at/aiger>
 - used in Hardware Model Checking Competition (HWMCC)
 - also used in the *structural track* in SAT competitions
 - many companies use similar techniques
- basic logical operators: *conjunction* and *negation*
- DAGs: nodes are conjunctions, negation/sign as *edge attribute*
bit stuffing: signs are compactly stored as LSB in pointer
- automatic sharing of isomorphic graphs, constant time (peep hole) simplifications
- *or even* SAT sweeping, full reduction, etc ... see ABC system from Berkeley



negation/sign are edge attributes
not part of node

$$x \text{ xor } y \equiv (\bar{x} \wedge y) \vee (x \wedge \bar{y}) \equiv \overline{\overline{(\bar{x} \wedge y)} \wedge \overline{(x \wedge \bar{y})}}$$

```
typedef struct AIG AIG;


struct AIG
{
    enum Tag tag;           /* AND, VAR */
    void *data[2];
    int mark, level;       /* traversal */
    AIG *next;             /* hash collision chain */
};

#define sign_aig(aig) (1 & (unsigned) aig)
#define not_aig(aig) ((AIG*)(1 ^ (unsigned) aig))
#define strip_aig(aig) ((AIG*)(~1 & (unsigned) aig))
#define false_aig ((AIG*) 0)
#define true_aig ((AIG*) 1)
```


assumption for correctness:


```
sizeof(unsigned) == sizeof(void*)
```

4-bit adder




8-bit adder





bit-vector of length 16 shifted by bit-vector of length 4




- Tseitin's construction suitable for most kinds of "model constraints"
 - assuming simple operational semantics: encode an interpreter
 - small domains: *one-hot encoding* large domains: *binary encoding*
- harder to encode *properties* or additional *constraints*
 - temporal logic / fix-points
 - environment constraints
- example for fix-points / recursive equations: $x = (a \vee y), \quad y = (b \vee x)$
 - has unique *least* fix-point $x = y = (a \vee b)$
 - and unique *largest* fix-point $x = y = true$ but unfortunately
 - only largest fix-point can be (directly) encoded in SAT otherwise need ASP

- given a set of literals $\{l_1, \dots, l_n\}$
 - constraint the *number* of literals assigned to *true*
 - $|\{l_1, \dots, l_n\}| \geq k$ or $|\{l_1, \dots, l_n\}| \leq k$ or $|\{l_1, \dots, l_n\}| = k$

- multiple encodings of cardinality constraints
 - naïve encoding exponential: *at-most-two* quadratic, *at-most-three* cubic, etc.
 - quadratic $O(k \cdot n)$ encoding goes back to Shannon
 - linear $O(n)$ parallel counter encoding [Sinz'05]
 - for an $O(n \cdot \log n)$ encoding see Prestwich's chapter in our Handbook of SAT

- generalization *Pseudo-Boolean* constraints (PB), e.g. $2 \cdot \bar{a} + \bar{b} + c + \bar{d} + 2 \cdot e \geq 3$
 actually used to handle MaxSAT in SAT4J for configuration in Eclipse

$$2 \leq |\{l_1, \dots, l_9\}| \leq 3$$



“then” edge downward, “else” edge to the right

```
// compile with: gcc -o ex2 ex2.c picosat.o
#include "picosat.h"
#include <stdio.h>
#include <assert.h>
int main () {
    int res, a, b;
    picosat_init ();
    picosat_add (1); picosat_add (2); picosat_add (0);
    picosat_add (-1); picosat_add (2); picosat_add (0);
    picosat_add (-2); picosat_add (1); picosat_add (0);
    assert (picosat_sat (-1) == 10); // SATISFIABLE
    a = picosat_deref (1); b = picosat_deref (2);
    printf ("v %d %d\n", a*1, b*2);
    picosat_assume (-a*1); assert (picosat_sat (-1) == 20); //UNSAT
    picosat_assume (-b*2); assert (picosat_sat (-1) == 20); //UNSAT
    return res;
}
```


```
static void block_current_solution (void) {
    int max_idx = picosat_variables (), i;

    // since 'picosat_add' resets solutions
    // need to store it first:
    signed char * sol = malloc (max_idx + 1);
    memset (sol, 0, max_idx + 1);

    for (i = 1; i <= max_idx; i++)
        sol[i] = (picosat_deref (i) > 0) ? 1 : -1;


    for (i = 1; i <= max_idx; i++)
        picosat_add ((sol[i] < 0) ? i : -i);
    picosat_add (0);

    free (sol);
}
```



- two ways to implement incremental SAT solvers
 - push / pop as in SMT solvers partial support in SATIRE, zChaff, PicoSAT
 - * clauses associated with context and pushed / popped in a stack like manner
 - * pop discards clauses of current context
 - most common: **assumptions** [ClaessenSörensson'03] [EénSörensson'03]
 - * allows to use set of literals as assumptions
 - * force SAT solver to first pick assumption as decisions
 - * more flexible, since assumptions can be reused
 - * assumptions are only valid for the next SAT call
- failed assumptions: sub set of assumptions inconsistent with CNF

- goal: reduce size of bit-vector constants in satisfying assignments
- refinement approach: for each bit-vector variable only use an “effective width”
 - example: 4-bit vector $[x_3, x_2, x_1, x_0]$ and effective width 2 use $[x_1, x_1, x_1, x_0]$
 - either encode from scratch with x_3 and x_2 replaced by x_1 (1)
 - or add $x_3 = x_1$ and $x_2 = x_1$ after push (2)
 - or add $a_x^2 \rightarrow x_3 = x_1$ and $a_x^2 \rightarrow x_2 = x_1$ and assume fresh literal a_x^2 (3)
- if satisfiable then a solution with small constants has been found
 otherwise increase eff. width of bit-vectors where it was used to derive UNSAT
 under-approximations not used then formula UNSAT “used” = “failed assumption”
- in (3) constraints are removed by forcing assumptions to the opposite value
 by adding a unit clause, e.g. $\neg a_x^2$ in next iteration



- clausal core (or unsatisfiable sub set) of an unsatisfiable formula
 - clauses used to derive the empty clause
 - may include not only original but also learned clauses
 - similar application as in previous under-approximation example
 - but also useful for diagnosis of inconsistencies
- variable core
 - sub set of variables occurring in clauses of a clausal core
- these cores are not unique and not necessary minimimal
- minimal unsatisfiable sub set (MUS) = clausal core where no clause can be removed

- PicoMUS is a MUS extractor based on PicoSAT
 - uses several rounds of clausal core extraction for preprocessing
 - then switches to assumption based core minimization using `picosat_failed_assumptions`
 - source code serves as a good example on how to use cores / assumptions
- new MUS track in SAT 2011 competition
 - with high- and low-level MUS sub tracks

```

c ex3.cnf          $ picosat ex3.cnf -c core
p cnf 6 10         s UNSATISFIABLE
1 2 3 0           $ cat core
1 2 -3 0          p cnf 6 9
1 -2 3 0          2 3 1 0
1 -2 -3 0         2 -3 1 0
4 5 6 0           -2 3 1 0
4 5 -6 0          -2 -3 1 0
4 -5 6 0          6 5 4 0
4 -5 -6 0         5 -6 4 0
-1 -4 0           6 4 -5 0
1 4 0             4 -6 -5 0
                  -1 -4 0

c ex4.cnf         $ picomus ex4.cnf mus
p cnf 6 11        s UNSATISFIABLE
1 2 3 0           $ cat mus
1 2 -3 0          p cnf 6 6
1 -2 3 0          1 2 3 0
1 -2 -3 0         1 2 -3 0
4 5 6 0           1 -2 3 0
4 5 -6 0          1 -2 -3 0
4 -5 6 0          -1 4 0
4 -5 -6 0         -1 -4 0
-1 -4 0
-1 4 0
-1 -4 0

```

- core extraction in PicoSAT is based on tracing proofs
 - enabled by `picosat_enable_trace_generation`
 - maintains “dependency graph” of learned clauses
 - kept in memory, so fast core generation
- traces can also be written to disk in various formats
 - RUP format by Allen Van Gelder (SAT competition)
 - or format of TraceCheck tool
- TraceCheck can check traces for correctness
 - orders clauses and antecedents to generate and check resolution proof
 - (binary) resolution proofs can be dumped

same as DIMACS except that we have additional quantifiers:

```
c SAT
p cnf 3 4
a 1 0
e 2 3 0
-1 -2 3 0
-1 2 -3 0
1 2 3 0
1 -2 -3 0
```

```
c UNSAT
p cnf 4 8
a 1 2 0
e 3 4 0
-1 -3 4 0
-1 3 -4 0
1 3 4 0
1 -3 -4 0
-2 -3 4 0
-2 3 -4 0
2 3 4 0
2 -3 -4 0
```

```
/* Create and initialize solver instance. */
QDPLL *qdpll_create (void);

/* Delete and release all memory of solver instance. */
void qdpll_delete (QDPLL * qdpll);

/* Ensure var table size to be at least 'num'. */
void qdpll_adjust_vars (QDPLL * qdpll, VarID num);

/* Open a new scope, where variables can be added by 'qdpll_add'.
   Returns nesting of new scope.
   Opened scope can be closed by adding '0' via 'qdpll_add'.
   NOTE: will fail if there is an opened scope already.
*/
unsigned int qdpll_new_scope (QDPLL * qdpll, QDPLLQuantifierType qtype);

/* Add variables or literals to clause or opened scope.
   If scope is opened, then 'id' is interpreted as a variable ID,
   otherwise 'id' is interpreted as a literal.
   NOTE: will fail if a scope is opened and 'id' is negative.
*/
void qdpll_add (QDPLL * qdpll, LitID id);

/* Decide formula. */
QDPLLResult qdpll_sat (QDPLL * qdpll);
```

- 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'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: $\delta' = \delta \cdot \frac{1}{f}$ δ typically increment by 5% – 11%
- use floating point representation of score
- “rescore” to avoid overflow in regular intervals
- EVSIDS linearly related to NVSIDS

- VSIDS score can be normalized to the interval $[0,1]$ as follows:
 - pick a decay factor f per conflict: typically $f = 0.9$
 - each variable is **punished by this decay factor** at every conflict
 - if a variable is **involved in conflict**, add $1 - f$ to score
 - s old score of one fixed variable before conflict, s' new score after conflict

$$s, f \leq 1, \quad \text{then} \quad s' \leq \overbrace{s \cdot f}^{\text{decay in any case}} + \underbrace{1 - f}_{\text{increment if involved}} \leq f + 1 - f = 1$$

- recomputing score of all variables at each conflict is costly
 - linear in the number of variables, e.g. millions
 - particularly, because number of involved variables \ll number of variables

consider again only one variable with score sequence s_n resp. S_n

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

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

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

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


[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
- mixed results as other variants VMTF, CMTF (var/clause move to front)

- 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

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
```

[Knuth'12]

$$(u_1, v_1) := (1, 1)$$

$$(u_{n+1}, v_{n+1}) := (u_n \ \& \ -u_n = v_n ? (u_n + 1, 1) : (u_n, 2v_n))$$


$(1, 1), (2, 1), (2, 2), (3, 1), (4, 1), (4, 2), (4, 4), (5, 1), \dots$


- phase assignment / direction heuristics:
 - 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
 - empirically works nice together with phase saving


[Van Der Tak, Heule, Ramos POS'11]

- in general *restarting does not change much*: since phases and scores saved
- assignment after restart can only differ if
 - before restarting
 - there is a decision literal d assigned on the trail
 - with smaller score than the next decision n on the priority queue
- in this situation backtrack **only** to decision level of d
 - simple to compute, particularly if decisions are saved separately
 - allows to skip many redundant backtracks
 - allows much higher restart frequency,
e.g. base interval 10 for reluctant doubling sequence (Luby)


- keeping all learned clauses slows down BCP kind of quadratically
 - so SATO and ReISAT 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* (FIFO) clauses
 - after reduction maximum number kept learned clauses is increased geometrically
- LBD (Glue) based (apriori!) prediction for usefulness [AudemardSimon’09]
 - LBD (Glue) = number of decision-levels in the learned clause
 - allows arithmetic increase of number of kept learned clauses
- freeze high PSM (dist. to phase assign.) clauses [AudemardLagniezMazureSais’11]









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







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




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




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



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




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



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

self subsuming resolution




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


- can be found by *graph traversal* in the reverse order of made assignments
 - *trail* respects this order
 - mark literals in conflict
 - traverse reasons of marked variables on trail in reverse order
- count *number unresolved variables* on current decision level
 - decrease counter if new reason / antecedent clause resolved
 - if counter=1 (only one unresolved marked variable left) then this node is a UIP
 - note, decision of current decision level is a UIP and thus a *sentinel*

```
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;
}
```





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




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


self subsuming resolution




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



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




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




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


algorithm of Allen Van Gelder in SAT'09 produces regular input resolution proofs directly

- 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

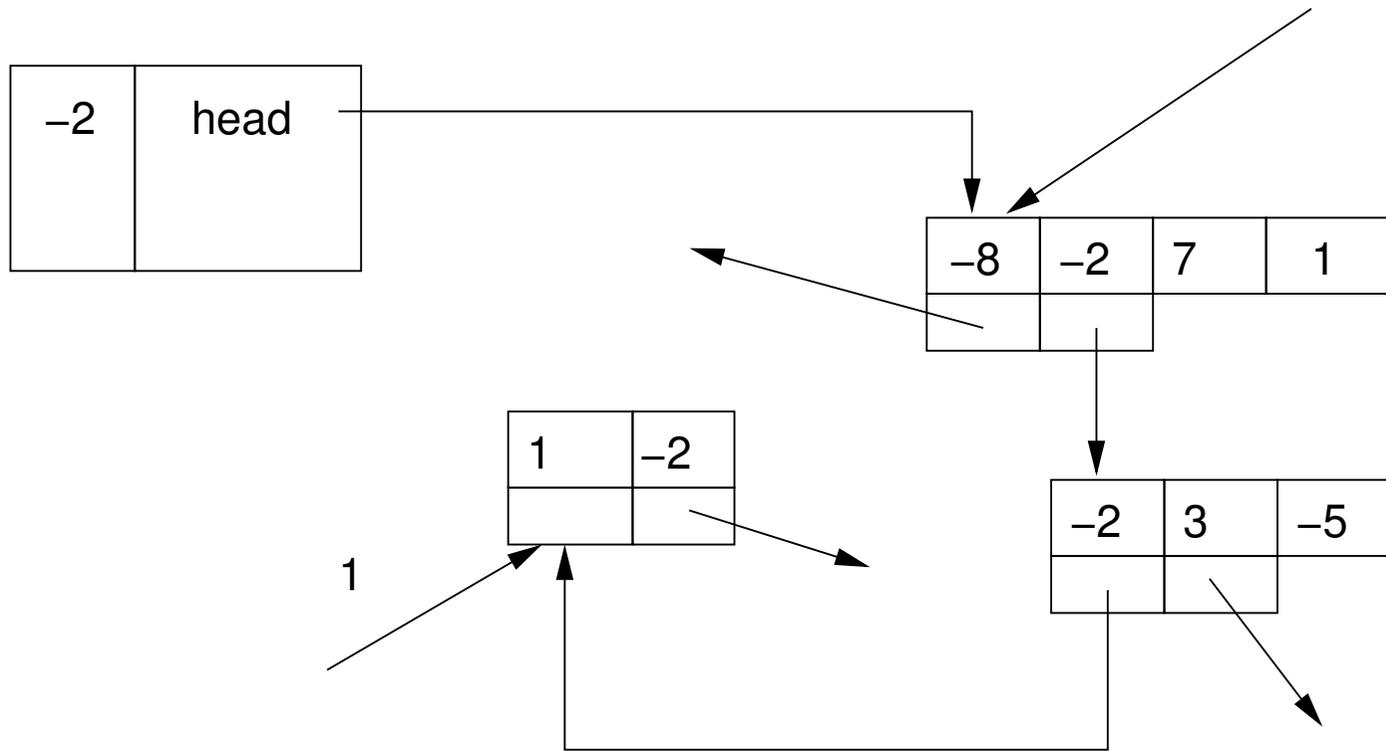




still seems to be best way for *real* sharing of clauses in multi-threaded solvers

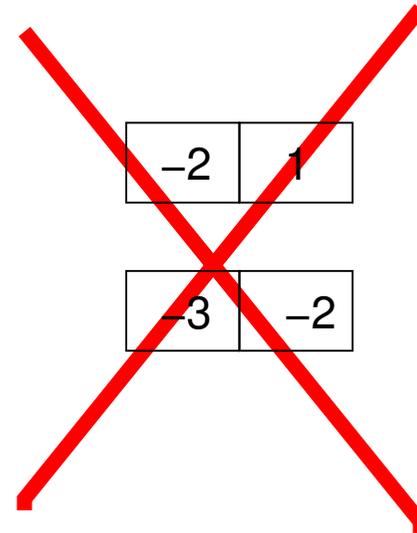
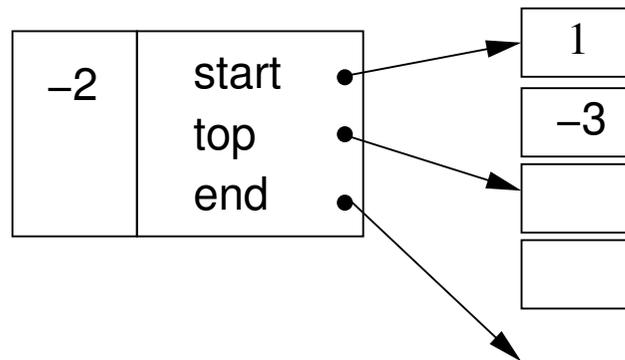


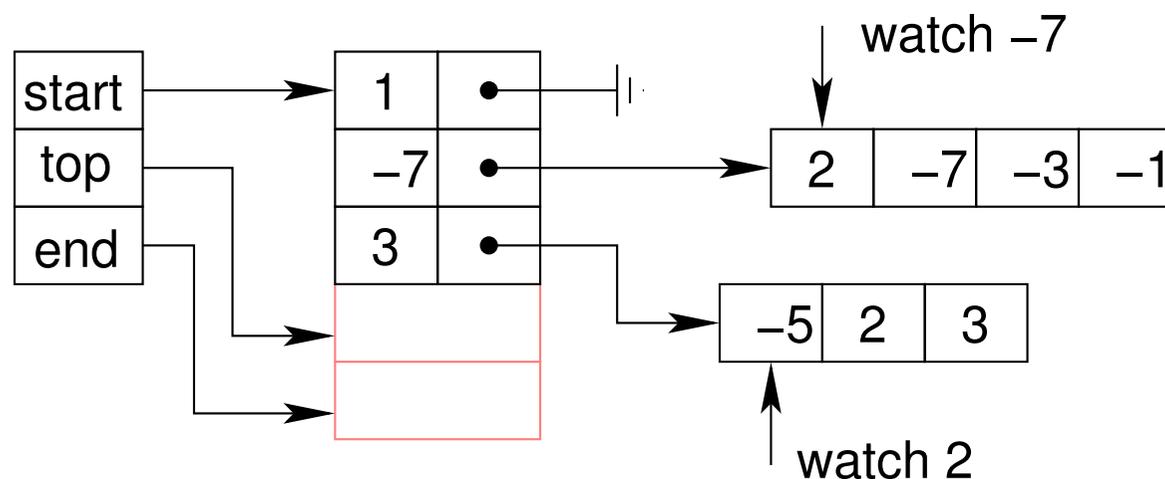
invariant: first two literals are watched



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
 - can be applied recursively (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]

[AnderssonBjessCookHanna DAC'02]

also in Oepir SAT solver, this is our reformulation

- for all literals l
 - for all clauses c in which l occurs (with this particular phase)
 - * assume the negation of all the other literals in c , assume l
 - * if BCP does not lead to a conflict continue with next literal in outer loop
 - if all clauses produced a conflict permanently assign $\neg l$

Correctness: Let $c = l \vee d$, assume $\neg d \wedge l$.

If this leads to a conflict $d \vee \neg l$ could be learned (but is not added to the CNF).

Self subsuming resolution with c results in d and c is removed.

If all such cases lead to a conflict, $\neg l$ becomes a pure literal.

Generalization of pure literals.

Given a partial assignment σ .

A clause of a CNF is “touched” by σ if it contains a literal assigned by σ .

A clause of a CNF is “satisfied” by σ if it contains a literal assigned to true by σ .

If all touched clauses are satisfied then σ is an “autarky”.

All clauses touched by an autarky can be removed.

Example: $(-1\ 2)(-1\ 3)(1\ -2\ -3)(2\ 5)\dots$ (more clauses without 1 and 3).

Then $\sigma = \{-1, -3\}$ is an autarky.

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

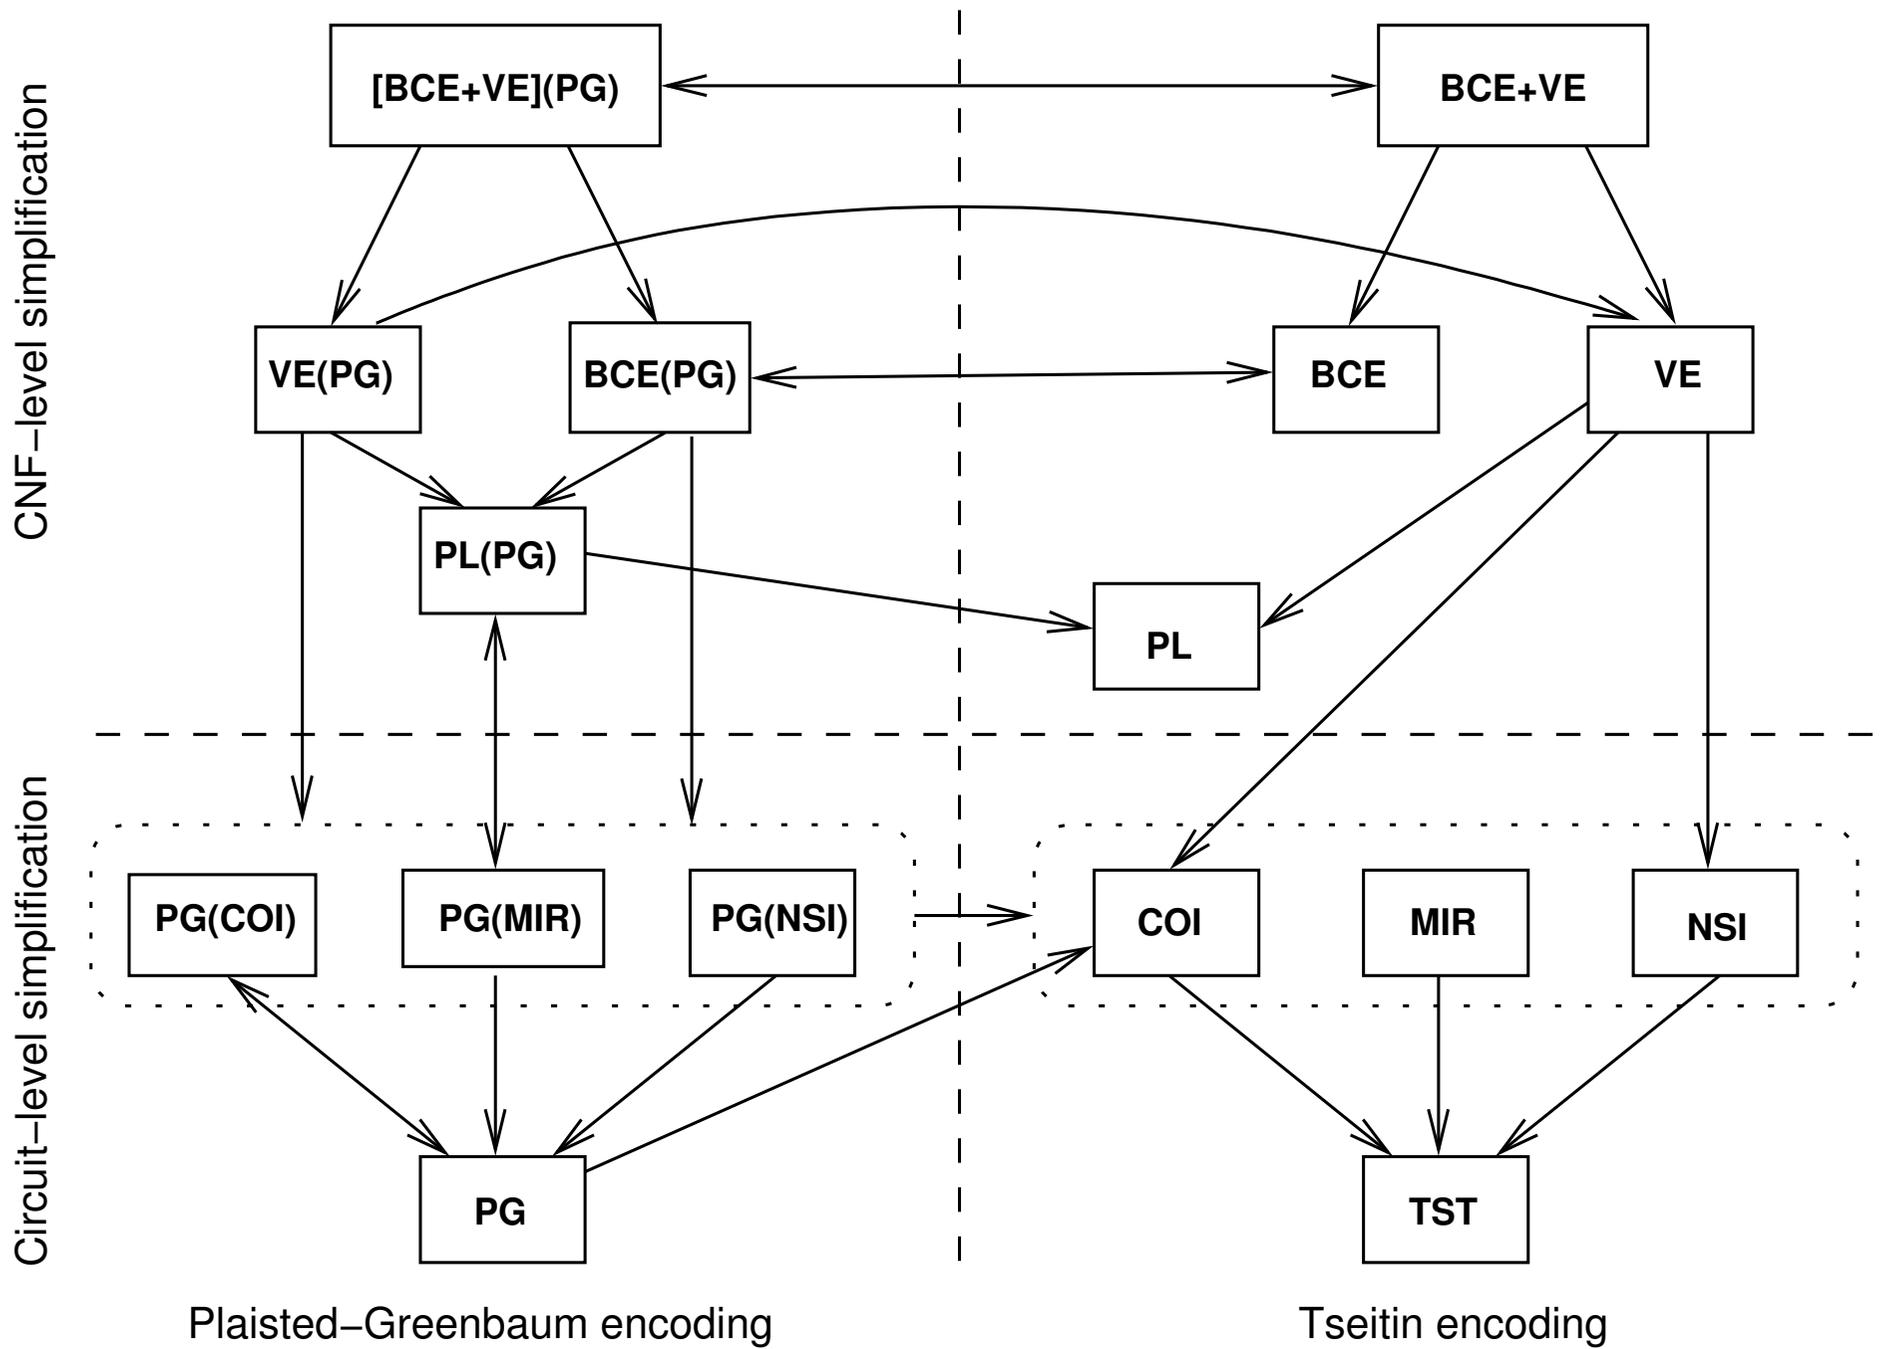
NSI Non-Shared Inputs reduction

PG Plaisted-Greenbaum polarity based encoding

TST standard Tseitin encoding

VE Variable-Elimination as in DP / Quantor / SATeLite

BCE Blocked-Clause-Elimination



Plaisted-Greenbaum encoding

Tseitin encoding

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]

$$\frac{\varphi[\rho \wedge C]\sigma}{\varphi \wedge C[\rho]\sigma}$$

STRENGTHEN

$$\frac{\varphi[\rho \wedge C]\sigma}{\varphi[\rho]\sigma}$$

FORGET

$$\frac{\varphi[\rho]\sigma}{\varphi[\rho \wedge C]\sigma} \mathcal{L}$$

LEARN

$$\frac{\varphi \wedge C[\rho]\sigma}{\varphi[\rho \wedge C]\sigma, l:C} \mathcal{W}$$

WEAKEN

\mathcal{L} is that $\boxed{\varphi \wedge \rho}$ and $\boxed{\varphi \wedge \rho \wedge C}$ are satisfiability-equivalent.

\mathcal{W} is that $\boxed{\varphi}$ and $\boxed{\varphi \wedge C}$ are satisfiability-equivalent.

“*resolution look-ahead*”

Clause R *asymmetric tautology* (AT) w.r.t. G iff $G \wedge \neg C$ refuted by BCP.

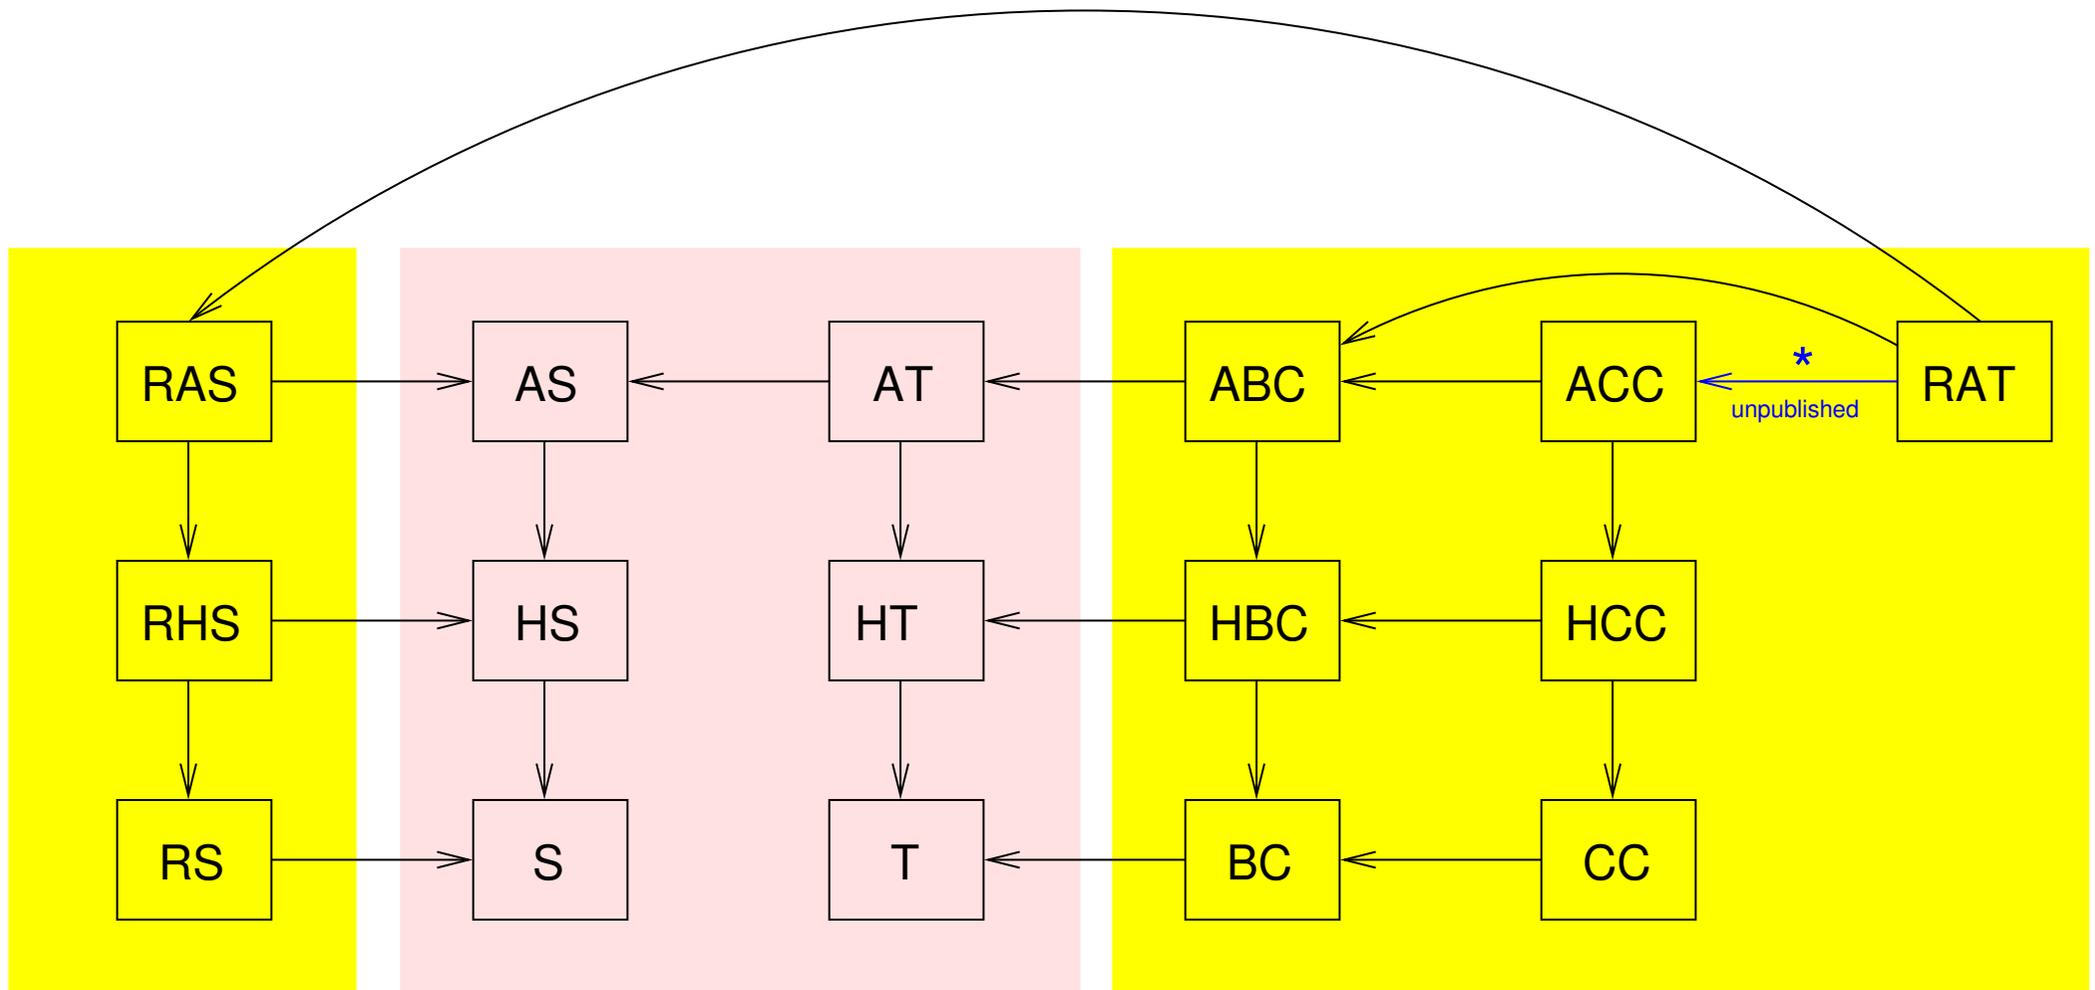
Given clause $C \in F$, $l \in C$.

Assume all resolvents R of C on l with clauses in F are AT w.r.t. $F \setminus \{C\}$.

Then C is called *resolution asymmetric tautology* (RAT) w.r.t. F on l .

In this case F is satisfiability equivalent to $F \setminus \{C\}$.

Inprocessing Rules with RAT simulate all techniques in current SAT solvers



logically equivalent

satisfiability equivalent

Asymmetric, **B**locked **C**lause, **C**overed **C**lause,
Hidden, **R**esolution, **S**ubsumed

- application level parallelism
 - run multiple “properties” at the same time
 - run multiple “engines” at the same time (streaming)
- portfolio solving
 - predict best solver through machine learning techniques SATzilla
 - run multiple solvers in parallel or sequentially (with/without “sharing”)
ManySAT, Plingeling, pfolio ...
- split search space
 - guiding path principle [ZhangBonacinaHsiang’96]
 - cube & conquer [HeuleKullmanWieringaBiere’11]
- low-level parallelism: parallelize BCP (threads, FPGA, GPU, ...) P complete

- use Look-Ahead at the top of the search tree, CDLC at the bottom
 - Look-Ahead solvers provide “good” global decisions but are slow
 - CDCL solvers are extremely fast based on local heuristics
- **when to switch from Look-Ahead to CDCL?**
 - limit decision height of Look-Ahead solver, e.g. 20 maximum tree height
 - avoid having too many branches closed by (slow) Look-Ahead
 - Concurrent Cube & Conquer runs CDCL and Look-Ahead in parallel
- open branches = cubes \Rightarrow solve in parallel
- solves hard instances, which none of the other approaches can