

# Using High Performance SAT and QBF Solvers

Armin Biere

Institute for Formal Models and Verification  
Johannes Kepler University  
Linz, Austria


Theorem Proving Tools for Program Analysis


Tutorial co-located with POPL 2011


Austin, Texas, USA


Monday, January 24, 2011

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**
  - 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

CDCL = Conflict Driven Clause Learning
  
- preprocessing is still a hot topic:
  - most practical solvers use SatELite style preprocessing [EénBiere'05] **DP**
  - *inprocessing* in fastest available solvers PrecoSAT, Lingeling, CryptoMiniSAT, ...



- 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 and arrays for HW/SW verification
  - our SMT solver Boolector ranked #1 in this category (SMT 2008 + 2009)

- *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)

MaxSAT

- 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

- look only for counter example made of  $k$  states    “ $k$ ” = the bound




- simple for *safety properties*     $p$  is invariantly true    (e.g.  $p = \neg B$ )

$$I(s_0) \wedge T(s_0, s_1) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \bigvee_{i=0}^k \neg p(s_i)$$

- harder for *liveness properties*     $p$  is eventually true

$$I(s_0) \wedge T(s_0, s_1) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \bigwedge_{i=0}^k \neg p(s_i) \wedge \exists l T(s_k, s_l)$$

- look only for counter example made of  $k$  states    " $k$ " = the bound













- simple for *safety properties*     $p$  is invariantly true    (e.g.  $p = \neg B$ )

$$I(s_0) \wedge T(s_0, s_1) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \bigvee_{i=0}^k \neg p(s_i)$$

- harder for *liveness properties*     $p$  is eventually true

$$I(s_0) \wedge T(s_0, s_1) \wedge \cdots \wedge T(s_{k-1}, s_k) \wedge \bigwedge_{i=0}^k \neg p(s_i) \wedge \bigvee_{l=0}^k T(s_k, s_l)$$

$s_0$ 

 $s_1$ 

 $s_2$ 

 $s_3$ 

 $s_4$ 

 $s_5$ 

 $s_6$ 

 $s_7$ 

 $s_8$ 

 $s_9$ 






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

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

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

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

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


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

$$\begin{aligned} \text{Negation:} \quad x \leftrightarrow \bar{y} &\Leftrightarrow (x \rightarrow \bar{y}) \wedge (\bar{y} \rightarrow x) \\ &\Leftrightarrow (\bar{x} \vee \bar{y}) \wedge (y \vee x) \end{aligned}$$


$$\begin{aligned} \text{Disjunction:} \quad x \leftrightarrow (y \vee z) &\Leftrightarrow (y \rightarrow x) \wedge (z \rightarrow x) \wedge (x \rightarrow (y \vee z)) \\ &\Leftrightarrow (\bar{y} \vee x) \wedge (\bar{z} \vee x) \wedge (\bar{x} \vee y \vee z) \end{aligned}$$

$$\begin{aligned} \text{Conjunction:} \quad x \leftrightarrow (y \wedge z) &\Leftrightarrow (x \rightarrow y) \wedge (x \rightarrow z) \wedge ((y \wedge z) \rightarrow x) \\ &\Leftrightarrow (\bar{x} \vee y) \wedge (\bar{x} \vee z) \wedge \overline{((y \wedge z) \vee x)} \\ &\Leftrightarrow (\bar{x} \vee y) \wedge (\bar{x} \vee z) \wedge (\bar{y} \vee \bar{z} \vee x) \end{aligned}$$

$$\begin{aligned} \text{Equivalence:} \quad x \leftrightarrow (y \leftrightarrow z) &\Leftrightarrow (x \rightarrow (y \leftrightarrow z)) \wedge ((y \leftrightarrow z) \rightarrow x) \\ &\Leftrightarrow (x \rightarrow ((y \rightarrow z) \wedge (z \rightarrow y))) \wedge ((y \leftrightarrow z) \rightarrow x) \\ &\Leftrightarrow (x \rightarrow (y \rightarrow z)) \wedge (x \rightarrow (z \rightarrow y)) \wedge ((y \leftrightarrow z) \rightarrow x) \\ &\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge ((y \leftrightarrow z) \rightarrow x) \\ &\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge (((y \wedge z) \vee (\bar{y} \wedge \bar{z})) \rightarrow x) \\ &\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge ((y \wedge z) \rightarrow x) \wedge ((\bar{y} \wedge \bar{z}) \rightarrow x) \\ &\Leftrightarrow (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee \bar{z} \vee y) \wedge (\bar{y} \vee \bar{z} \vee x) \wedge (y \vee z \vee x) \end{aligned}$$

- 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)
```

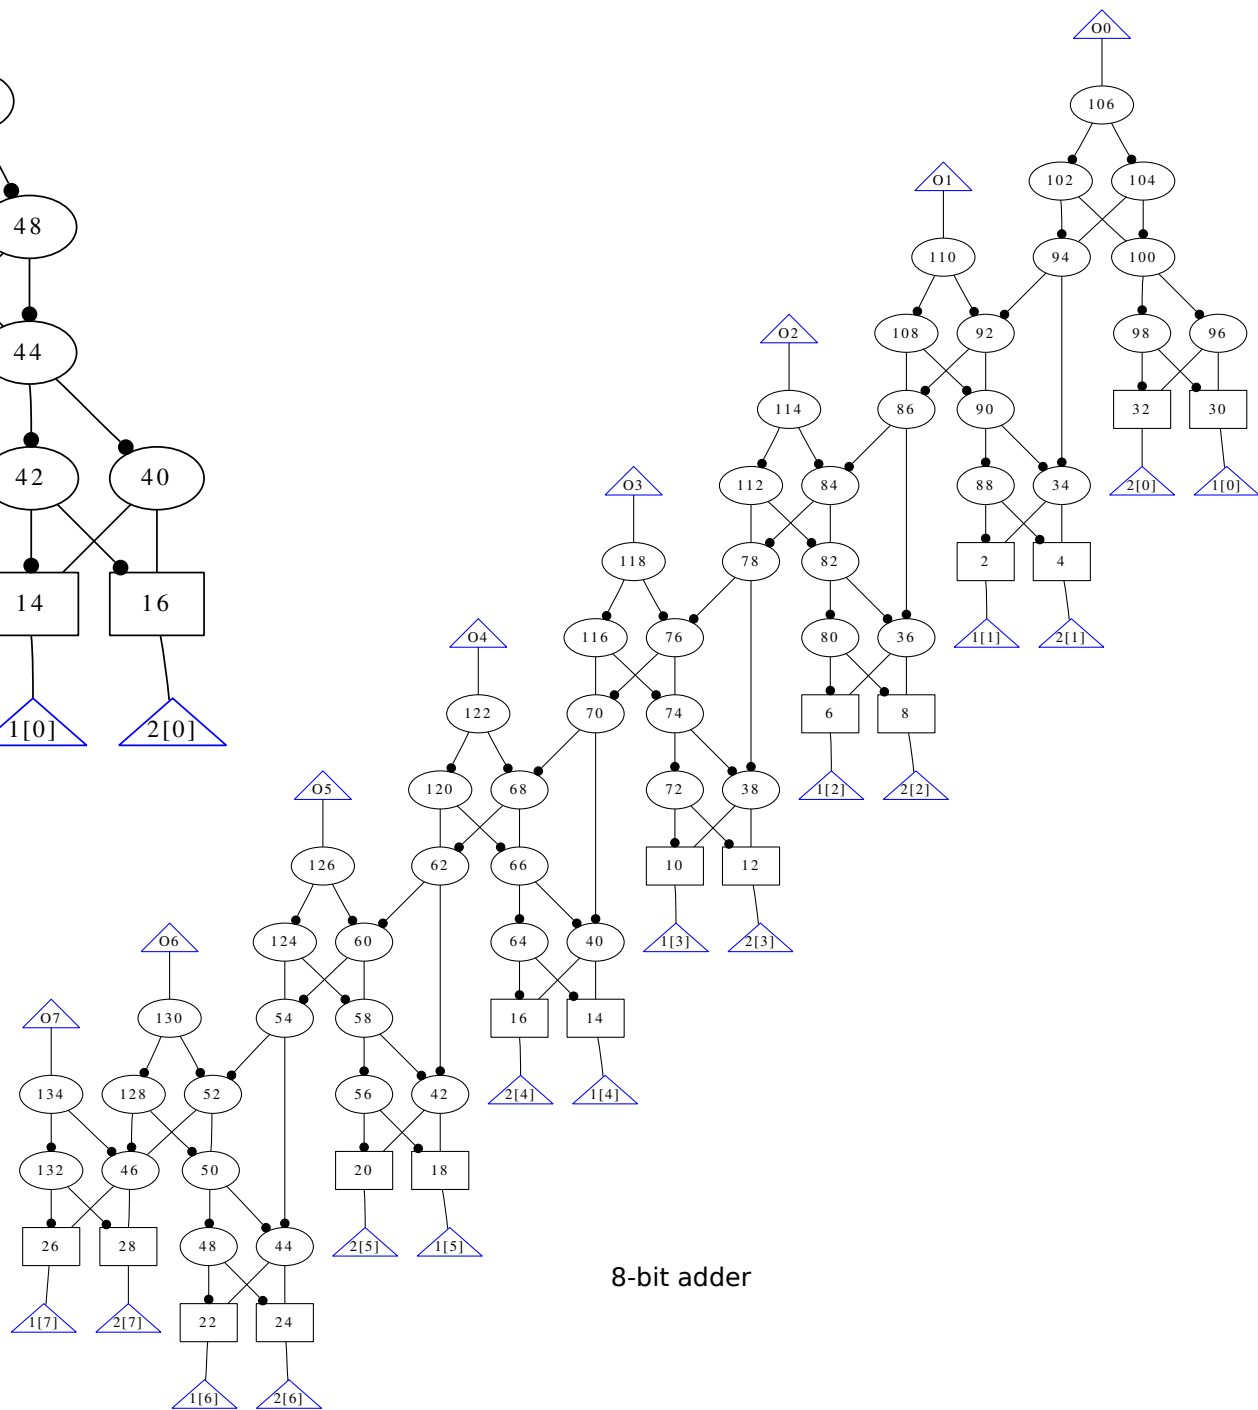
assumption for correctness:


```
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 minimal
- 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 this year's 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 *qdp11_create (void);

/* Delete and release all memory of solver instance. */
void qdp11_delete (QDPLL * qdp11);

/* Ensure var table size to be at least 'num'. */
void qdp11_adjust_vars (QDPLL * qdp11, VarID num);

/* Open a new scope, where variables can be added by 'qdp11_add'.
   Returns nesting of new scope.
   Opened scope can be closed by adding '0' via 'qdp11_add'.
   NOTE: will fail if there is an opened scope already.
*/
unsigned int qdp11_new_scope (QDPLL * qdp11, 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 qdp11_add (QDPLL * qdp11, LitID id);

/* Decide formula. */
QDPLLResult qdp11_sat (QDPLL * qdp11);
```

- extraction of “certificates”
  - satisfying assignment to outer-most existential variables
  - resp. in general skolem-functions for satisfiable instances
  - falsifying assignment to outer-most universal variables
- incremental QBF solving
- API for preprocessing / inprocessing
- beside steady progress **more scalability**

- SAT and QBF
- conjunctive normal form (CNF) and (Q)DIMACS format
- encoding of models and logical constraints
- PicoSAT, PicoMUS, TraceCheck, DepQBF
- examples and use cases
- APIs

- how SAT can handle millions of variables routinely
- whether there will still be progress in SAT
- AIGER format, proof (trace) formats
- the full API of PicoSAT (see `picosat.h` for more details)
- skipped API of other Solvers, in particular (Crypto)MiniSAT, SAT4J
- tricky issues with incremental pre/in-processing such as “freezing variables”
- compact QBF encodings [JussilaBiere’06]

and topics have attracted researchers from various disciplines. Logic, planning, scheduling, operations research and combinatorial optimization, as well as on the theme of complexity, and much more, they all are connected

SAT stems from actual solving: The increase in power of modern SAT solvers has been phenomenal. It has become the key enabling technology of both computer hardware and software. Bounded Model Checking (BMC) is now probably the most widely used model checking technique. BMC finds just satisfying instances of a Boolean formula obtained by unrolling a sequential circuit and its specification in linear temporal logic. The problem of software verification is a much more difficult problem on the frontier of computer science. A promising approach for languages like C with finite word-length integers is in BMC but with a decision procedure for the theory of bit-vectors. Modern SAT solvers and procedures for bit-vectors that I am familiar with ultimately make use of simple complex formulas.

More complicated theories, like linear real and integer arithmetic, are also being handled by SAT solvers. Most of them use powerful SAT solvers in an essential way.

SAT solving is a key technology for 21st century computer science. I expect that research on all theoretical and practical aspects of SAT solving will be extremely active in the coming years and will lead to many further advances in the field.

Edmund Clarke

Edmund Clarke is a University Professor of Computer Science and Professor of Electrical and Computer Engineering at Carnegie Mellon University, is one of the initiators and main contributors to the field of SAT solving, for which he also received the 2007 ACM Turing Award.

Edmund Clarke was one of the first researchers to realize that SAT solving has the potential to become one of the most important technologies in model checking.

# HANDBOOK of satisfiability



Editors:  
Armin Biere  
Marijn Heule  
Hans van Maaren  
Toby Walsh

Frontiers in Artificial Intelligence and Applications

# HANDBOOK of satisfiability

Editors:  
Armin Biere  
Marijn Heule  
Hans van Maaren  
Toby Walsh

ISBN 978-1-58603-929-5



9 781586 039295 >



blocked clause  $C \in F$ all clauses in  $F$  with  $\bar{l}$ fix a CNF  $F$ 

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

$$(a \vee b \vee l)$$

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

since all resolvents of  $C$  on  $l$  are tautological  $C$  can be removed

**Proof**

assignment  $\sigma$  satisfying  $F \setminus C$  but not  $C$

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

**COI** Cone-of-Influence reduction

**MIR** Monotone-Input-Reduction

**NSI** Non-Shared Inputs reduction

---

**PG** polarity based encoding

[PlaistedGreenbaum'86]

**TST** standard Tseitin encoding

---

**VE** Variable-Elimination as in DP / Quantor / SATeLite

**BCE** Blocked-Clause-Elimination

