



Practical SAT Solving

Lecture 5 – Conflict-Driven Clause Learning Markus Iser, Dominik Schreiber | May 19, 2025



Overview

Recap.

Lecture 4: Classic Heuristics and Modern SAT Solving 1:

- Decision Heuristics, Restart Strategies, Phase Saving
- Modern SAT Solving 1: Conflict Analysis / Clause Learning

Today's Topic: Modern SAT Solving 2

- Efficient Unit Propagation
- Clause Forgetting
- Modern Decision Heuristics



Conflict-driven Clause Learning (CDCL) Algorithm

Last Time

- Classic Decision Heuristics
- Restart Strategies
- Clause Learning
- Non-Chronological Backtracking

Today

- Efficient Unit Propagation
- Clause Forgetting
- Modern Decision Heuristics

Algorithm 1: CDCL(CNF Formula *F*, &Assignment $A \leftarrow \emptyset$)

- 1 if not PREPROCESSING then return UNSAT
- 2 while A is not complete do
- 3 UNIT PROPAGATION
- 4 **if** A falsifies a clause in F **then**
 - if decision level is 0 then return UNSAT

else

5

6

7

8

9

- (clause, level) \leftarrow CONFLICT-ANALYSIS
- add clause to *F* and backtrack to level
 - continue
- 10 if RESTART then backtrack to level 0
- 11 **if** CLEANUP **then** forget some learned clauses
- 12 BRANCHING

13 return SAT



Hot Paths in CDCL Solvers	S
---------------------------	---

heat	\varnothing per sec. ¹	
Clause Access		Unpredictable memory access: most expensive
Iterate Occurrences		Predictable memory access: array of pointers (hardware prefetching)
Propagation	$\sim 10^{6}$	Access occurrence-list of yet unpropagated literal
Decision	$\sim 10^3$	
Conflict	$\sim 10^3$	Learn a clause $ ightarrow$ more to check for propagation
Restart	$\sim 10^{-1}$	
Cleanup		Forget some learned clauses $ ightarrow$ less to check for propagation



Examp	Example: Unit Propagation with Full Occurrence Lists								
Trail		Occurrence Lists Formula							
level	value	reason	idx.	occurrences	addr.	claus	е		
1	а		а	* 1	*1	а	b	С	
			$\neg a$	*2 *3	*2	$\neg a$	b	$\neg C$	
			b	*1 *2	*3	$\neg a$	$\neg b$	С	
			$\neg b$	*3					
			С	*3 *1					
			¬ <i>C</i>	*2					



Example: Unit Propagation with Full Occurrence Lists

Irall		
level	value	reason
1	а	

Occurrence Lists				
occurre	ences			
*1				
*2	*3			
*1	*2			
*3				
*3	*1			
*2				
	nce Lists occurre *1 *2 *1 *3 *3 *3			

Formula

addr.	clause	9		
* 1	а	b	С	
*2	$\neg a$	b	$\neg C$	
*3	$\neg a$	$\neg b$	С	



Example: Unit Propagation with Full Occurrence Lists

Trail			Occurr	ence Lists	Formula	a		
level	value	reason	idx.	occurrences	addr.	claus	e	
1	а		а	*1	*1	а	b	С
2	С		$\neg a$	*2 *3	*2	$\neg a$	b	¬ <i>C</i>
			b	*1 *2	*3	$\neg a$	$\neg b$	С
			$\neg b$	*3				
			С	*3 *1				
			$\neg c$	*2				



Example: Unit Propagation with Full Occurrence Lists

Trail			Occurrence Lists		Formula	9		
level	value	reason	idx.	occurrences	addr.	clause	9	
1	а		а	* 1	*1	а	b	С
2	С		$\neg a$	*2 *3	*2	$\neg a$	b	¬ <i>C</i>
			b	*1 *2	*3	$\neg a$	$\neg b$	С
			$\neg b$	*3				
			С	*3 *1				
			¬ <i>C</i>	*2				



Example: Unit Propagation with I	Full Occurrence Lists
----------------------------------	-----------------------

Trail			Occurrence Lists		Formula	a		
level	value	reason	idx.	occurrences	addr.	clause	•	
1	а		а	*1	*1	а	b	С
2	С		$\neg a$	*2 *3	*2	$\neg a$	b	¬ <i>C</i>
2	b	*2	b	*1 *2	*3	$\neg a$	$\neg b$	С
			$\neg b$	*3	-			
			С	*3 *1	-			
			$\neg c$	*2				



Motivation: Hot Path					
heat	Ø per sec. ²	Idea: Reduced occurrence tracking by only keeping the following invariant:			
Clause Access		Each yet unsatisfied clause is watched by, i.e., in the occurrence list of, two of its unassigned literals. Beasoning: less literals watched \rightarrow shorter occurrence lists \rightarrow less clause			
Iterate Occurrences		accesses \rightarrow fast unit propagation			
Propagation	\sim 10 ⁶				



Motivation: Hot Path						
heat	Ø per sec.²	Idea: Reduced occurrence tracking by only keeping the following invariant:				
Clause Access		Each yet unsatisfied clause is watched by, i.e., in the occurrence list of, two of its unassigned literals.				
Iterate Occurrences		Reasoning: less literals watched \rightarrow shorter occurrence lists \rightarrow less clause accesses \rightarrow fast unit propagation Why do two watched literals per clause suffice?				
Propagation	$\sim 10^{6}$					



Motivation: Hot Pa	Notivation: Hot Path								
heat	Ø per sec.²	Idea: Reduced occurrence tracking by only keeping the following invariant:							
Clause Access		Each yet unsatisfied clause is watched by, i.e., in the occurrence list of, two of its unassigned literals.							
		Reasoning: less literals watched \rightarrow shorter occurrence lists \rightarrow less clause accesses \rightarrow fast unit propagation							
Iterate Occurrences		Why do two watched literals per clause suffice?							
Propagation	\sim 10 ⁶	vvny does one watched illeral per clause not suffice?							



Motivation: Hot Pa	Notivation: Hot Path							
heat	Ø per sec.²	Idea: Reduced occurrence tracking by only keeping the following invariant:						
Clause Access		Each yet unsatisfied clause is watched by, i.e., in the occurrence list of, two of its unassigned literals.						
		Reasoning: less literals watched \rightarrow shorter occurrence lists \rightarrow less clause accesses \rightarrow fast unit propagation						
Iterate Occurrences		Why do two watched literals per clause suffice?						
Propagation	\sim 10 ⁶	 Why does one watched literal per clause not suffice? How do we keep that invariant? (Branching?, Backtracking?) 						



Example: Unit Propagation with Two Watched Literals										
Trail	ail Two Watched Literals			Formula	a					
level	value	reason	idx.	occu	rrences		addr.	claus	е	
			а	* 1			* 1	а	b	С
			$\neg a$	*2	*3		*2	$\neg a$	b	$\neg c$
			b	*1	*2		*3	$\neg a$	$\neg b$	С
			$\neg b$	*3						
			С							
			$\neg C$							



Trail		Two Wa	atched Literals	Formula	Formula			
level	value	reason	idx.	occurrences	addr.	claus	е	
1	а		а	*1	* 1	а	b	С
			$\neg a$	*2 *3	*2	$\neg a$	b	$\neg C$
			b	*1 *2	*3	$\neg a$	$\neg b$	С
			$\neg b$	*3				
			С					
			$\neg C$					



Example: Unit Propagation with Two Watched Literals

value a	reason
а	

Two Watched Literals						
idx.	occurrences					
а	*1					
$\neg a$	*3					
b	*1 *2					
$\neg b$	*3					
С						
$\neg C$	*2					

Formula addr. clause *1 a b c *2 ¬c b ¬a *3 ¬a ¬b c



Example: Unit Propagation with Two Watched Literals

Trail		
level	value	reason
1	а	

Two Watched Literals							
idx.	occurrences						
а	*1						
$\neg a$							
b	*1 *2						
$\neg b$	*3						
С	*3						
$\neg C$	*2						

Formula

addr.	clause		
* 1	а	b	С
*2	¬ <i>C</i>	b	$\neg a$
*3	С	$\neg b$	<i>¬</i> a



Example: Unit Propagation with Two Watched Literals									
Trail	rail Two Watched Literals				Formula	a			
level	value	reason	idx.	occurrences	addr.	claus	е		
1	а		а	* 1	*1	а	b	С	
2	С		$\neg a$		*2	$\neg C$	b	$\neg a$	
			b	*1 *2	*3	С	$\neg b$	$\neg a$	
			$\neg b$	*3					
			С	*3					
			$\neg c$	*2					



Example: Unit Propagation with Two Watched Literals									
Trail	rail Two Watched Literals				Formula	a			
level	value	reason	idx.	occurrences	addr.	claus	е		
1	а		а	* 1	*1	а	b	С	
2	С		$\neg a$		*2	$\neg C$	b	$\neg a$	
2	b	*2	b	*1 *2	*3	С	$\neg b$	$\neg a$	
			$\neg b$	*3					
			С	*3					
			$\neg c$	*2					



Two Watched Literals: Optimizations					
heat	Ø per sec. ³	Invariant: Each yet unsatisfied clause is watched by two of its unassigned literals.			
Clause Access		→ Reduced Load in Occurrence Tracking			
Iterate Occurrences		\rightarrow Alternative: Store watched literals the first two in clause Note: What happens if clauses are kept in shared memory for parallel solving?			
		Optimization 2: Also keep a literal of each clause directly in occurrence list \rightarrow Skip clause access if that literal is satisfied			
Propagation	\sim 10 ⁶				





- Hottest path in CDCL solvers
- Two watched literals per clause suffice (for unit propagation and conflict detection)

Further optimizations:

- Invariant: the two first literals in each clause are the watched ones
- **Blocking literals:** keep a literal of each clause directly in occurrence list

Next Up

Clause Forgetting



Motivation

Clause learning is most important pruning strategy in CDCL solvers.⁴

Problem:

- Slows down unit propagation
- Risk of running out of memory

Solution:

- Periodically forget some learned clauses
- Keep only "the best" learned clauses

⁴"Empirical Study of the Anatomy of Modern SAT Solvers", Katebi et al., 2013



Motivation

Clause learning is most important pruning strategy in CDCL solvers.⁴

Problem:

- Slows down unit propagation
- Risk of running out of memory

Solution:

- Periodically forget some learned clauses
- Keep only "the best" learned clauses
- How to figure out which learned clauses are "the best"?

⁴"Empirical Study of the Anatomy of Modern SAT Solvers", Katebi et al., 2013



Periodic Clause Forgetting: Heuristics

Clause Size

Keep short clauses

Least Recently Used (LRU)

Keep clauses which where reasons in recent conflicts: clause activity (EMA)



Periodic Clause Forgetting: Heuristics

Clause Size

Keep short clauses

Least Recently Used (LRU)

Keep clauses which where reasons in recent conflicts: clause activity (EMA)

Literal Block Distance (LBD)

Keep clauses with a low number of decision levels⁵

⁵Predicting Learnt Clauses Quality in Modern SAT Solvers, Audemard & Simon (IJCAI 2009)



Periodic Clause Forgetting: Heuristics

Clause Size

Keep short clauses

Least Recently Used (LRU)

Keep clauses which where reasons in recent conflicts: clause activity (EMA)

Literal Block Distance (LBD)

Keep clauses with a low number of decision levels⁵

Why is low LBD good?

⁵Predicting Learnt Clauses Quality in Modern SAT Solvers, Audemard & Simon (IJCAI 2009)



Forgetting Heuristic: Literal Block Distance (LBD)

"Impact of Community Structure on SAT Solver Performance", Newsham et al., SAT 2014

Take home: LBD correlates with number of touched communities



Image Source: "Community Structure in Industrial SAT Instances", Ansotegui et al., AIJ 2019



Clause Forgetting: Modern Hybrid Approach

Manage clauses differently in three tiers

Tier	Strategy	Description	
core	LBD	Permanently store clauses of LBD $\leq k$ (core-cut value, 3 in practice)	
mid-tier	LRU	Clauses stay here if used in recent conflicts	
local	LRU	Keep fixed number of clauses (say 5000) of highest activity	

History

- core and local tier introduced in SWDiA5BY (Chanseok Oh, 2014)
- mid-tier introduced in CoMinisatPS (Chanseok Oh, 2015)
- "Between SAT and UNSAT: The Fundamental Difference in CDCL SAT" (Chanseok Oh, 2015)
- Note: The award-winning SAT solver MapleCOMSPS (2016) is a CoMinisatPS fork



initial layout, recently active variables after 1000 conflicts





initial layout, recently active variables after 1690 conflicts





initial layout, recently active variables after 3090 conflicts





initial layout, recently active variables after 5000 conflicts







14/24

core after 52500 conflicts



15/24









after 300000 conflicts









So far

- Efficient Unit Propagation
- Clause Forgetting Heuristics

Next Up

Modern Decision Heuristics



VSIDS Heuristic

Implemented in most CDCL solvers. First presented in SAT solver Chaff.⁶

Always select variable with highest score for branching. Scores are updated after each conflict.

Initialize variable score (with zero or use some static heuristic)

New conflict clause c: score is incremented for all variables in c

Periodically, divide all scores by a constant

⁶Chaff: Engineering an efficient SAT solver (Moskewicz et al., 2001)



Example: Score Update after Conflict

Formula:		Scores before:	Scores after:
$\{x_1, x_4\}, \{x_1, \overline{x_3}, \overline{x_8}\}, \{x_1, \overline{x_3}, \overline{x_8}\}, \{x_1, $	$\{x_1, x_8, x_{12}\}, \{x_2, x_{11}\}, $	4 : <i>x</i> ₈	4 : <i>x</i> ₈ , <i>x</i> ₇
$\{\overline{x_7}, \overline{x_3}, x_9\}, \{\overline{x_7}, x_8, \overline{x_9}\}$	$\overline{\mathbf{y}}$, { \mathbf{x}_7 , \mathbf{x}_8 , $\overline{\mathbf{x}_{10}}$ }	$3: x_1, x_7$	3 : <i>x</i> ₁
$\{\mathbf{X}_7, \mathbf{X}_{10}, \overline{\mathbf{X}_{12}}\}$	(new learned clause)	2 : <i>x</i> ₃	2 : <i>x</i> ₃ , <i>x</i> ₁₀ , <i>x</i> ₁₂
		1 : $x_2, x_4, x_9, x_{10}, x_{11}, x_{12}$	$1: x_2, x_4, x_9, x_{11}$



Example: Score Update after Conflict

Formula:	Scores before:	Scores after:
$\{x_1, x_4\}, \{x_1, \overline{x_3}, \overline{x_8}\}, \{x_1, x_8, x_{12}\}, \{x_2, x_{11}\}, \{x_3, x_{12}\}, \{x_4, x_{12}\}, \{x_2, x_{11}\}, \{x_3, x_{12}\}, \{x_4, x_{12}\}, \{x_4, x_{12}\}, \{x_5, x_{$	4 : <i>x</i> ₈	4 : <i>x</i> ₈ , <i>x</i> ₇
$\{\overline{\textbf{X}_7}, \overline{\textbf{X}_3}, \textbf{X}_9\}, \{\overline{\textbf{X}_7}, \textbf{X}_8, \overline{\textbf{X}_9}\}, \{\textbf{X}_7, \textbf{X}_8, \overline{\textbf{X}_{10}}\}$	$3: x_1, x_7$	3 : <i>x</i> ₁
$\{x_7, x_{10}, \overline{x_{12}}\}$ (new learned clause)	2 : <i>x</i> ₃	2 : <i>x</i> ₃ , <i>x</i> ₁₀ , <i>x</i> ₁₂
	$1: x_2, x_4, x_9, x_{10}, x_{11}, x_{12}$	1 : <i>x</i> ₂ , <i>x</i> ₄ , <i>x</i> ₉ , <i>x</i> ₁₁

- VSIDS leads to more "focused" search
- prefers variables that occurred in recent conflicts
- tends to find smaller unsatisfiable subsets



Common implementation: Binary Heap

Heap Operation	Complexity	Callee
insert_with_priority	$\mathcal{O}(\log n)$	Backtracking
pull_highest_priority_element	$\mathcal{O}(\log n)$	Branching
increase_key / bump_variable	$\mathcal{O}(\log n)$	Conflict Analysis
decay	$\mathcal{O}(n)$	[Periodic] ⁷

⁷Periodically divide scores to give priority to recently learned clauses



Historic Implementations

Chaff (2001)

- decay: half scores every 256 conflicts
- sort priority queue after each decay only

Minisat (2003): Exponential VSIDS (EVSIDS)

Idea: Exponential decay of scores s(v) with damping factor 0 < f < 1

 $s'(v) := egin{cases} f \cdot s(v) + (1-f) & ext{if } v ext{ is to be bumped} \ f \cdot s(v) & ext{otherwise} \end{cases}$

Berkmin (2002)

divide scores by 4

bump all literals in implication graph

Theory: Exponential Moving Average (EMA)

Implementation: Score increment by g^i , with *i* denoting the conflict-index and $g = \frac{1}{f}$ (no decay)

Reason-side Bumping: Also bump variables in the reason clauses of the conflict

Evaluating CDCL Variable Scoring Schemes (Biere & Fröhlich, 2015)



solved SAT competition 2014 application track instances (ordered by time)



Alternative and Hybrid Approaches

Alternatives

Siege (2004): Variable Move To Front (VMTF)

HaifaSAT (2008): Clause Move To Front (CMTF)

Recent Hybrid Approaches

• Warmup Phase:

- MapleCOMSPS (2016): use Learning Rate-based Branching (LRB) in *initial* period, then switch to VSIDS
- Maple_LCM_Dist (2017): use Distance Heuristic (Dist.) in *initial* period, then switch to VSIDS

Reinforcement Learning: Kissat_MAB (2021)

- Two-armed Bandid switches between VSIDS and Conflict History-Based (CHB) Heuristic
- Reward function favors variables that contribute to learning "good" clauses



Modern Branching: Local Search Intergration

Better Decision Heuristics in CDCL through Local Search and Target Phases (Cai et al., 2022)

Branching:

- Target Phases: cache and use the phases which led to the previously largest assignment
- Integrate Sprints of Local Search: use unit-propagation to complete the assignment (ignoring all conflicts)
- Rephasing: save the best assignment found during local search for phase selection (diversification)

Ordering:

Import Statistics: use frequency of appearing in unsatisfied clauses to modify the variables VSIDS score



Recap

Recap

- Efficient Unit Propagation
- Clause Forgetting
- Modern Decision Heuristics

Next Time

Preprocessing

Algorithm 2: CDCL(CNF Formula *F*, &Assignment $A \leftarrow \emptyset$)

- 1 if not PREPROCESSING then return UNSAT
- 2 while A is not complete do
- 3 UNIT PROPAGATION
- 4 if A falsifies a clause in F then
 - if decision level is 0 then return UNSAT

else

5

6

7

8

9

- (clause, level) \leftarrow CONFLICT-ANALYSIS
- add clause to *F* and backtrack to level
- continue
- 10 **if** RESTART **then** backtrack to level 0
- 11 **if** CLEANUP **then** forget some learned clauses
- 12 BRANCHING

13 return SAT

