

**MaxSAT**  
**Current Solution Techniques**

Fahiem Bacchus  
University of Toronto

# MaxSat

- ▶ MaxSat is a formalism for expressing Boolean **optimization** problems expressed in CNF.
  - ▶ Hard clauses: must be satisfied
  - ▶ Soft clauses each with a weight, falsifying these incurs a cost equal to their weight.
- ▶ MaxSat: find a truth assignment that **satisfies all of the hard clauses** while falsifying **a minimum weight** of soft clauses. (Equivalently satisfying a maximum weight of soft clauses).
- ▶ MaxSat solvers are effective in a growing range of applications.

# MaxSat Applications


probabilistic inference	[Park, 2002]
design debugging	[Chen, Safarpour, Veneris, and Marques-Silva, 2009] [Chen, Safarpour, Marques-Silva, and Veneris, 2010]
maximum quartet consistency	[Morgado and Marques-Silva, 2010]
software package management	[Argelich, Berre, Lynce, Marques-Silva, and Rapicault, 2010] [Ignatiev, Janota, and Marques-Silva, 2014]
Max-Clique	[Li and Quan, 2010; Fang, Li, Qiao, Feng, and Xu, 2014; Li, Jiang, and Xu, 2015]
fault localization	[Zhu, Weissenbacher, and Malik, 2011; Jose and Majumdar, 2011]
restoring CSP consistency	[Lynce and Marques-Silva, 2011]
reasoning over bionetworks	[Guerra and Lynce, 2012]
MCS enumeration	[Morgado, Liffiton, and Marques-Silva, 2012]
heuristics for cost-optimal planning	[Zhang and Bacchus, 2012]
optimal covering arrays	[Ansótegui, Izquierdo, Manyà, and Torres-Jiménez, 2013b]
correlation clustering	[Berg and Järvisalo, 2013; Berg and Järvisalo, 2016]
treewidth computation	[Berg and Järvisalo, 2014]
Bayesian network structure learning	[Berg, Järvisalo, and Malone, 2014]
causal discovery	[Hyttinen, Eberhardt, and Järvisalo, 2014]
visualization	[Bunte, Järvisalo, Berg, Myllymäki, Peltonen, and Kaski, 2014]
model-based diagnosis	[Marques-Silva, Janota, Ignatiev, and Morgado, 2015]
cutting planes for IPs	[Saikko, Malone, and Järvisalo, 2015]
argumentation dynamics	[Wallner, Niskanen, and Järvisalo, 2016]




# Annual MaxSat Evaluations

- ▶ Each year we receive new maxsat benchmarks in which maxsat is being used to solve new problems.
  - ▶ Not surprising as many problems involve optimization.
- ▶ The improvement in performance in MaxSat solvers has been impressive.

# Improvements in MaxSat Solving WEIGHTED (2008-2016)



# Improvements in MaxSat Solving UNWEIGHTED (2008-2016)



# Problem Sizes

- ▶ Largest problems solved in 2017 MaxSat Evaluation, >6,000,000 variables and > 13,000,000 clauses (solved by MaxHS in < 800 sec.)
- ▶ MaxSat is considerably harder than SAT, SAT solvers can solve bigger problems

# Outline

- ▶ Formalizing the MaxSat problem and some necessary notation.
- ▶ Solvers that use SAT to solve a sequence of relaxations.
- ▶ Implicit hitting set solvers that exploit both SAT and Integer Programming technology.



## Necessary Formal Notions

# MaxSat

- ▶ A MaxSat instance  $\mathbf{F}$  is a CNF where the clauses are partitioned into HARD clauses and SOFT clauses:  
 $\mathbf{F} = \mathbf{H} \cup \mathbf{S}$
- ▶ Associated with each soft clause  $c \in \mathbf{S}$  is a integer weight  $\mathbf{wt}(c) > 0$ .
  - ▶ Some solvers can handle floating point weights.
- ▶ For any set of soft clauses  $A \subseteq \mathbf{S}$ ,  $\mathbf{wt}(A)$  is the sum of the weights of the clauses in  $A$ .

$$\mathbf{wt}(A) = \sum_{c \in A} \mathbf{wt}(c)$$

# MaxSat

- ▶ A truth assignment  $\pi$  that satisfies  $\mathbf{H}$  is called a **feasible solution**:  $\pi \models \mathbf{H}$
- ▶ **Cost** of a feasible solution  $\pi$  is the weight of the set of soft clauses it falsifies.

$$cost(\pi) = \sum_{\pi \not\models c} wt(c)$$

- ▶ A feasible solution is an **optimal solution** if it has minimum cost amongst all feasible solutions

$$\forall \sigma. \sigma \models \mathbf{H} \rightarrow cost(\sigma) \geq cost(\pi)$$

- ▶ Solving a MaxSat instance  $F$  means finding an optimal solution.

# Cores

- ▶ A set of soft clauses  $\kappa \subseteq \mathbf{S}$  is a **core** of  $\mathbf{F}$  if

$\kappa \cup \text{hard}(\mathbf{F})$  is **UNSAT**

- ▶ Note that in SAT a core is a subset of clauses that are UNSAT. In MaxSat we always have to satisfy the hard clauses, so more useful to define cores relative to the hard clauses.

# Cores via Assumptions

- ▶ The current best performing algorithms for MaxSat need to extract cores.
- ▶ Currently most accessible way to do this is to use SAT with assumptions...built into most SAT solvers.
- ▶ Assumptions must be a set of literals.
- ▶ **SAT\_ASSUME(H, Asmp)**
  - ▶ Sat solve the CNF H under the assumption that every literal in Asmp is true.
  - ▶ Return SAT and  $\pi$  if  $\pi \models H$  and makes every literal in Asmp true
  - ▶ Return UNSAT and a conflict clause  $(\neg l_1, \neg l_2, \dots, \neg l_k)$  implied by H where each  $l_i \in \text{Asmp}$ .
    - ▶ At least one of the subset of assumptions  $\{l_1, l_2, \dots, l_k\}$  must be falsified in every model of H.

# Cores via Assumptions

- ▶ Since assumptions are limited to literals, to extract MaxSat cores we must have literals that imply the satisfaction of soft clauses
- ▶ Unit soft clauses, we can use its literal as an assumption
  - ▶ (z) we can use the literal z as an assumption.
- ▶ Non unit soft clauses **C**, we need to create a new literal b, add to the hard clauses the condition  $b \rightarrow C$ , and then use b as an assumption.
  - ▶  $(x, y, z) \rightarrow (\neg b, x, y, z)$
  - ▶  $H = H \cup \{(\neg b, x, y, z)\}$
  - ▶ Use b as an assumption

# Conversion to Unit Softs

- ▶ In fact this “getting the formula ready for using assumptions” transformation is the same as converting it to a new MaxSat instance which has **only unit soft clauses**.
- ▶  $F = H \cup S \rightarrow F^b = H^b \cup S^b$
- ▶ Let  $S^{>1} = \{c_i \mid c_i \in S \text{ and is non-unit}\}$
- ▶  $H^b = H \cup \{(c_i \vee \neg b_i) \mid c_i \in S^{>1}\}$  with new variable  $b_i$ 
  - ▶ Make non-unit softs into assumable hard clause
- ▶  $S = (S \setminus S^{>1}) \cup \{(b_i) \mid c_i \in S^{>1}\}$ 
  - ▶ Add assumption literals as softs.

# Conversion to Unit Softs

- ▶ This new "only unit softs" instance is equivalent:
- ▶ Every feasible model of  $\mathbf{F}$ ,  $\pi$ , can be extended to a feasible model of  $F^b$  with the same cost.
  - ▶ For each new variable  $b^i$  in  $F^b$  set  $b^i$  to TRUE iff  $\pi \models c^i$
- ▶ Every feasible model of  $F^b$ ,  $\pi^b$  restricted to the variables of  $F$  is a feasible model  $\pi$  of  $F$ .  $\text{cost}(\pi) \leq \text{cost}(\pi^b)$
- ▶ **Hence, an optimal model of  $F^b$  restricted to the variables of  $F$  is an optimal model of  $F$ .**



# Cores via Assumptions with unit softs.

- ▶ With  $F^b$  we can easily use assumptions to extract cores.
- ▶ **SAT\_ASSUME( $H^b$ ,  $Asmp$ )**  
where  $Asmp$  are the literals of a set of unit softs.
- ▶ If UNSAT, we get a conflict clause  $(\neg l_1, \neg l_2, \dots, \neg l_k)$  where each  $\neg l_i$  falsifies the unit soft  $(l_i)$
- ▶ The negation of literals in this conflict  $\{l_1, l_2, \dots, l_k\}$  are a **core**.

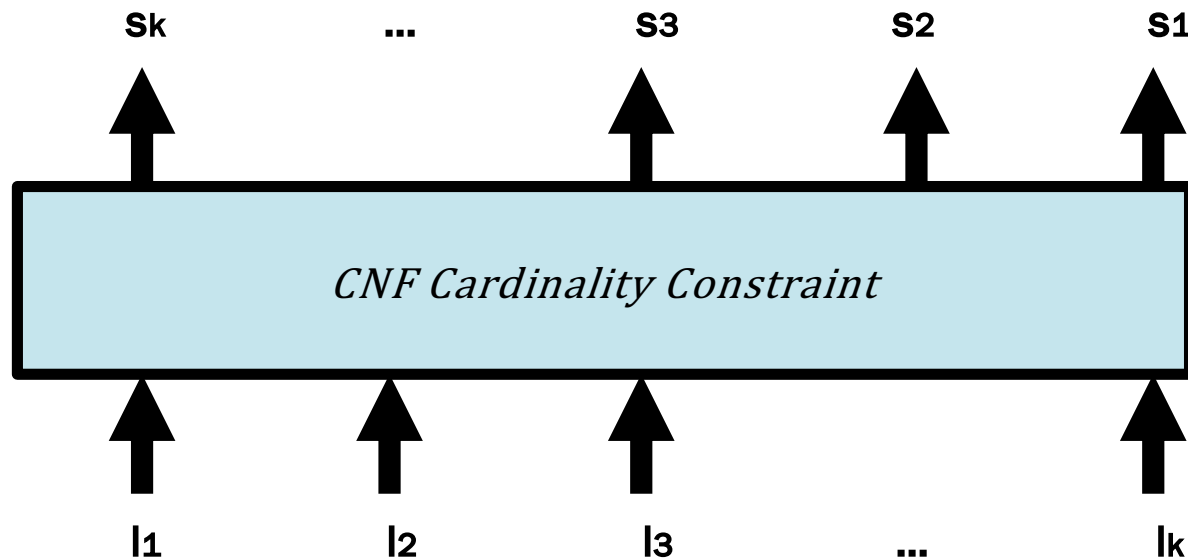


# Unhiding the use of Assumptions

- ▶ Note that most MaxSat algorithms in the literature are specified by abstracting away from the assumption mechanisms.
  - ▶ An interface with the SAT solvers is assumed that returns cores in a black box manner.
- ▶ But I think this can make some of the details these algorithms a bit more difficult to understand.
- ▶ So in this talk explain these algorithms in terms of assumptions to the SAT solver.


# Summation Circuits

- ▶ MaxSat solvers using sequences of relaxations exploit cardinality constraints.
- ▶ These use CNF encodings of summation circuits that output a unary representation of the sum of their inputs.



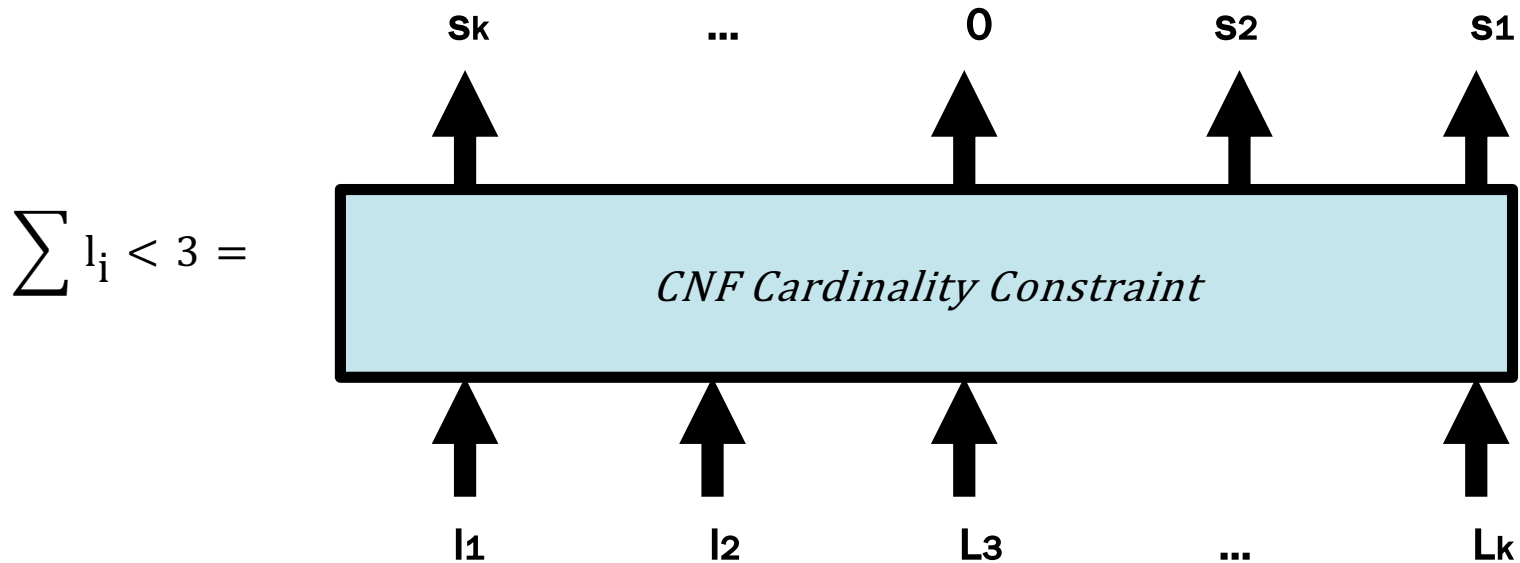
# Summation Circuits

- ▶ The output is a string of 0's followed by 1's with (as many 1's as there are true inputs).
- ▶ The clauses of the summation circuit ensure that
  - ▶  $s_i \leftrightarrow |_1 + |_2 \dots + |_k \geq i$
  - ▶  $\neg s_i \leftrightarrow |_1 + |_2 \dots + |_k < i$



# Cardinality Constraints

- ▶ So we can impose the constraint  $\sum l_i < k$  by adding the clauses of the summation circuit along with the assumption  $\neg s_k$



**Solving MaxSat by SAT solving a sequence of relaxation**

# Sat Solving a sequence of relaxations

- ▶ A sequence of SAT instances are created (usually incrementally) where each instance is more relaxed so that it allows a larger weight of soft clauses to be falsified.
- ▶ The amount of additional weight that can be falsified is restricted so that the first time the formula becomes SAT the resulting satisfying models are optimal solutions to the original MaxSat formula.
  - ▶ i.e., there is no lesser relaxation that is satisfiable.
- ▶ We can also go the other direction starting with most relaxed and working down until we reach a relaxation that is UNSAT.
  - ▶ This works ok, but not as well (so far)

# CARDINALITY LAYER

- ▶ First we start off by ensuring that the input MaxSat  $F$  has only **unit soft clauses**. (Using the  $F^b$  conversion discussed before).
- ▶ Then the needed relaxations are achieved by adding a **CARDINALITY LAYER** to the **hard clauses of  $F$** .
- ▶ The **inputs** of the **CARDINALITY LAYER** are the **negations** of the unit soft clause literals--limiting how many of these that can be true limits how the number of soft clauses that can be falsified.
- ▶ The **outputs** of the **CARDINALITY LAYER** can be set by assumptions to restrict how many and which groups of softs that can be falsified.



# UNWEIGHTED INSTANCES

- ▶ First we assume that all soft clauses have the same weight—so solving MaxSat is the same as minimizing the number of falsified softs.
- ▶ Later we will show the technique for lifting to the weighted case (i.e., differing weights).
  - ▶ This technique is simple but it seems to lead to some inefficiencies.

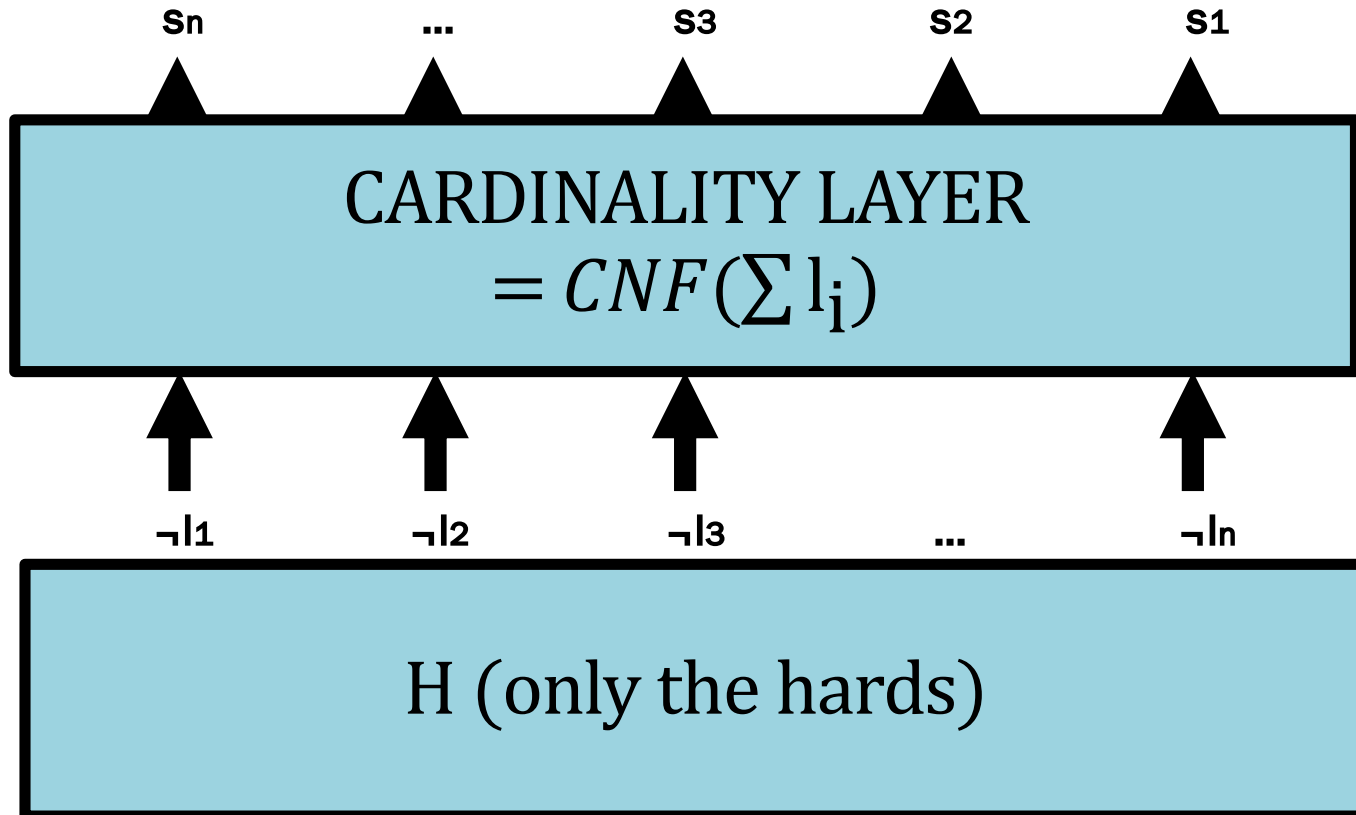
# Example

- ▶ Standard reduction of MaxSat to a sequence of SAT problems
  - ▶ Is the formula satisfiable falsifying 0 softs?
  - ▶ Is the formula satisfiable falsifying 1 soft?
  - ▶ ...
  - ▶ Is the formula satisfiable falsifying  $k$  softs?
- ▶ Stop as soon as the answer is yes...the truth assignment (excluding the new variables in the cardinality layer) is an optimal solution.

# Cardinality Layer Simple Example

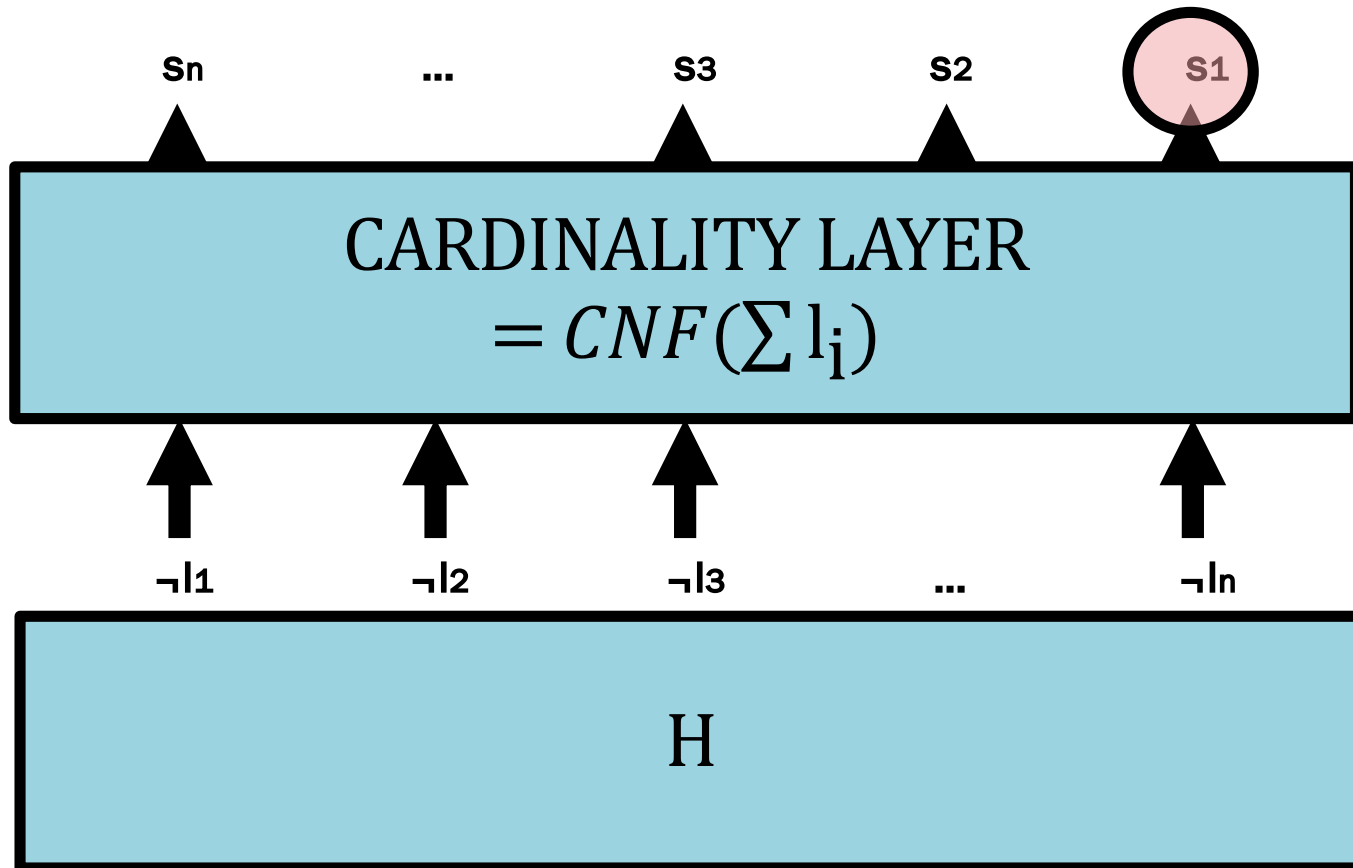
$F = H \cup S$

$S = \{ (l_1), (l_2), \dots, (l_n) \}$




# Cardinality Layer Simple Example

On the  $i$ -th iteration solve  $\text{SAT\_ASSUME}(H \cup \text{Card}, \neg s_i)$




# Cardinality Layer Simple Example

On the  $i$ -th iteration solve  $\text{SAT\_ASSUME}(H \cup \text{Card}, \neg s_i)$



# Cardinality Layer Simple Example

On the  $i$ -th iteration solve  $\text{SAT\_ASSUME}(H \cup \text{Card}, \neg s_i)$




# Cardinality Layer

- ▶ In these algorithms the unit soft clauses are not part of the SAT solver's CNF.
  - ▶ Their literals serve as inputs to the cardinality layer.
- ▶ When we set various various literals in the cardinality layer as assumptions these assumptions restrict the allowed T/F settings of unit soft clauses.

# Cardinality Layer Simple Example

Feasible settings of  $l_1$  to  $l_n$  falsify at most 2 soft clauses

$\binom{n}{2} + \binom{n}{1} + \binom{n}{0}$  feasible settings





# Cardinality Layer

- ▶ Given some assumptions  $A$ , let  $CARD(A)$  be all of the set of T/F settings of the soft clauses allowed by the cardinality layer under  $A$ .
- ▶ Since the softs are units,  $CARD(A)$  corresponds to a set of truth assignments to  $\{l_1, l_2, \dots, l_n\}$  (the literals in the SOFTS).
- ▶ To refute  $A$ , the SAT solver must refute every  $\rho \in CARD(A)$ . That is, the SAT solver must verify that  $HARDS \cup \rho$  is UNSAT for every  $\rho$ .
  - ▶ If  $A$  is refuted then the softs set to true by every  $\rho$  form a core of the input MaxSat Formula  $F$ . (Fahiem Bacchus, Nina Narodytska, Cores in Core Based MaxSat Algorithms: An Analysis. SAT 2014).

# Cardinality Layer

- ▶ Clause learning can allow the SAT solver to refute entire subsets of  $\text{CARD}(A)$  rather than having to refute each  $\rho \in \text{CARD}(A)$  individually.
- ▶ The learnt clauses can contain the new variables added to construct the Cardinality Layer. This can support learning more powerful clauses that speed up the refutation.
- ▶ Different proposed algorithms construct differently structured Cardinality Layers.
- ▶ **Understanding more precisely how the structure of the Cardinality layer affects the efficiency of refuting  $\text{CARD}(A)$  is an important open question.**

# MSU3 using Incremental Cardinality Constraints

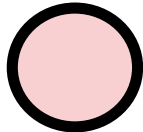
- ▶ Used in OpenWBO.
- ▶ R. Martins, S. Joshi, V. Manquinho & I. Lynce  
“Incremental Cardinality Constraints for MaxSAT”  
CP-2014
- ▶ (Original MSU3) J. Marques-Silva & J. Planes. “On Using Unsatisfiability for Maximum Satisfiability” CoRR, 2007

# MSU3 Incremental


- Consider  $F = H \cup S$  where
- $S = \{(x_1), (x_2), \dots, (x_8)\}$
- $H = \sum_{i=1}^8 x_i \leq 4$
- We want to satisfy all 8 literals  $x_i$  but at most 4 of these literals can be satisfied—at least 4 must be falsified.
- Every set of 5 or more soft clause is a core
- What will MSU3 do on this input formula?

# MUS3 Incremental

Start with Cardinality Constraint Layer with inputs only


 Literals assumed FALSE

First SAT solve is  $\text{SAT\_ASSUME}(H, \{x_1, x_2, \dots, x_8\})$ ,  
i.e., try to falsify zero softs.




# MUS3 Incremental

- UNSAT; say we get the core  $(\neg X_1, \neg X_2, \neg X_3, \neg X_4, \neg X_5)$
- Add one to overall cost
- Build summation network over softs in core




# MUS3 Incremental

- We know that the sum must be at least one—set  $s_1 = 1$ .
- Now **SAT\_ASSUME(H U Card,  $\{\neg s_2, x_6, x_7, x_8\}$ )**,
- Allow one soft to be falsified among  $x_1$ — $x_5$



# MUS3 Incremental

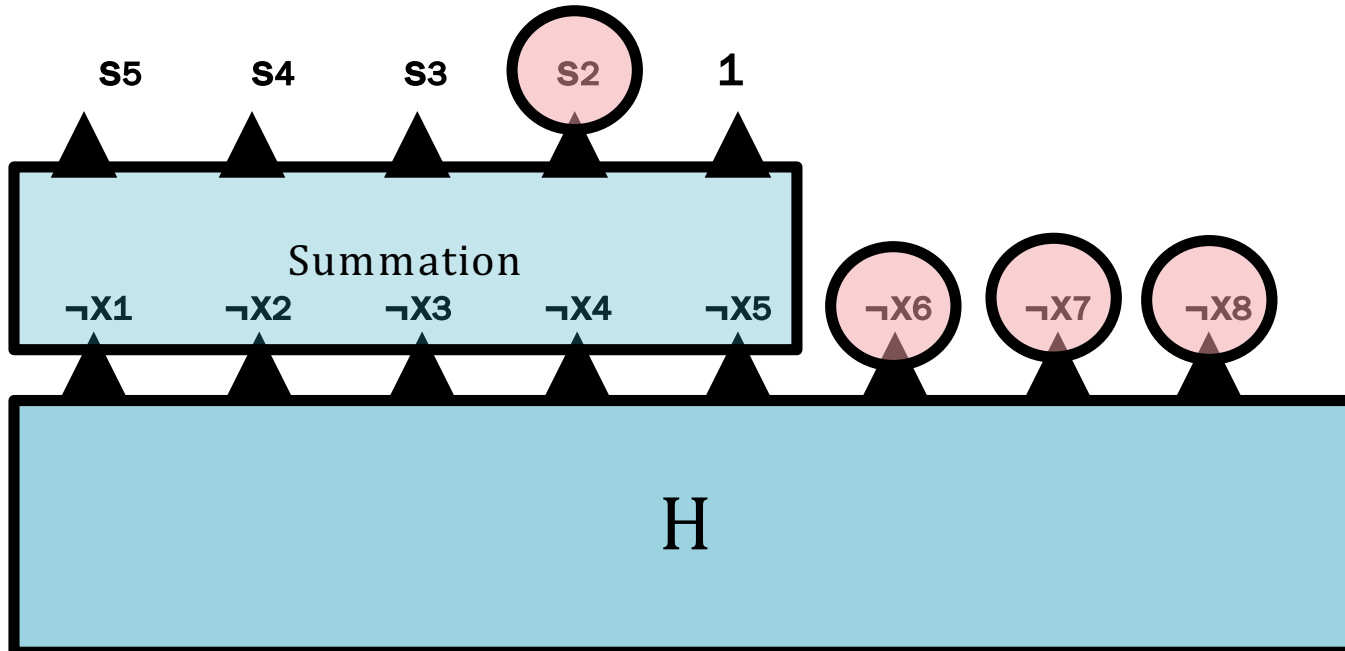
- Get another core, add 1 to overall cost.
- Get another conflict,  $(s_2, \neg x_6)$ 
  - ▶ Either we must falsify two of  $x_1 \dots x_5$  or one of  $x_1 \dots x_5$  and  $x_6$





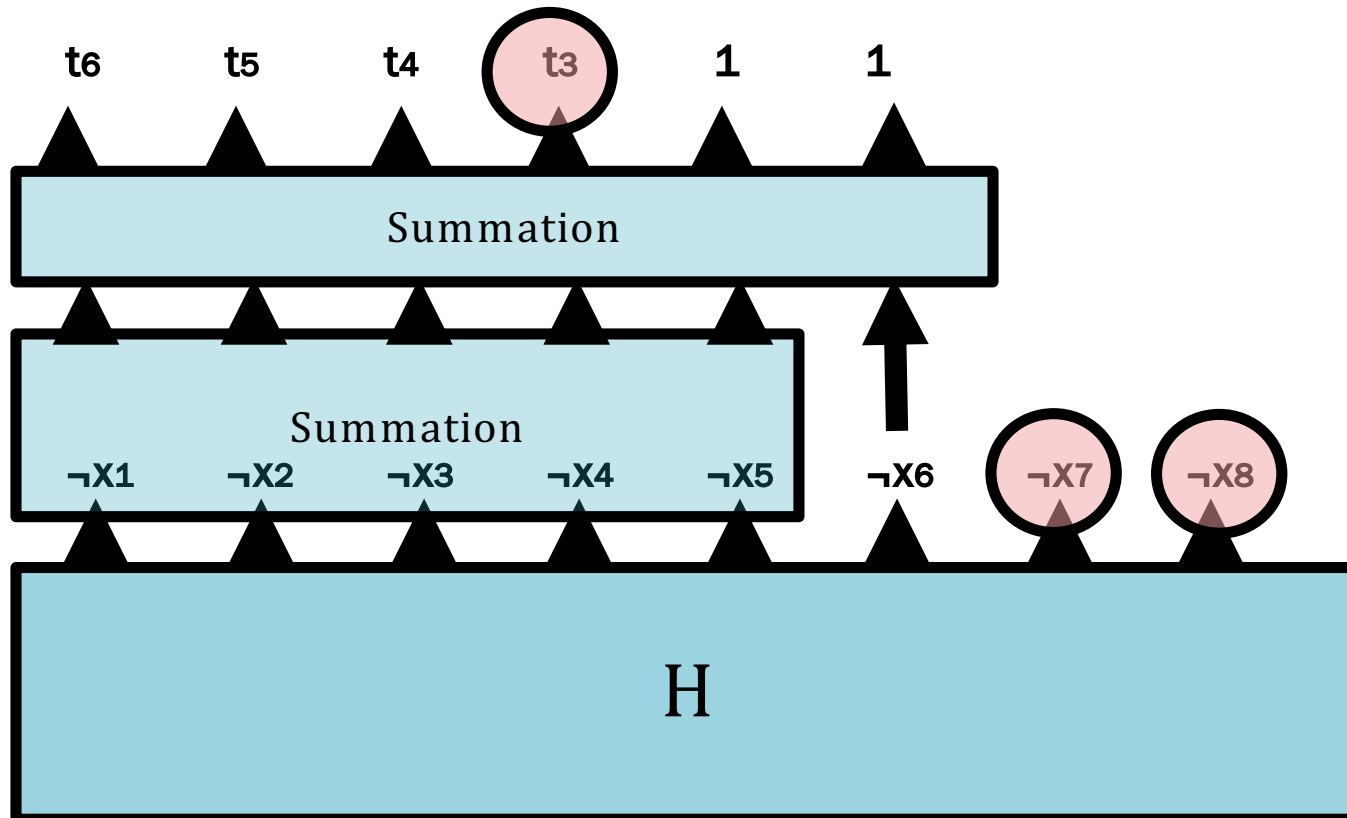
# MUS3 Incremental

- Now the incremental part. MUS3-Incremental reuses the previous summation constraint and its variables **subsuming** them into a new summation constraint that accounts for the new core ( $s_2, \neg X_6$ )



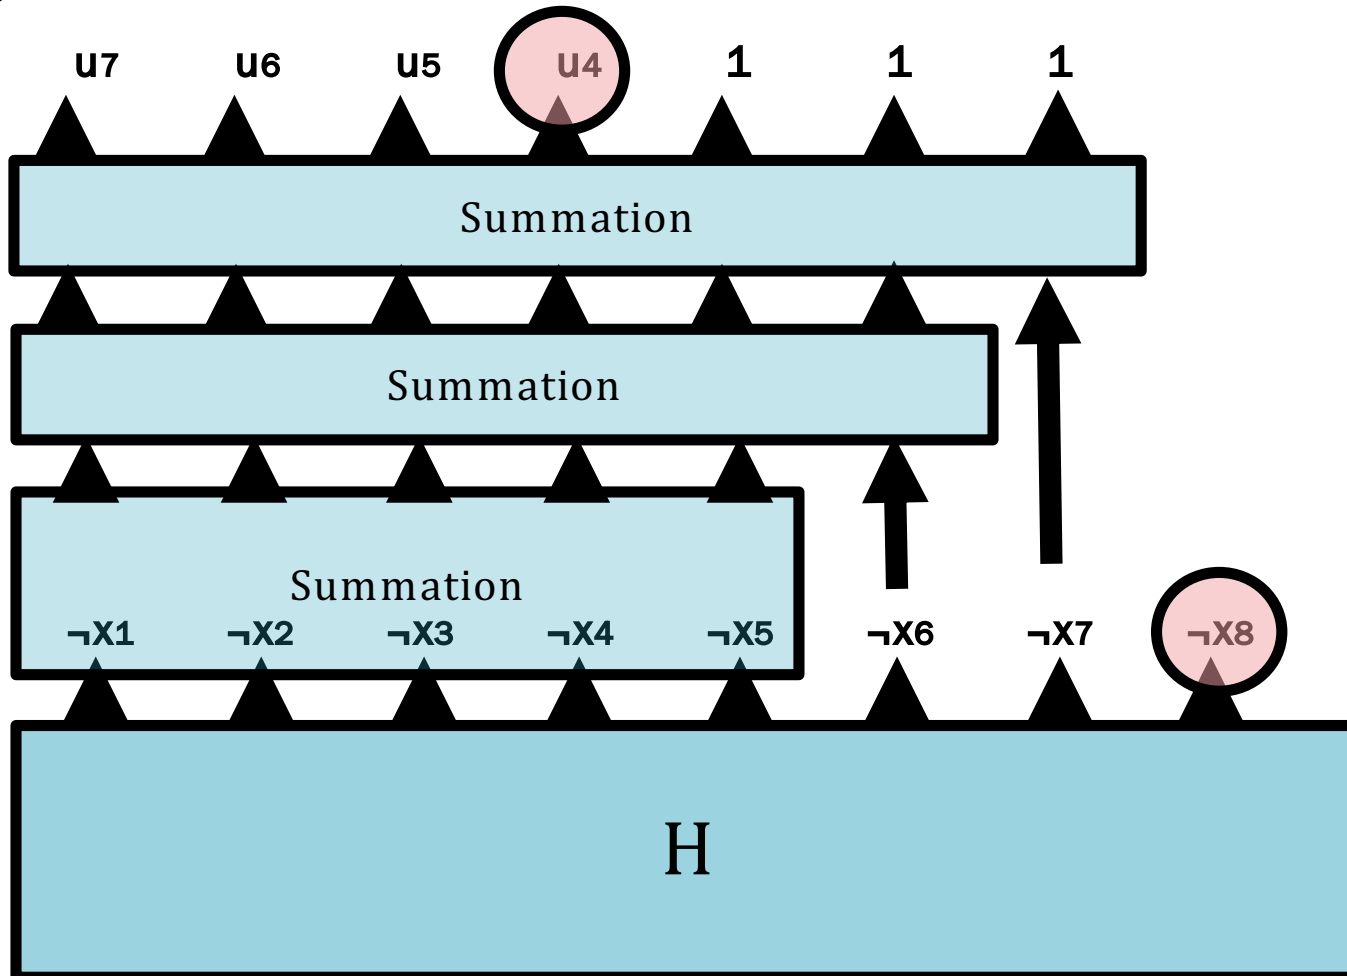
# MUS3 Incremental

- We know that the sum must be at least 2—set  $t_1$  and  $t_2$  to 1.
- Now **SAT\_ASSUME(H U Card,  $\{\neg t_3, x_7, x_8\}$ )**,
- Allow two softs to be falsified among  $x_1$ — $x_6$



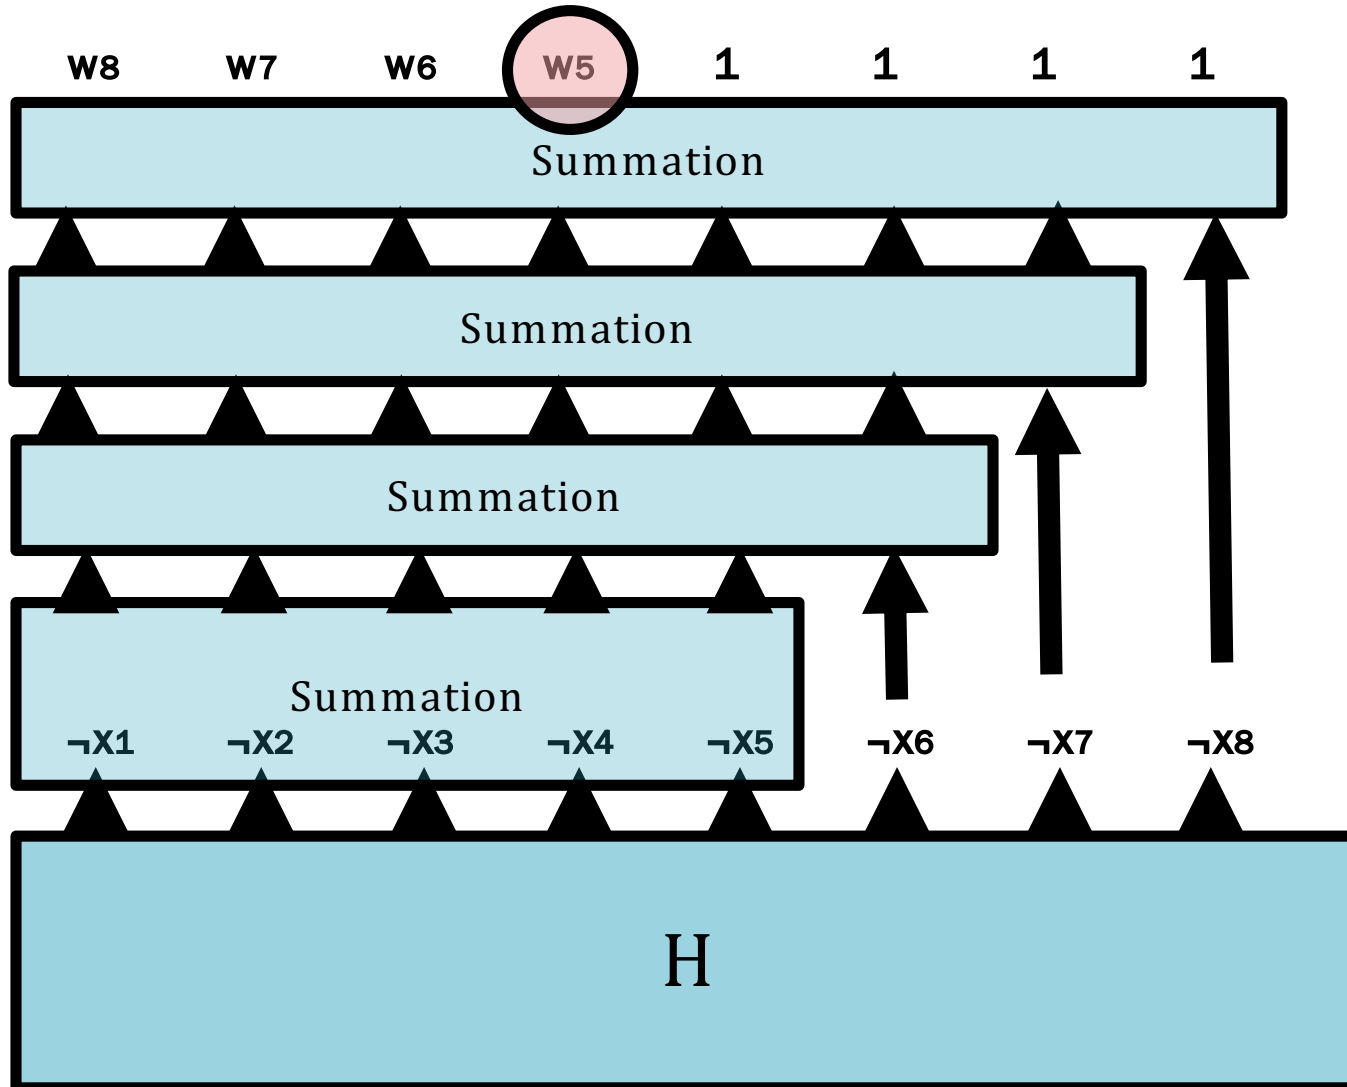
# MUS3 Incremental

- Another core ( $t_3, \neg x_7$ ), add 1 to the overall cost.
- Now **SAT\_ASSUME(H U Card, { $\neg u_4, x_8$ })**, allow up to 3 falsified softs




# MUS3 Incremental

- Another core ( $u_4, \neg x_8$ ); Add one to overall cost. Then  $\text{SAT\_ASSUME}(H \cup \text{Card}, \{\neg w_5\})$



# MUS3 Incremental

- SAT – optimal solution with overall cost = accumulated overall cost



# OLL

- ▶ Used in RC2.
- ▶ A. Morgado, C. Dodaro, and J. Marques-Silva. **Core-guided MaxSAT with soft cardinality constraints**. CP 2014
- ▶ Uses summation circuits like MUS3 but links up the summation circuits in a more flexible way.

# OLL

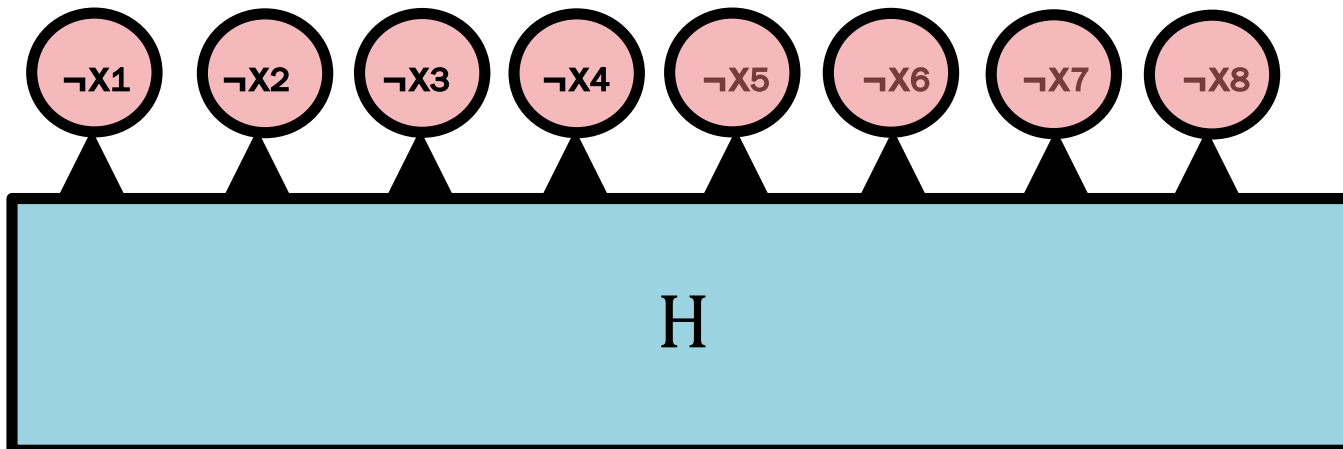
With same example

Start with Cardinality Constraint Layer with inputs only

 Literals assumed FALSE


First SAT solve is  $\text{SAT\_ASSUME}(H, \{x_1, x_2, \dots, x_8\})$ ,

i.e., try to falsify zero softs.



# OLL

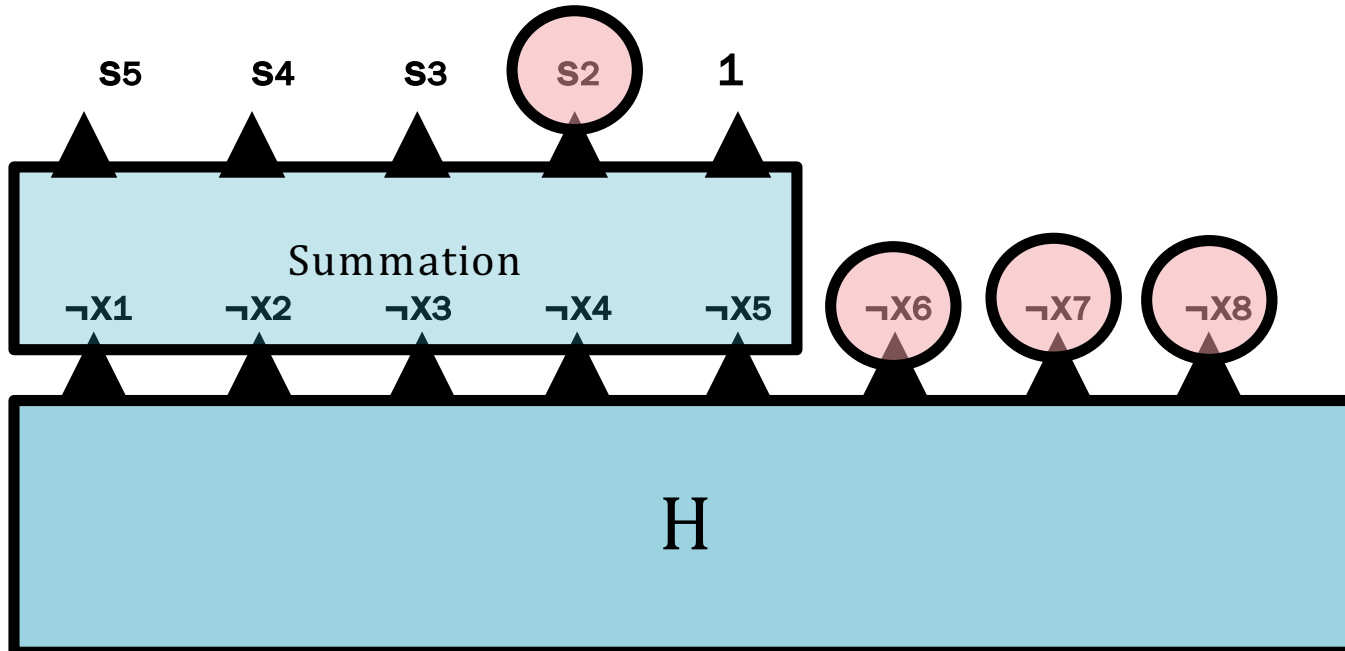
- UNSAT; say we get the core  $(\neg X_1, \neg X_2, \neg X_3, \neg X_4, \neg X_5)$
- Build summation network over softs in core






# OLL

- We know that the sum must be at least one—set  $s_1 = 1$ .
- Now **SAT\_ASSUME(H U CARD, { $\neg s_2, x_6, x_7, x_8$ })**,
- Allow one soft to be falsified among  $x_1$ — $x_5$




# OLL

- Get another core
- Get another conflict, e.g.,  $(s_2, \neg x_6)$

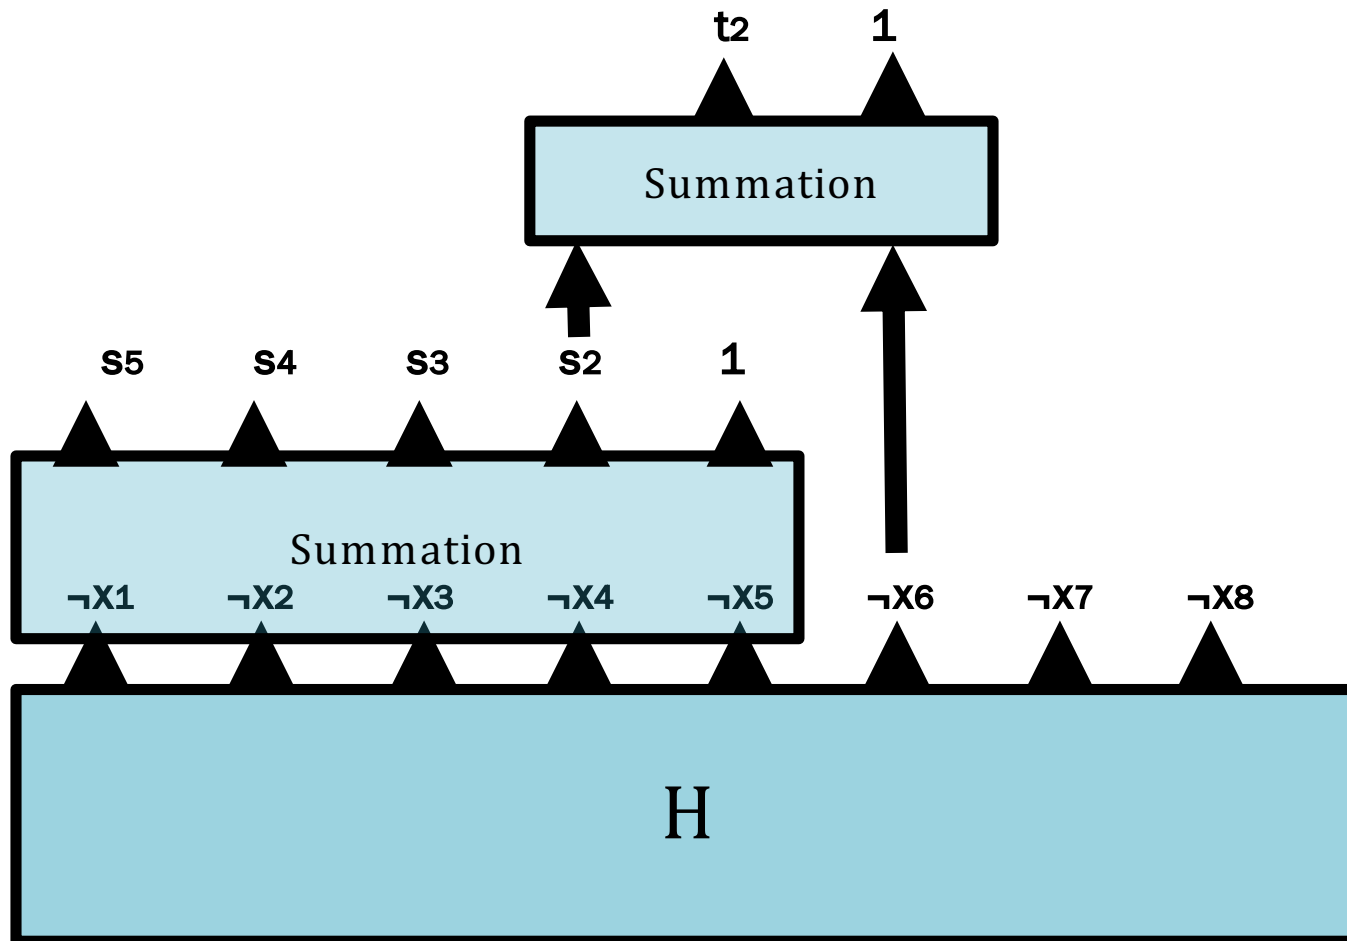


# OLL

- Like MSU3 we create a new summation constraint but unlike MSU3 we do not subsume the prior summation constraint into the new one. Only the literals of the conflict.

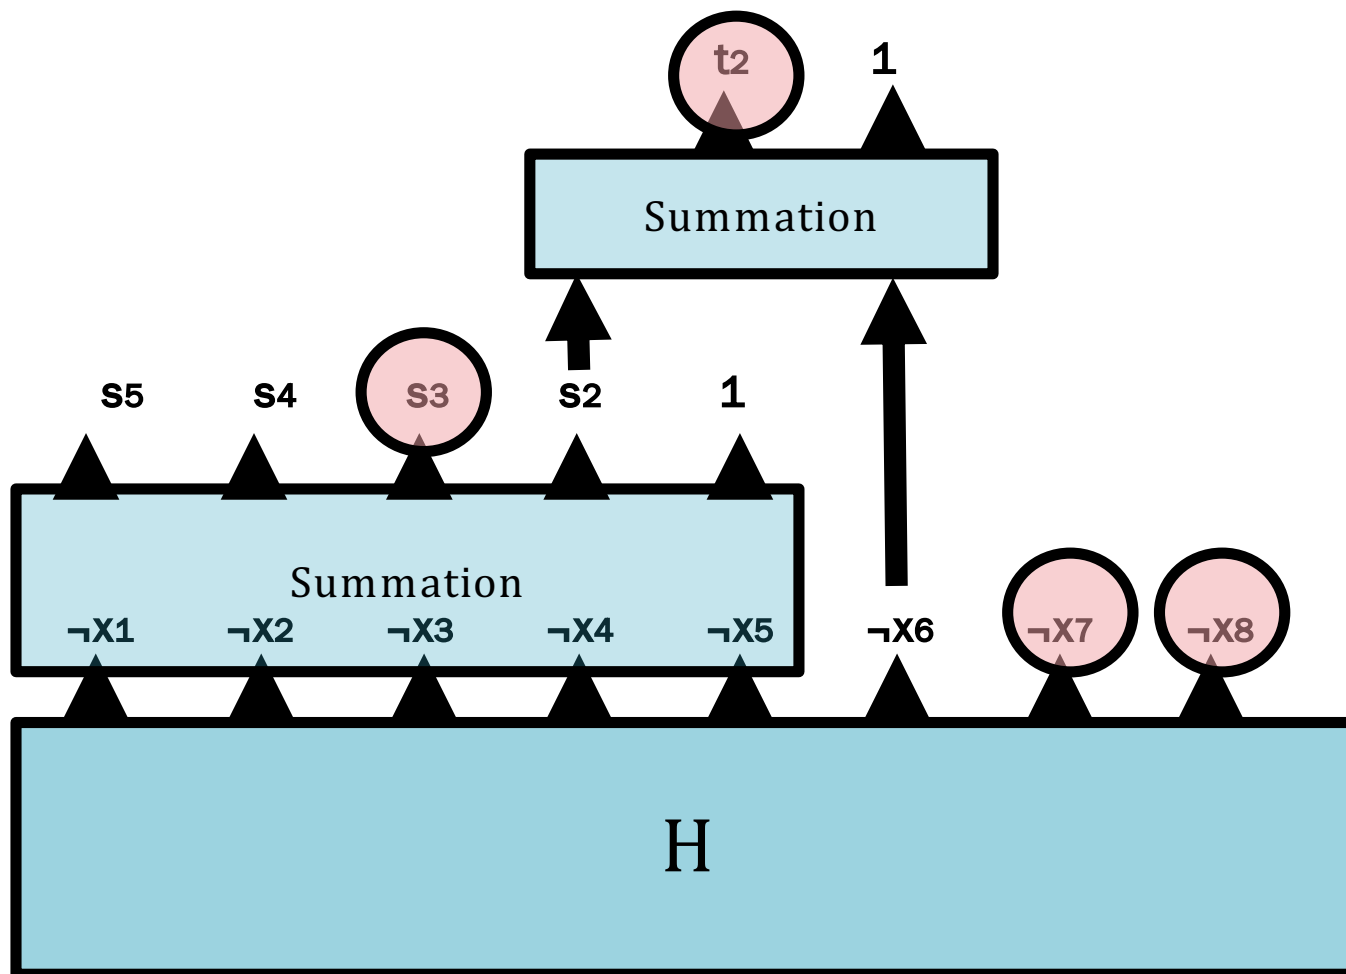


# OLL



# OLL

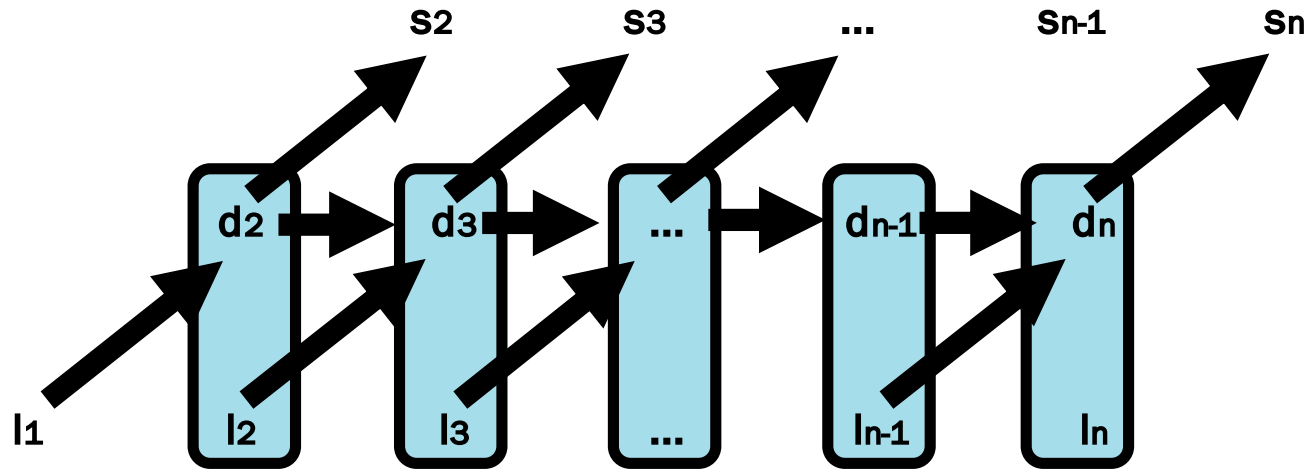
- Then **SAT\_ASSUME(H U CARD, { $\neg s_3, \neg t_2, x_7, x_8$ })**...continue in this way



# PMRES

- ▶ Used in Eva500a.
- ▶ Narodytska, N., Bacchus, F. **Maximum satisfiability using core-guided MaxSAT resolution.** AAI 2014
- ▶ Does not use summation circuits, rather it uses a circuit that detects if more than one literal is true.


# PM-Res



- ▶ The  $d_i$  and  $s_i$  variables are new.
- ▶ We make  $d_i$  if  $d_{i-1}$  or  $l_{i-1}$  are true.
  - ▶  $d_i$  encodes that at least one of the first  $i-1$  inputs is true.
- ▶  $d_i \wedge l_i \rightarrow s_i$  one of the first  $i-1$  inputs was true and  $l_i$  is true (i.e., sum of  $l_1 \dots l_i$  is  $> 1$  AND  $l_i$  is true)

# PM-Res

- ▶ Draw this as below.





# PMRes


Using same example

**Start with Cardinality Constraint Layer with inputs only**

 Literals assumed FALSE

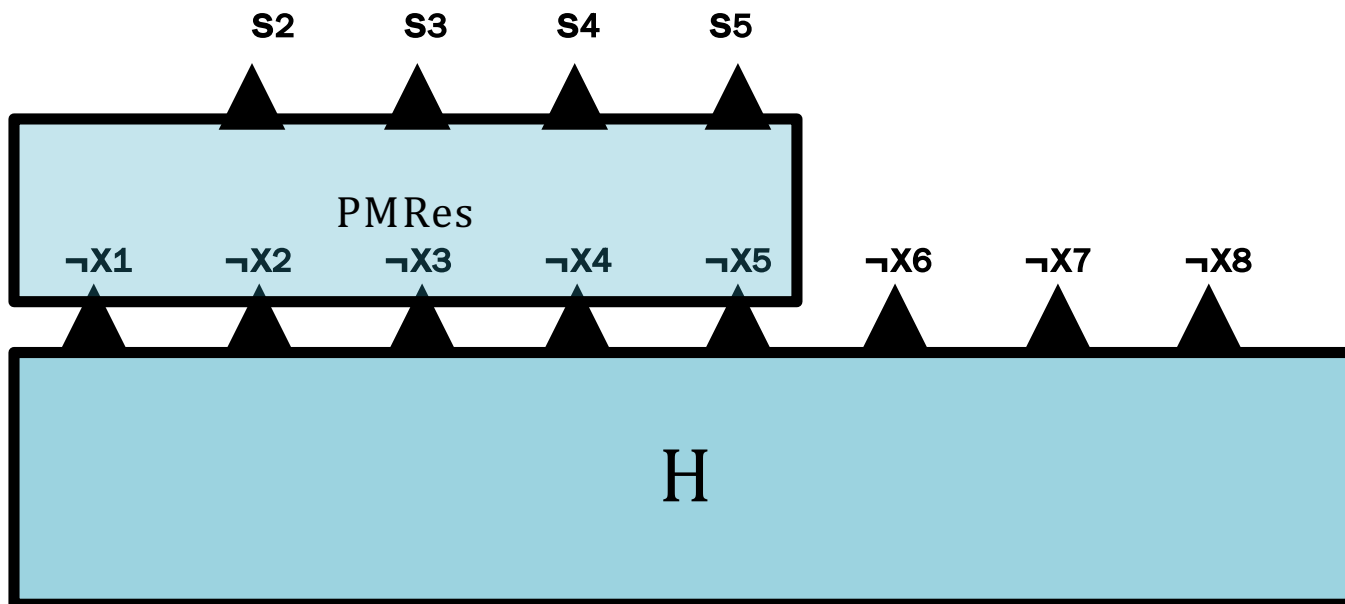
**First SAT solve is  $\text{SAT\_ASSUME}(H, \{x_1, x_2, \dots, x_8\})$ ,**

I.e., try to falsify zero softs.



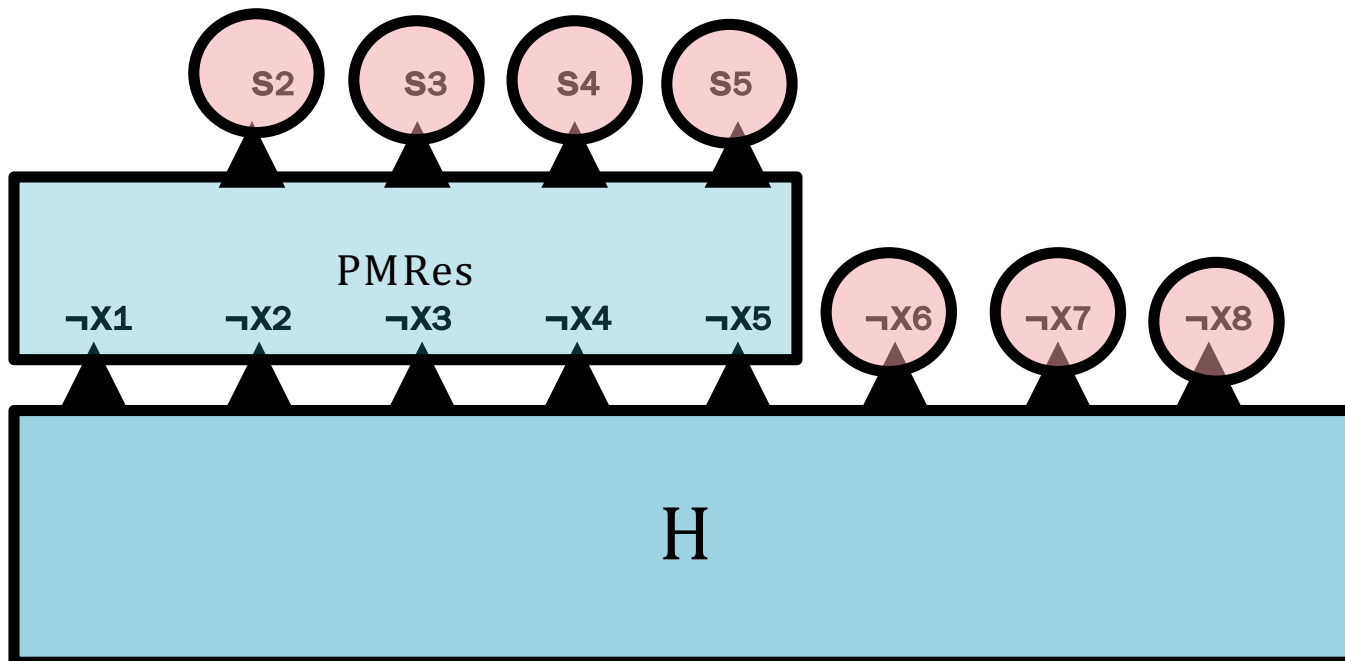
# PMRes

- UNSAT; say we get the core  $(\neg X_1, \neg X_2, \neg X_3, \neg X_4, \neg X_5)$
- Build PMRes circuit over softs in core.
- In the new formula (with the PMRes circuit) we always assume the negation of all outputs of the cardinality layer




# PMRes

- Now **SAT\_ASSUME(HUCard, { $\neg S_2, \neg S_3, \neg S_4, \neg S_5, X_6, X_7, X_8$ })**,
- $\rightarrow$  We can falsify at most one of the softs  $X_1, \dots, X_5$  but no other softs.
- UNSAT (we must falsify at least 4)



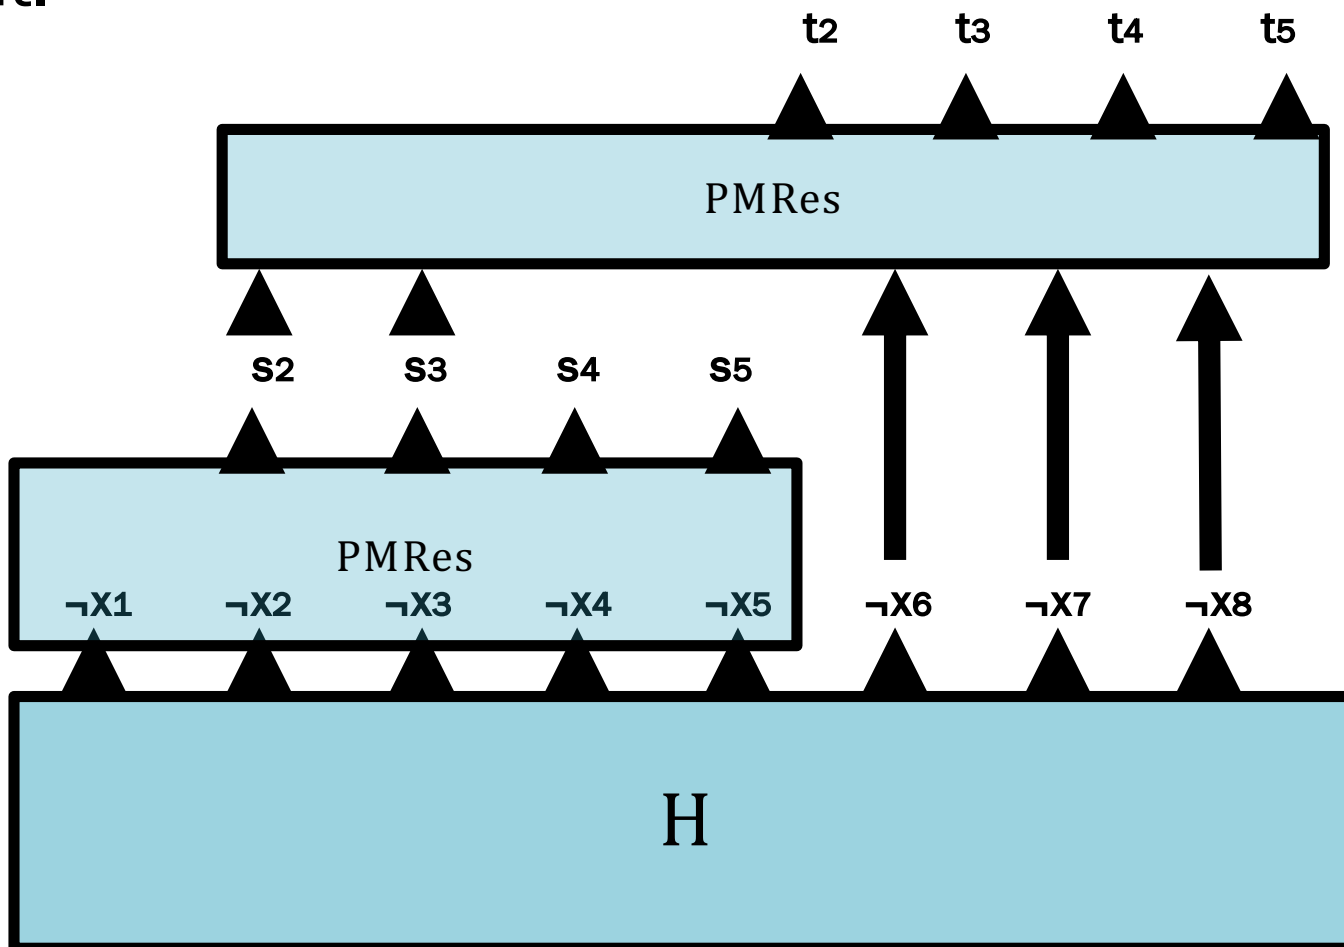
# PMRes

- Various conflicts, say **SAT\_ASSUME(H, { $\neg S_2, \neg S_3, \neg S_4, \neg S_5, x_6, x_7, x_8$ })** returns the core  $(x_6, x_7, x_8, \neg S_2, \neg S_3)$
- PMRes then builds a new PMRes circuit with these literals as input.




# PMRes

- new core ( $x_6, x_7, x_8, \neg s_2, \neg s_3$ )
- PMRes then builds a new PMRes circuit with these as input.




# PMRes

- $\text{SAT\_ASSUME}(H, \{\neg s_4, \neg s_5, \neg t_2, \neg t_3, \neg t_4, \neg t_5\})$
- It gets complex...




# PMRes

- But subject to these assumptions at most two original softs can be falsified.



# PMRes

- PMRes continues in this way until it achieves SAT.






# Dealing with Weights.

- ▶ These algorithms employ the technique of **clause cloning** to deal with weighted softs.
- ▶ When a core is found e.g.,  $\{l_1, l_2, l_3, l_4\}$  where the soft clauses ( $l_i$ ) have different weights.
  - ▶ Let  $\min W_t$  be the minimum weight among the softs in the core.
  - ▶ Add  $\min W_t$  to the overall cost (instead of 1)
  - ▶ Make a copy of those literals that have cost greater than  $\min W_t$ . Give these copies weight = original weight -  $\min W_t$

# Weighted instances

Say  $S = \{(x_1), (x_2), (x_3), (x_4)\}$  with  $wt(x_i) = i$

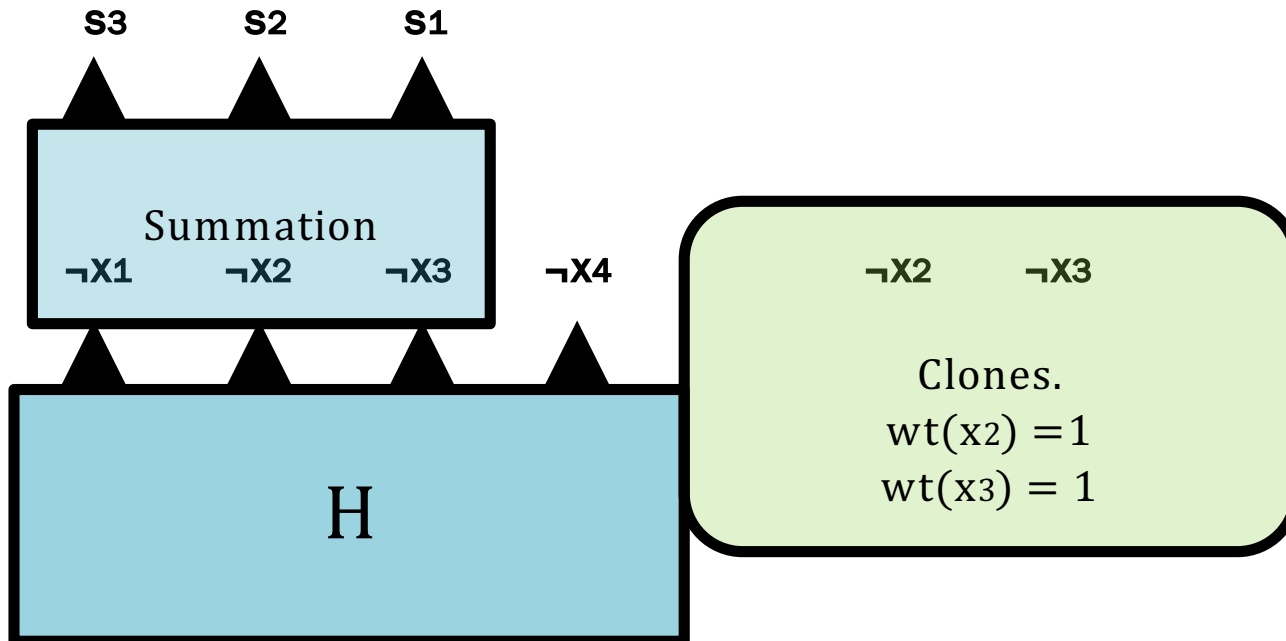
First SAT solve is  $SAT\_ASSUME(H, \{x_1, x_2, x_3, x_4\})$ ,



# Weighted instances

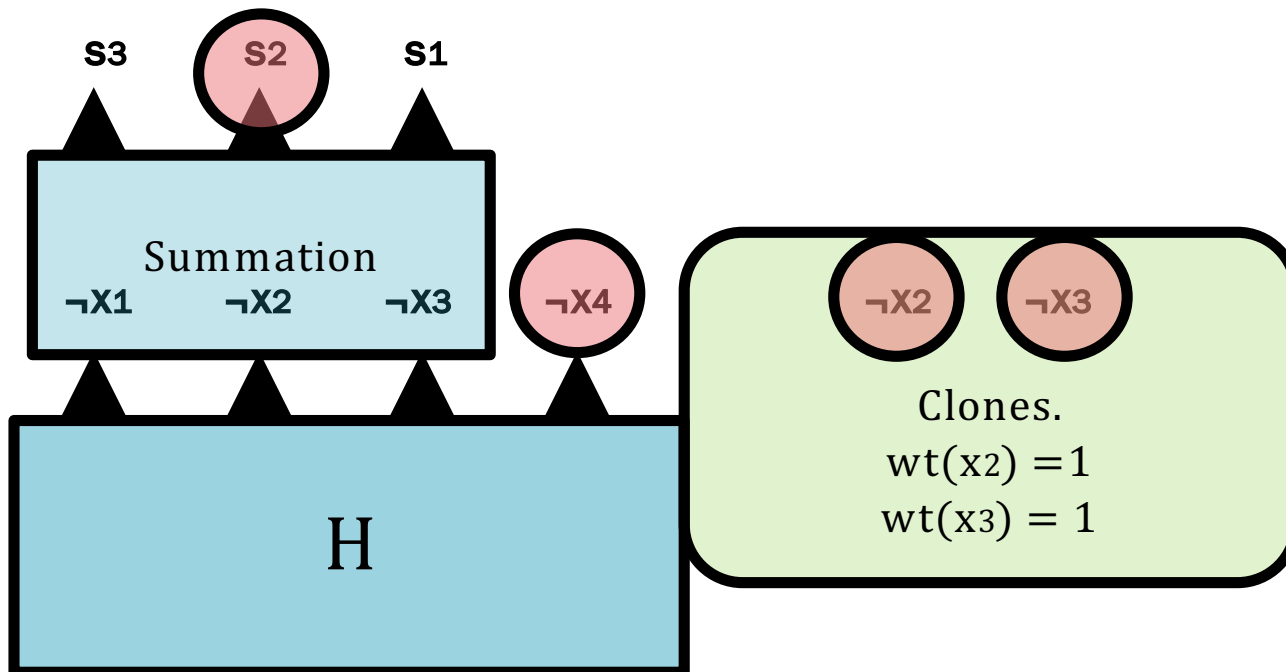
UNSAT; say we get the core  $(x_1, x_2, x_3)$

As before build summation network over softs in core.  
 $\text{minWt}=1$ , so we must clone  $x_2$  and  $x_3$ .



# Weighted instances

Now  $\text{SAT\_ASSUME}(H, \{\neg s_2, x_4, x_2, x_3\})$ ,



# Solving a sequence of relaxations

- ▶ The approach can be very effective. Much more effective than the naive approach of restricting the sum over all falsified soft clauses.
- ▶ The discovered cores are exploited in a non-trivial manner to constraint the search for the set of soft clauses that can be falsified.

# Solving a sequence of relaxations

- ▶ The structure of the Cardinality layer becomes quite complex, and although clear empirical differences can be observed among the different ways of constructing the Cardinality layer, there is no real understanding of this.
- ▶ Instance independent question:
  - ▶ From each UNSAT result we extract a conflict clause. Are the current ways of constructing the cardinality layer fully exploiting the information in those conflicts? How can we formalize this question.
- ▶ Instance dependent question:
  - ▶ Each way of constructing the cardinality layer ends up posing a different type of query to the SAT solver. Do instances with particular structure work better with particular ways of constructing that layer?

# Solving a sequence of relaxations

- ▶ As the cardinality layer grows the SAT solver has a harder and harder time solving it.
- ▶ The approach of clause cloning for dealing with weighted instances is limited.
  - ▶ In the 2018 MaxSat Evaluation a OLL solver RC2 solved more weighted instances among those tested than any other solver.
  - ▶ But clearly less effective the implicit hitting set solver MaxHS when run on a larger benchmark suite.
  - ▶ On instances where the # of distinct weights  $> 3$ , the difference is more profound.
- ▶ On unweighted instances (all softs have the same weight) sequence of sat approaches are generally superior to implicit hitting set approaches.

## Implicit Hitting Set Solvers



# Implicit Hitting Set (IHS) Approach to MaxSat

- ▶ First developed by [Davies PhD]
- ▶ IHS solvers utilize both a SAT solver and an Integer Program solver (IP)

# Implicit Hitting Set (IHS) Approach to MaxSat

- ▶ Unlike the previous approaches, IHS solvers never modify the input MaxSat  $\mathbf{F}$ .
  - ▶ The SAT solver is always run on instances that are no more complex than the input  $\mathbf{F}$ .
- ▶ The cores exploited by IHS solvers are cores of the input formula  $\mathbf{F}$  (not cores of  $\mathbf{F}$  augmented by cardinality constraints).
- ▶ All numeric reasoning about weights is delegated to an Integer Programming solver (e.g., CPLEX)
  - ▶ designed for optimization
  - ▶ weights can be floating point numbers
  - ▶ the underlying LP + Cuts approach is very powerful

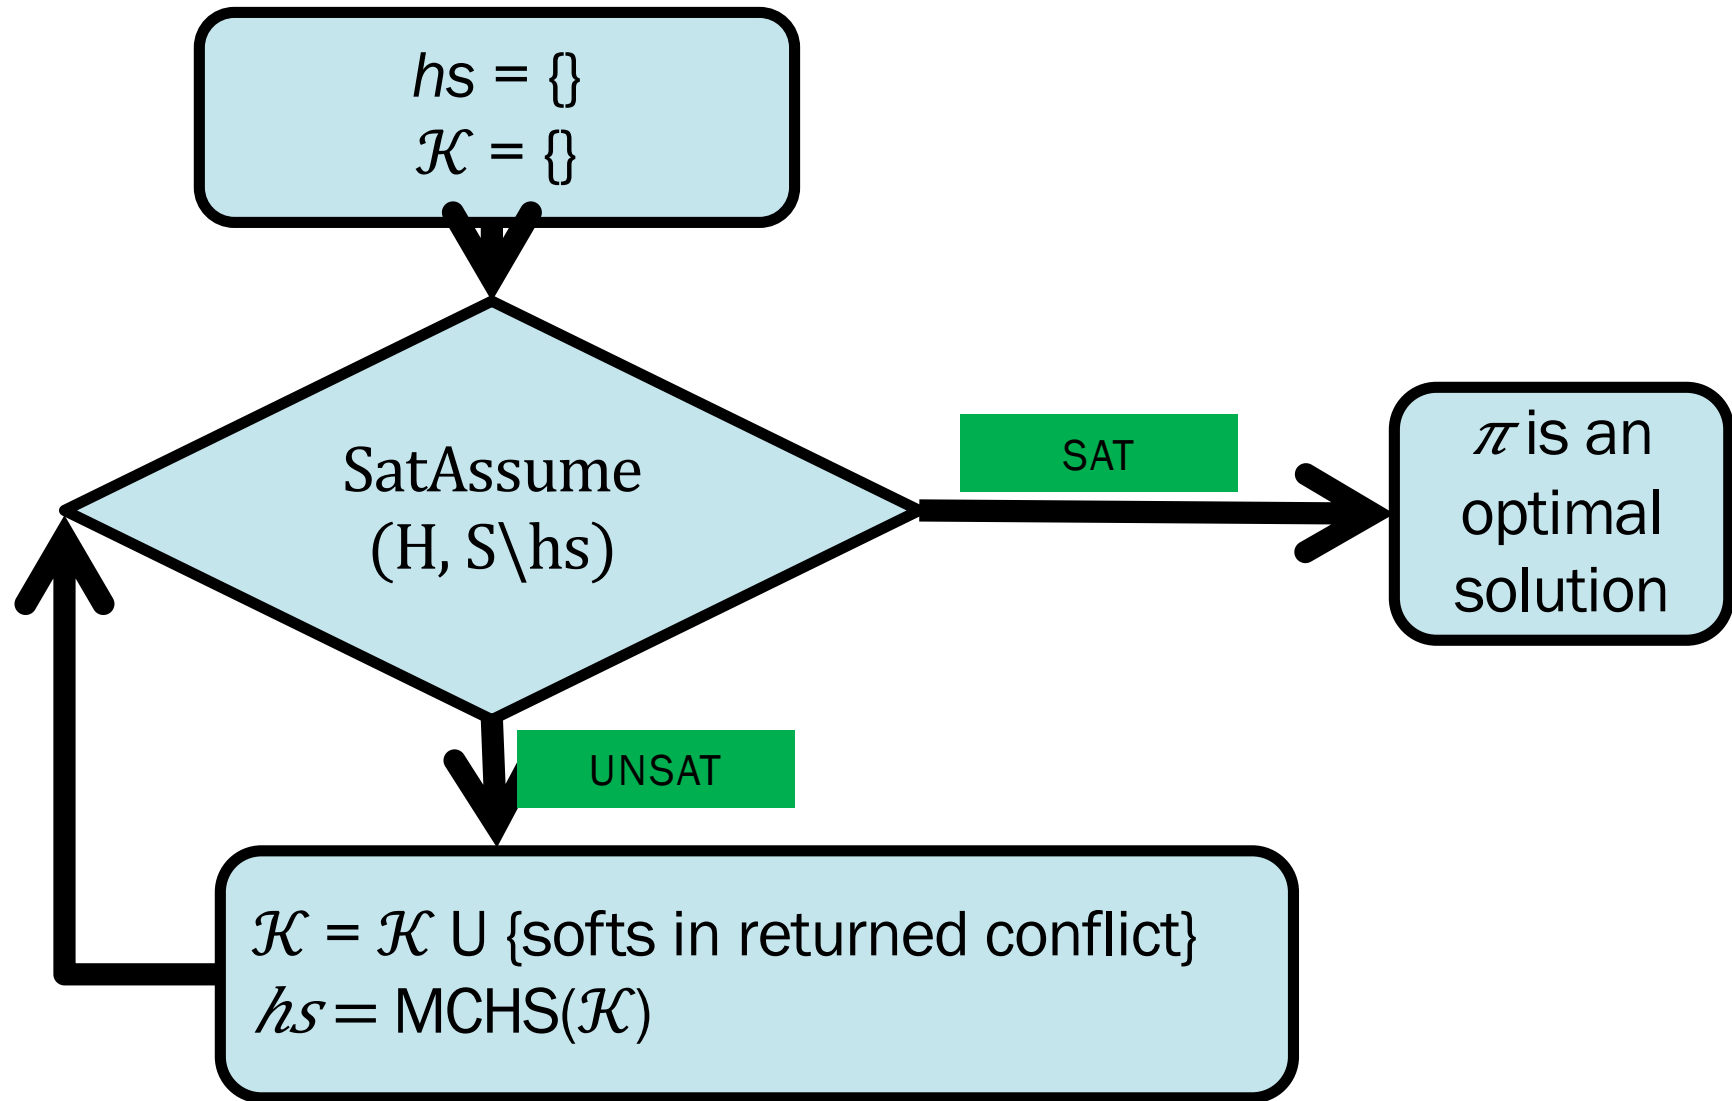
# Cores and hitting sets

- ▶ Remember
  - ▶ A set of soft clauses  $\kappa \subseteq \mathbf{S}$  is a **core** of  $\mathbf{F}$  if  $\kappa \cup \mathbf{H}$  is **UNSAT**
  - ▶ Feasible solutions satisfy the hard clauses  $\mathbf{H}$
- ▶ Let  $\mathbf{K}$  be any set of cores of  $\mathbf{F}$  and  $\pi$  any feasible solution.  **$\pi$  must falsify** at least one soft clause of every core in  $\mathbf{K}$ .
- ▶ Let  $A = \{c \mid \pi \not\models c\}$  be the set of clauses falsified by  $\pi$
- ▶ Then  $A$  is a hitting set of  $\mathbf{K}$  (non-empty intersection with every member of  $\mathbf{K}$ ).

# Cores and hitting sets

- ▶ Let  $\text{MCHS}(K)$  be a minimum cost hitting set of  $K$ —this is a set of soft clauses.
- ▶ For every feasible solution  $\pi$   
$$\text{cost}(\pi) = \text{wt}(A) \geq \text{wt}(\text{MCHS}(K))$$
- ▶ The weight of a minimum cost hitting set of **any** set of cores is a **lower bound** on the cost of an optimal solution.
- ▶ Therefore, for **any** set of cores  $K$  and **any** feasible solution  $\pi$  if  $\text{cost}(\pi) = \text{wt}(\text{MCHS}(K))$ ,  $\pi$  must be an **optimal solution**.
- ▶ This leads to a simple algorithm for finding an optimal solution.

# The IHS Algorithm



# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

*$\pi$  satisfies  $H$  and all soft clauses except possibly the softs in  $hs$ . So  $cost(\pi) \leq wt(MCHS(\mathcal{K}))$*

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = MCHS(\mathcal{K})$

# The IHS Algorithm

If UNSAT the SAT solver the conflict returned is a core

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{MCHS}(\mathcal{K})$

# The IHS Algorithm

The returned core must be new, not previously in  $\mathcal{K}$ —the new core contains no softs from  $hs$ , but every core in  $\mathcal{K}$  contains a soft of  $hs$ .

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{MCHS}(\mathcal{K})$



# The IHS Algorithm

This process must terminate as there are only a finite number of cores.

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

UNSAT


$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{MCHS}(\mathcal{K})$

# IP solver used to compute MCHS (when an MCHS is needed)

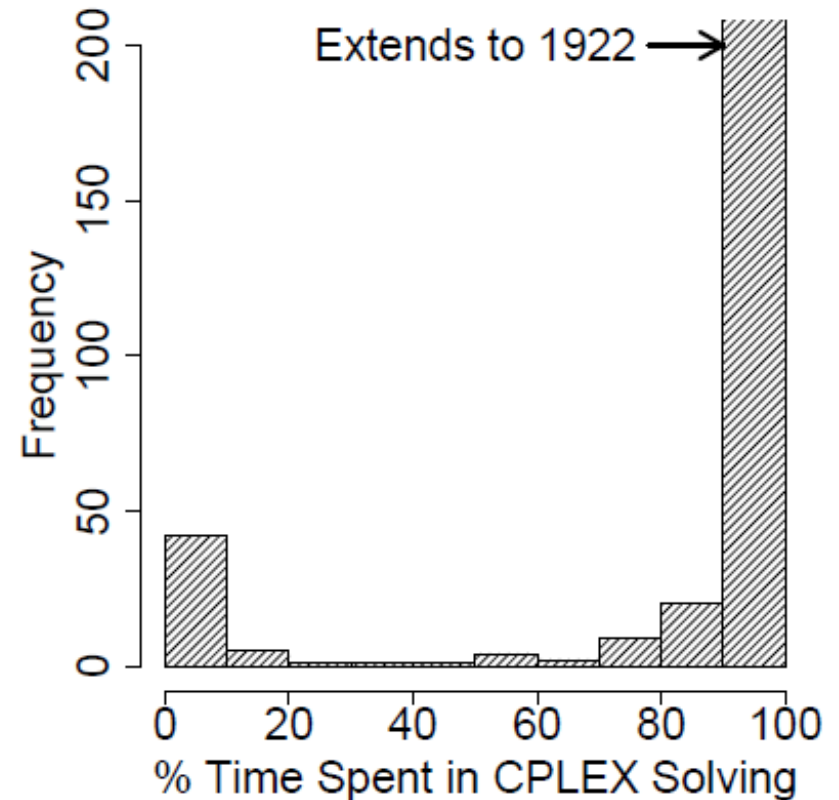
- ▶ The MCHS (aka, set-cover) problem is an NP-Hard optimization problem. But in practice it can usually be solved efficiently by an integer programming solver.
  - ▶ Typically IBM's CPLEX is used
  - ▶ Seems to be by far the most effective way of finding an MCHS.
- ▶ The approach is most closely related to the implicit hitting set formalism of Karp. It can also be viewed as being a CEGAR or a logic based Benders approach.
- ▶ The direct encoding of MaxSat into an IP is far less effective than this hybrid approach.

# The IHS Algorithm

- ▶ This basic IHS algorithm is not that effective.



Solved



Unsolved

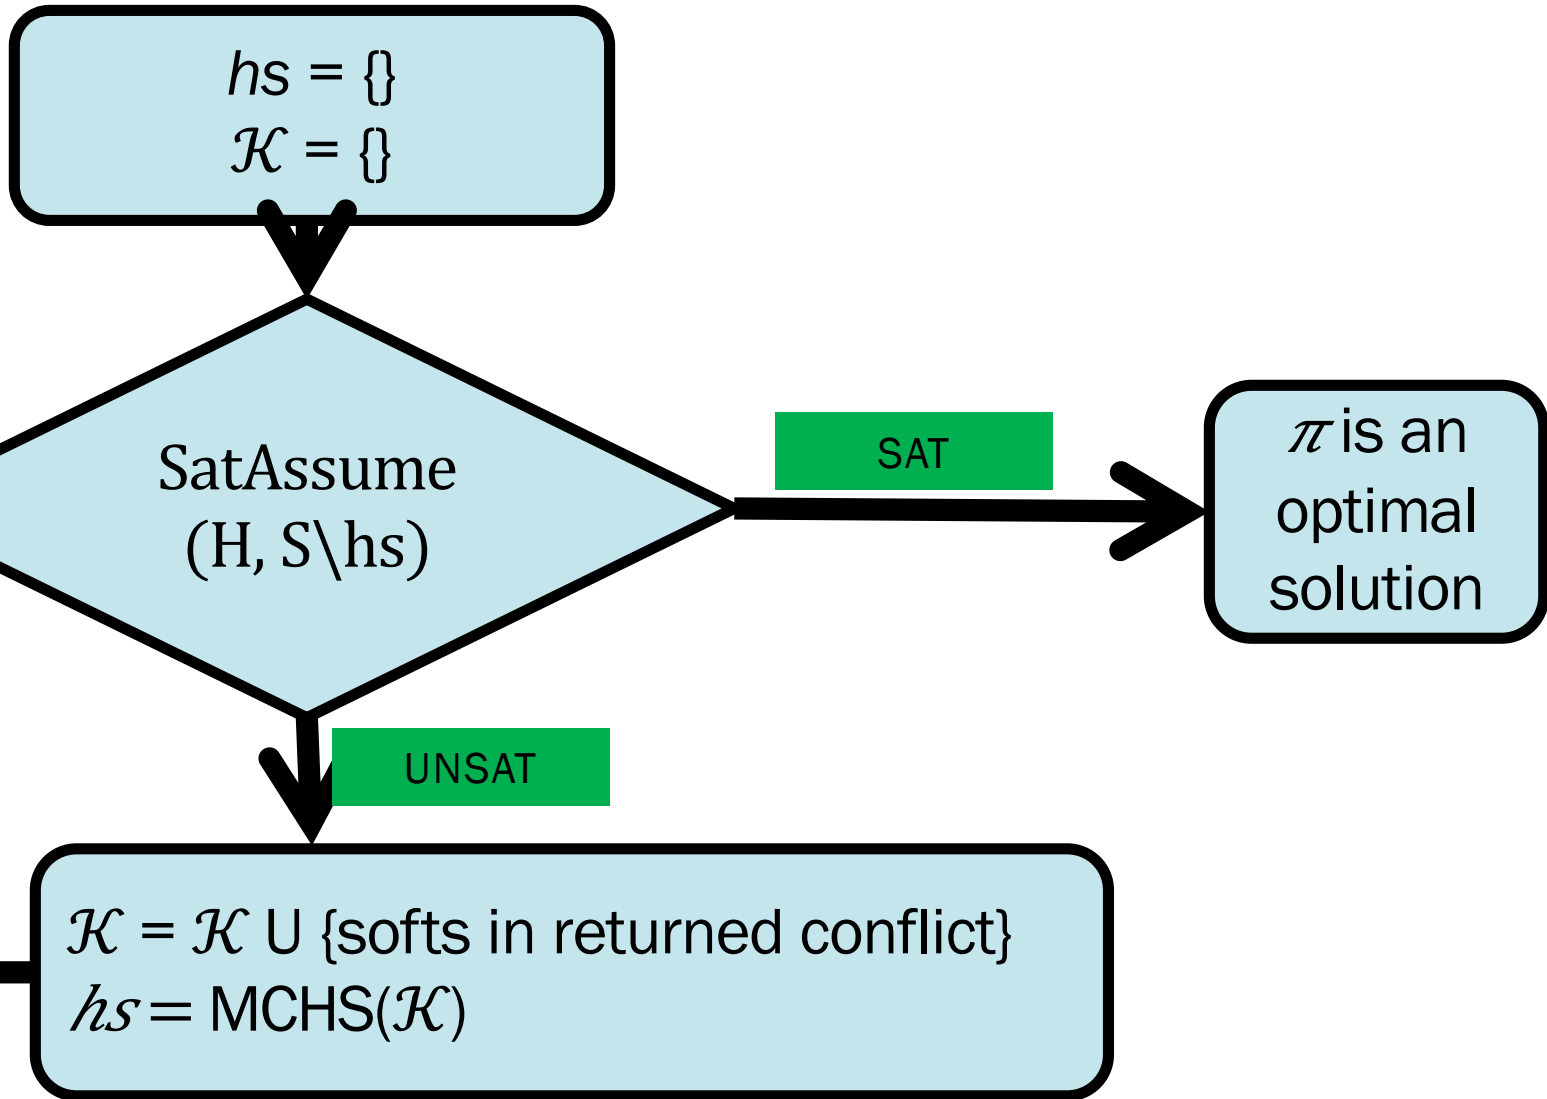
# The IHS Algorithm

- ▶ The IP hitting set model is being incrementally improved by adding new cores.
- ▶ Generally many cores have to be accumulated before the IP model is strong enough to yield hitting sets whose removal yields SAT.
- ▶ Always computing an MCHS on these “too weak” IP models becomes very expensive.

# The IHS Algorithm

- ▶ Various are employed techniques to improve the IP model more quickly
- ▶ Most importantly computing an MCHS can be delayed and performed only occasionally. Much cheaper to compute non-minimum hitting sets can be used instead.
- ▶ This leads to a different formulation of IHS algorithms (this formulation is what is used in current IHS solvers).

# The IHS Algorithm



# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{MCHS}(\mathcal{K})$

Use a non  
minimum cost  
hitting set  
instead.

# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

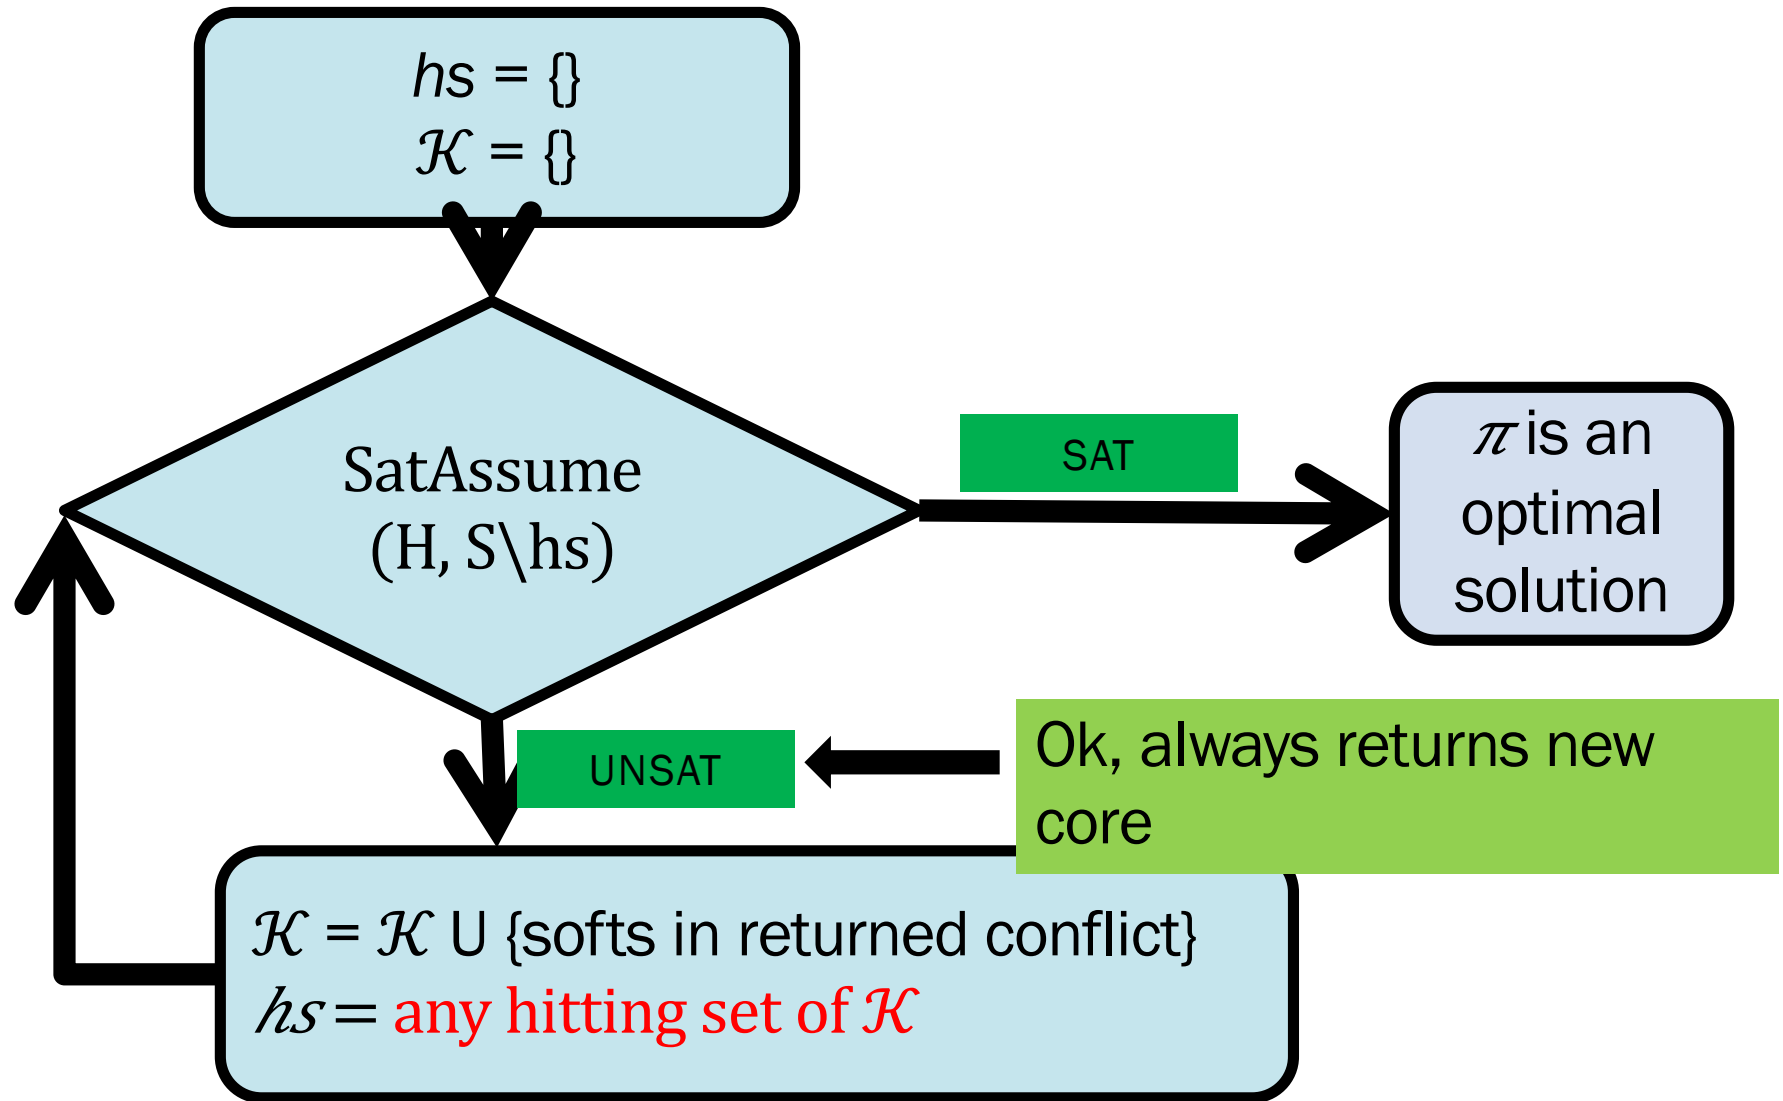
UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$

Use a non  
minimum cost  
hitting set  
instead.



# The IHS Algorithm



# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

$\pi$  is an  
optimal  
solution

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$

But now, we  
cannot conclude  
 $\pi$  is optimal

# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

However  $\pi$   
might be lower  
cost model than  
we have seen  
before

If  $\pi$  is the cheapest  
model found  
install as new  
incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$

# The IHS Algorithm

We must continue

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

If  $\pi$  is the cheapest  
model found  
install as new  
incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$

# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

If  $\pi$  is the cheapest  
model found  
install as new  
incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$

We must continue

Make sure that  
we don't cycle  
returning the  
same  $hs$  as  
before!

# The IHS Algorithm

$hs = \{\}$   
 $\mathcal{K} = \{\}$

SatAssume  
( $H, S \setminus hs$ )

SAT

To terminate we must occasionally compute an MCHS.


If  $\pi$  is the cheapest model found install as new incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$

To terminate we must occasionally compute an MCHS.

$hs = \{\}$   
 $\mathcal{K} = \{\}$



SAT


If  $\pi$  is the cheapest model found install as new incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$   
Occasionally (via some policy):  
 $hs = \text{MCHS}(\mathcal{K})$

MCHS provides a lower bound!

$hs = \{\}$   
 $\mathcal{K} = \{\}$



SAT

If  $\pi$  is the cheapest model found install as new incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$   
 $hs = \text{any hitting set of } \mathcal{K}$   
Occasionally (via some policy):  
 $hs = \text{MCHS}(\mathcal{K}); \text{LB} = \text{wt}(hs)$



$hs = \{\}$   
 $\mathcal{K} = \{\}$

Lower bound meets  
upper bound  
becomes new  
termination condition.

SatAssume  
(H, S \setminus hs)

SAT

If  $\pi$  is the cheapest  
model found install as  
new incumbent.

If  $LB \geq \text{cost}(\text{incumbent})$   
return incumbent

UNSAT

$\mathcal{K} = \mathcal{K} \cup \{\text{softs in returned conflict}\}$

$hs = \text{any hitting set of } \mathcal{K}$

Occasionally (via some policy):

$hs = \text{MCHS}(\mathcal{K}); LB = \text{wt}(hs)$

If  $LB \geq \text{cost}(\text{incumbent})$  return incumbent

# IHS Algorithm

- ▶ As long as computing an MCHS is never “starved” (i.e., always eventually we compute the MCHS) the algorithm must terminate.
- ▶ Maintaining an UB model also allows the IP technique of reduced cost fixing to be exploited
  - ▶ Fahiem Bacchus, Antti Hyttinen, Matti Jarvisalo, and Paul Saikko; **Reduced Cost Fixing in MaxSAT, CP 2018**

# IHS on our Example.

- $F = H \cup S$
- $H = \sum_{i=1}^{2n} x_i < n$       Generalized to  $2n$  unit softs.
- $S = \{(x_1), (x_2), \dots, (x_{2n})\}$
- Optimal solution has cost  $n+1$ .
- [Davies 2013] The IHS approach needs to accumulate at least  $\binom{2n}{n}$  cores in  $\mathcal{K}$  before  $\text{wt}(\text{MCHS}(\mathcal{K})) = n+1$ .
  - For any set of cores  $X$  of size with  $|X| < \binom{2n}{n}$ ,  $\text{MCHS}(X) < n + 1$
- So the IHS approach has to make an exponential number of calls to the SAT solver in this case.



# Weighted Case (Intuition)

- $H = \left\{ \begin{array}{cccc} (\neg y_{11} & \dots & \neg y_{1k} & \neg Z) \\ \vdots & \vdots & \vdots & \\ (\neg y_{m1} & \dots & \neg y_{mk} & \neg Z) \end{array} \right\}$
- $S = \{(y_{11}), \dots, (y_{mk}), (z)\}$
- Weight of all softs is 2 except weight of  $\text{wt}(\mathbf{z})=2^{k-1}$ .
- ▶ Every clause specifies a core. IHS would find  $k$  cores, all by unit prop.
- ▶ In the  $i$ -th iteration ( $1 \leq i \leq k-1$ ) the MCHS would have weight  $2^i$ , and would consist of one  $y_{ij}$  from each of the  $i$  cores found so far.
- ▶ At step  $k$ , the MCHS would be computed to be  $\{(z)\}$ , with cost  $2^{k-1}$ .
- ▶ With  $hs = \{(z)\}$  we get SAT.

## Weighted Case.

- $$H = \left\{ \begin{array}{cccc} (\neg y_{11} & \dots & \neg y_{1k} & \neg Z) \\ \vdots & \vdots & \vdots & \\ (\neg y_{m1} & \dots & \neg y_{mk} & \neg Z) \end{array} \right\}$$
- $S = \{(y_{11}), \dots, (y_{mk}), (Z)\}$
- Weight of all softs is 2 except weight of  $\text{wt}(\mathbf{z})=2^{k-1}$ .
- ▶ The sequence of sat approaches with clause cloning would do something more complex.
- ▶ Each UNSAT call would add a cardinality constraint over the variables in the core, and would clone  $z$ .
- ▶ So the formula gets harder to solve—not clear if this could lead to a non-polynomial separation.

# Implicit Hitting Set Solvers

- ▶ The SAT solving episodes are much simpler—they involve restrictions of the original MaxSat formula rather than augmentations of that formula.
- ▶ In practice, like other CEGAR approaches, only a few thousand cores need to be generated before the MCHS lower bound meets the optimal cost.
- ▶ But there are other cases where the number of cores required is too large, making both finding them and solving the MCHS too expensive.
- ▶ Currently it is usually a more effective way of dealing with a diverse collection of weights.
  - ▶ It would be good get some theoretical results here.

# Implicit Hitting Set Solvers.

- ▶ As with the sequence of SAT approaches, it is not clearly understood how to guide the SAT solver to find good cores (cores that move the MCHS bound up more rapidly).

**Thank you**