


# To Link or not to Link? That is a Watch List!

Robin Coutelier  

TU Wien, Vienna, Austria

---

## Abstract

The two-watched literal scheme is a powerful method used in state-of-the-art SAT solvers to reduce the number of clauses checked during Boolean constraint propagation. In this paper, we explore the representation of watch lists using a linked list data structure. We explain why this representation intuitively has advantages over the traditional array-based representation. We then empirically evaluate the performance of this representation and explain why it is not used in practice when combined with the blocker optimization. Both paradigms were implemented in the NapSAT SAT solver. Based on the implementation process, we detail difficulties raised by a linked list representation. Experimenting with this alternative representation shows insights into cache behaviors. However, we conclude that it should not be used for general-purpose SAT solvers.

**2012 ACM Subject Classification** Automated Reasoning, Constraint and logic programming

**Keywords and phrases** Watcher Lists, SAT, CDCL

**Digital Object Identifier** 10.4230/LIPIcs.POS.2024.1

**Supplementary Material** NAPSAT (*Source Code GitHub*): <https://github.com/RobCoutel/NapSAT>

**Acknowledgements** We thank the reviewers for their valuable feedback. As a reviewer pointed out, a similar approach was presented in [1] and [7]. This work is therefore not novel, but rather an independent re-discovery of an existing method. The author acknowledges support from the ERC Consolidator Grant ARTIST 101002685; the TU Wien Doctoral College TrustACPS; the FWF SpyCoDe SFB projects F8504; the WWTF Grant ForSmart 10.47379/ICT22007. NapSAT is a project started at the University of Liège under the supervision of Prof. Pascal Fontaine and continued at the TU Wien under the supervision of Prof. Laura Kovács. We thank the latter for proofreading the paper.

## 1 Introduction

Modern SAT solvers are based on Conflict Driven Clause Learning (CDCL) [9]. During a typical execution of the CDCL algorithm, most of the computation time is spent on Boolean constraint propagation (BCP). BCP searches for literals that can be implied by the current assignment and a clause via unit propagation. There may be a large number of clauses to inspect during the propagation of a literal  $\ell$ . To reduce this number, the *two watched literals scheme* was introduced [7]. As the name suggests, the idea is to watch each clause with two literals  $\ell_1, \ell_2$ . Each literal  $\ell$  is associated with a watch list  $WL(\ell)$ , which contains all the clauses watched by  $\ell$ . Provided that certain invariants are maintained on the watched literals, this allows one to only check the clauses in the list  $WL(\neg\ell)$  when propagating  $\ell$ , greatly reducing the number of clause visits during CDCL.

This research was conducted without the knowledge of the literature on linked watch lists. Reviewers pointed out that this idea is not novel and has been documented in [1] but was experimented on in [7] before. In this paper, we explore for a third time the representation of watch lists using a linked list data structure. We evaluate our work with a practical implementation in the NAPSAT SAT solver. We detail the implementation difficulties raised by the linked list representation and empirically evaluate the performance of using watch lists via arrays and linked lists. We conclude that the linked list representation is not suitable for general applications.




© Robin Coutelier;  
licensed under Creative Commons License CC-BY 4.0

Pragmatics of SAT.



Leibniz International Proceedings in Informatics  
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 1:2 To link or not to link?



■ **Figure 1** Representation of clauses in NAPSAT. Clauses are stored as a fixed size structure containing a pointer to the literals, the size  $s$  of the clause, and a blocker literal  $b$ .

## 45 2 NapSAT Background

46 We assume that the reader is familiar with the CDCL algorithm and the basic data structures  
47 used in SAT solvers [4, 7]. In this section, we introduce the data structures used in NAPSAT  
48 relevant to understanding the rest of the paper.

49 NAPSAT is a CDCL-based SAT solver written in C++. The code is available on GitHub<sup>1</sup>  
50 and GitLab<sup>2</sup>, and consists of approximately  $\sim 5,800$  loc, among which the core corresponds  
51 to  $\sim 2,100$  loc<sup>3</sup>. NAPSAT supports different variations of chronological backtracking [3]. In  
52 this paper, however, we will only consider the classical non-chronological variant, as it is the  
53 most common in practice.

54 NAPSAT is an experimental solver that is not intended to compete with state-of-the-art  
55 solvers yet. The author therefore acknowledges that some techniques and representations may  
56 not be standard and comparable to more sophisticated solvers. However, most arguments  
57 presented in this paper are general enough to be applied to more advanced solvers. In the  
58 future, we plan to integrate the methods suggested by the reviewers to improve the NAPSAT.  
59 We shall mark discussion on other solvers in a “remark” environment.

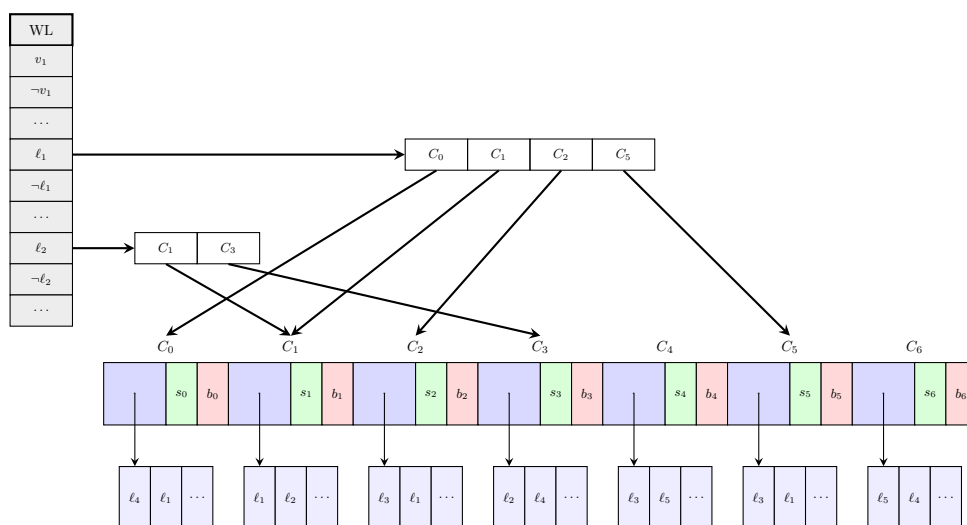
60 In NAPSAT, literals are represented as 32-bits unsigned integers whose least significant  
61 bit indicates the polarity. Clauses are stored as a fixed-size structure containing a pointer  
62 to the literals, the size  $s$  of the clause, and a blocker literal  $b$ . The blocker must be in the  
63 clause and is used to quickly check if the clause is satisfied. Each clause takes 16 bytes of  
64 memory (8 for the pointers to literals, 4 for the size, and 4 for the blocker.). All clauses are  
65 stored in a contiguous vector, allowing us to identify them with a 32-bits unsigned integer  
66 instead of a 64-bits pointer. Figure 1 illustrates the representation of NAPSAT clauses. We  
67 assume that the watched literals are the first two literals in the clause as in [1, 4, 7].

68 ► **Remark 1.** A main difference between NAPSAT and state-of-the-art solvers is the location  
69 of the blockers. In NAPSAT, the blocker is stored in the clause data structure, while in  
70 state-of-the-art solvers, the blocker is stored in the watch list [2, 4]. The advantage of storing  
71 the blocker in the clause is that it allows one to share it between the two watch lists. However,  
72 since most state-of-the-art solvers seem to agree on storing the blocker in the watch list,  
73 we will also implement and evaluate this representation in NAPSAT in the future. In the  
74 following, we consider the data structures used in NAPSAT but discuss how our arguments  
75 can be applied to different techniques.

<sup>1</sup> <https://github.com/RobCoutel/NapSAT>

<sup>2</sup> <https://gitlab.uliege.be/smt-modules/sat-library>

<sup>3</sup> as of commit df5a9ca4



■ **Figure 2** Array-based representation of watch lists. The watch list of  $l_1$  is  $\{C_0, C_1, C_2, C_5\}$  and the watch list of  $l_2$  is  $\{C_1, C_3\}$ .

### 76 3 Watch Lists Representations

#### 77 3.1 Array-Based Representation


78 The traditional representation of watch lists is an array of pointers to an array of references  
 79 to clauses (potentially with blockers). For example, Figure 2 shows the watch lists of  $l_1$  and  
 80  $l_2$ . The watch list of  $l_1$  is  $[C_0, C_2, C_1, C_5]$  and is simply represented as an array of indices to the  
 81 clauses. There exist variations of this representation, such as the use of pointers to the  
 82 clauses instead of indices, as in CADICAL [2].

83 For this paper, we will consider indices as on Figure 2. This representation has the  
 84 advantage of being very simple and flexible. It is easy to implement and does not complicate  
 85 the maintenance of the code base.

86 The complexity of removal of an arbitrary element in the list is  $O(n)$ , where  $n$  is the  
 87 number of clauses in the watch list. However, the solver seldom removes arbitrary clauses  
 88 from the list. More often, a clause is removed as it is inspected during propagation. In  
 89 this case, the complexity is  $O(1)$  since we know the index of the clause in the watch list.  
 90 There are two main approaches to removing an element from a list. Either (i) we swap the  
 91 clause with the last clause in the list and decrement the size of the list, or (ii) we shift all  
 92 the elements after the clause to the left, as in MiniSat [4]. Both approaches are sensible,  
 93 since most of the time, the propagation will continue until the end, and the shifting naturally  
 94 happens during the propagation. In the array-based implementation of NAPSAT, we use  
 95 the first approach (i). Watching a new clause is done in  $O(1)$  time if the watch list is not  
 96 full, and  $O(n)$  otherwise, by simply pushing the clause at the end of the list (maybe after  
 97 reallocation if the list is full).

#### 98 3.2 Linked List Representation

99 When using linked watch lists, we remove one level of indirection. The watch list of a literal  $l$   
 100 is now a simple reference to one of the clauses watched by  $l$ . The clause data structure is now  
 101 enhanced with two clause references, one for each watched literal. Exploring the watch list



■ **Figure 3** Representation of watch lists using the linked list data structure. The watch list of  $\ell_1$  is  $\{C_0, C_1, C_2, C_5\}$  and the watch list of  $\ell_2$  is  $\{C_1, C_3\}$ .

102 of a literal  $\ell$  now only requires following the first pointer of the clause if  $\ell$  is the first watched  
 103 literal, and the second pointer otherwise. This representation is shown in Figure 3. Clauses  
 104 are now 24 bytes long, however, this space is saved from the list of references, compared to  
 105 the array-based representation.

106 Similarly to the array representation, removal of an arbitrary clause from the list is  $O(n)$ ,  
 107 but removal during propagation is  $O(1)$ . It simply requires connecting the previous clause  
 108 to the next clause in the list. Doing this efficiently and elegantly is however not trivial since  
 109 the connection might be either with the first or the second pointer of the previous clause.  
 110 For example, in Figure 3, if the solver removes  $C_1$  from the watch list of  $\ell_1$ , we connect the  
 111 second pointer of  $C_0$  to  $C_2$  since  $\ell_1$  is the second watched literal of  $C_0$  and the first watched  
 112 literal of  $C_1$ . Watching a new clause  $C$  by  $\ell$  is always done in  $O(1)$  time, by connecting  $C$   
 113 to  $WL(\ell)$  and updating  $WL(\ell)$  to  $C$ , thereby pushing  $C$  at the front of the list.

114 **Why linked lists?** Using linked lists has several advantages over the array-based represen-  
 115 tation. First, it allows us to reduce the dereference level by one. Indeed, there is no longer  
 116 a need for a pointer to the array of clause references. A reference to the first clause in the  
 117 list is sufficient. Furthermore, in NAPSAT, when exploring the list, the clause must be  
 118 dereferenced anyway, there does not seem to be an extra cost of using this technique.

119 ► **Remark 2.** The previous statement is however not true when using the blocker technique  
 120 as in MINISAT [4] or CADICAL [6] where the blockers are stored in the watch lists. A  
 121 blocked clause does not need to be dereferenced. In that regard, comparing linked lists and  
 122 arrays in NAPSAT should yield better results than comparing linked lists and arrays in  
 123 other solvers. Since linked lists perform worse than arrays in NAPSAT (Section 4), we can  
 124 expect that the difference would be even more significant in other solvers.

125 Second, it is more memory efficient. When elements are moved from one watch list to  
 126 another, the array-based representation might require to reallocate memory in the destination  
 127 list, and some memory might be wasted in the origin list. With linked lists, the memory is  
 128 allocated during the creation of the clause, and the only lost memory is the pointers at the  
 129 end of the list (which are not larger than the field remembering the size of the array). A  
 130 simple and not very rigorous evaluation with Valgrind [8] on 10 problems with 150 variables  
 131 from the uniform random 3-SAT satisfiable problems of SATLIB [5] (`uf150-01.cnf` to  
 132 `uf150-010.cnf`) showed us that the linked list scheme uses about 13% less memory than the

■ **Table 1** Intuitive comparison of the array and linked list representation of watch lists.

Aspect	Array		Linked list	
Dereference level	2 levels	(−)	1 level	(+)
Memory usage	Extensible	(−)	Fixed	(+)
Insertion	$O(1)$ or $O(n)$	(−)	$O(1)$	(+)
Arbitrary removal	$O(n)$	(=)	$O(n)$	(=)
Removal during propagation	$O(1)$	(=)	$O(1)$	(=)
Bookkeeping overhead	Low	(+)	High	(−)
Code complexity	Low	(+)	Medium	(−)

array-based scheme. We do not desire to make any strong claim about this, nor spend more time on this evaluation, but it is an interesting observation. Furthermore, the additional cost of adding two 32 bits integers to the clause data structure is negligible. Trying to artificially double the size of the clauses (32 bytes) using the `alignas` keyword in C++, without any benefit, had only a limited impact on the runtime of the solver.

► **Remark 3.** The PICO SAT paper [1] observes a similar save in memory for the linked list representation.

Finally, singly linked lists can easily be extended to doubly linked lists that allow  $O(1)$  removal of any clause in the list. This could be interesting in a context where clauses are often removed by the user.

**Implementation complications.** The main difficulty of implementing a linked list scheme, aside from the necessary bookkeeping, is simplifying the clause set. Indeed, the array-based approach allows having a clause in more than two watch lists at a time, and then cleaning it up in post-processing. This is not possible with the linked list representation. In particular, NAPSAT supports different backtracking strategies that maintain different invariants on the watched literals [3]. This makes the implementation of the linked list representation more complex.

Table 1 summarizes the comparison of the observations made in this section. We denote by (+) (resp. (−)) when a feature benefits (resp. harms) the respective representation, and (=) when the feature is equivalent in both representations.


## 4 Empirical Evaluation

Figure 4 shows the difference of runtime between both implementations<sup>4</sup> of watch lists in NAPSAT. Our experiments were conducted on the random 3-SAT instances of the SATLIB [5] with 250 variables. While the measurements were conducted on a single core of an Intel Core i7-10750H of a laptop with 16GB of DDR4 RAM, the results are significant enough to assert that they are not due to noise and that the linked list representation is significantly slower than the array-based representation: measured at 2.21 for SAT instances and 2.91 for UNSAT instances.

Note that the entropy in the results is also because changing the representation modifies the behavior of the solver. In particular, the watch lists change order differently in both implementations. When using linked lists, the removal of an element in the list is stable, that is, the order of the list does not change. However, when removing an element in the array

<sup>4</sup> commits 0f300b57 and df5a9ca4 for linked lists and array respectively

## 1:6 To link or not to link?



(a) SAT instances. The linked list representation runs on average 2.21 times longer than the array implementation. (b) UNSAT instances. The linked list representation runs on average 2.91 times longer than the array implementation.

■ **Figure 4** Comparison of the runtime of the array and linked list representation of watch lists on the random 3-SAT instances of the SATLIB with 250 variables. All runs were performed with default options of NAPSAT (non-chronological backtracking with clause deletion).

165 implementation, the last element of the list is moved to the position of the removed element.  
166 Furthermore, when inserting a clause in the watch list, the array implementation pushes it  
167 to the end of the list, while the linked list implementation pushes it to the beginning of the  
168 list. This has an impact on which conflicts are found by the solver. However, this effect is  
169 quasi-random, and we can observe similar trends on all sizes of the SATLIB.

170 The **Linked list** and **Array** lines of Table 2 also show that the linked list scheme does  
171 not scale well. The runtime ratio increases with the size of the problems.

## 172 5 Practical Difficulties of Using Linked Lists

173 A seemingly insignificant drawback of the linked list scheme is that it is not possible to  
174 explore the watch list without dereferencing the pointer to the literals of the clause. Since  
175 the solver needs to know which branch to choose, it needs to know the order of the two  
176 watched literals. To do so, either a copy must be kept inside of the clause data structure,  
177 further increasing the bookkeeping overhead, or the literals must be dereferenced. This  
178 largely negates the advantage of the blocker. Indeed, when the blocker is satisfied, we wish  
179 to avoid checking the literals. However, in our linked list representation, this is not possible.

180 ► **Remark 4.** In the PICOSAT paper [1], it was suggested to use a bit in the link to store  
181 this information. For example, a link to the clause  $C$  in a watch list of literal  $\ell$  would be  
182 marked with a bit set to 0 if  $\ell$  is the first watched literal of  $C$ , and 1 otherwise. This would  
183 allow us to avoid dereferencing the literals. However, swapping literals in the clause is now  
184 a problem since we would need to update the other link as well.

185 To test this claim, we have ensured that the literals were artificially dereferenced in  
186 NAPSAT before checking the blocker. The results are shown in Table 2. We can see that  
187 the cost of dereferencing the literals is significant and is responsible for a nonnegligible part

■ **Table 2** Comparison of the average runtime of the different watch list representations on the uniform random 3-SAT instances of the SATLIB.

	uf200	uuf200	uf225	uuf225	uf250	uuf250
<b>Linked list</b>	0.28 s	0.75 s	1.78 s	5.10 s	15.00 s	52.20 s
<b>Array</b>	0.20 s	0.44 s	1.10 s	2.63 s	6.80 s	17.92 s
<b>Array with dereference</b>	0.17 s	0.46 s	1.16 s	2.92 s	8.52 s	24.52 s

■ **Table 3** Final comparison of the array and linked list representation of watch lists. Empirically, the array representation is faster and easier to implement.

Aspect	Array	Linked list
Dereference level	2 levels (−)	1 level (+)
Memory usage	Extensible (−)	Fixed (+)
Bookkeeping overhead	Low (+)	High (−)
Code complexity	Low (+)	Medium (−)
Usefulness of blockers	High (+)	Low (−)

188 of the slowdown. However, there seems to be more to it. The rest of the performance  
 189 drop is probably due to the increased bookkeeping overhead. For example, swapping the  
 190 watched literals is a common practice in SAT solvers, and can be done in propagation with  
 191 4 assembly instructions using xor operations without branching. However, in the linked list  
 192 representation, we cannot simply swap the literals, we also need to conditionally update  
 193 the clause references. The branching and the additional memory accesses slow down the  
 194 propagation. This leads us to believe that copying the literals to avoid dereferencing them  
 195 might not be a good idea. For further investigation, we would check the impact of the order  
 196 of clauses in the lists on the runtime of the solver.

197 ► **Remark 5.** As opposed to this work, PICOSAT [1] results are more promising. It might be  
 198 because PICOSAT does not use blockers, or because the linked list representation is more  
 199 efficient in PICOSAT than in NAPSAT. It would be interesting to investigate this further.

## 200 6 Conclusion

201 This paper studied the benefits and challenges of implementing watch lists using arrays and  
 202 linked lists. Table 3 summarizes the comparison between both representations of watch  
 203 lists. While linked lists have some merits, our experiments show that they do not pay off  
 204 in practice. The blocker is such a powerful tool that saving a bit of memory is not worth  
 205 negating its benefits. We therefore do not recommend using linked lists for general-purpose  
 206 SAT solvers.

## 207 References

- 208 1 Armin Biere. Picosat essentials. *J. Satisf. Boolean Model. Comput.*, 4(2-4):75–97, 2008.  
 209 URL: <https://doi.org/10.3233/sat190039>, doi:10.3233/SAT190039.
- 210 2 Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximillian Heisinger. CaDiCaL, Kissat,  
 211 Paracooba, Plingeling and Treengeling entering the SAT Competition 2020. In *Proc. of SAT  
 212 Competition 2020 – Solver and Benchmark Descriptions*, volume B-2020-1 of *Department of  
 213 Computer Science Report Series B*, pages 51–53. University of Helsinki, 2020.
- 214 3 Robin Coutelier, Mathias Fleury, and Laura Kovács. Lazy reimplication in chronological  
 215 backtracking. In *To Appear in the proceeding of SAT, 2024*.

## 1:8 To link or not to link?

- 216 **4** Niklas Eén and Niklas Sörensson. An extensible sat-solver. In *SAT*, volume 2919  
217 of *Lecture Notes in Computer Science*, pages 502–518. Springer, 2003. doi:10.1007/  
218 978-3-540-24605-3\\_37.
- 219 **5** Holger H Hoos and Thomas Stützle. SATLIB: An Online Resource for Research on SAT.  
220 *Sat*, 2000:283–292, 2000.
- 221 **6** Norbert Manthey. Radical modification–watch sat. *SAT COMPETITION 2021*, page 28,  
222 2021.
- 223 **7** Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik.  
224 Chaff: Engineering an Efficient SAT Solver. In *DAC*, pages 530–535. ACM, 2001. doi:  
225 10.1145/378239.379017.
- 226 **8** Nicholas Nethercote and Julian Seward. Valgrind: a framework for heavyweight dynamic  
227 binary instrumentation. In *Proceedings of the ACM SIGPLAN 2007 Conference on Program-*  
228 *ming Language Design and Implementation, San Diego, California, USA, June 10-13, 2007*,  
229 pages 89–100. ACM, 2007. doi:10.1145/1250734.1250746.
- 230 **9** João P. Marques Silva and Karem A. Sakallah. GRASP - a new search algorithm for  
231 satisfiability. In *ICCAD*, pages 220–227. IEEE Computer Society / ACM, 1996. doi:  
232 10.1109/ICCAD.1996.569607.