

# Project Report

## Parallel DPLL SAT Solver

Project Name: BoolSAT

---

---

Name: Kenneth Kiprotich  
UNI: kk3302  
Date: 12/21/2021

---

---

## 1 Introduction

The Boolean Satisfiability Problem involves determining whether variables of a given Boolean formula can be replaced using True or False in a way that satisfies the Boolean formula, i.e. leads the Boolean formula to evaluate to True. If the Boolean evaluates to true, it is SAT (otherwise, UNSAT). SAT is NP-Complete, making it one of the interesting problems in theoretical computer science.

For this project, I sought to explore existing SAT-solving algorithms and find ways to parallelize them in Haskell, then compare the performance of sequential approaches against the parallel one. There are several algorithms which use different heuristics. I mainly focused on the DPLL algorithm.

## 2 DPLL Algorithm

The DPLL (Davis-Putnam-Logemann-Loveland) Algorithm makes use of backtracking search, unit propagation, literal elimination, and recursive checks on different variable assignments to determine boolean satisfiability.

### 2.1 Unit Propagation

This is a step to simplify the provided CNF-SAT problem at every iteration. It involves searching for a unit clause and (1) removing instances of other clauses that contain the literal in the unit clause and (2) removing instances of the literal's complement from other clauses in the formula. This works since the problem is in Conjunctive Normal Form, so if a literal value is assigned to True, all the clauses that contain that literal value will automatically evaluate to True, and the literal's complement adds no meaning to the formula. Unit propagation reduces the search space significantly.

### 2.2 Literal Elimination

This is the removal of literals that do not contain any complements in the set of clauses provided. Such literals are usually called pure literals. They are eliminated from all the clauses since they can be assigned to True without having effect on the eventual satisfiability result of the formula.

## 2.3 Determining SAT or UNSAT

The DPLL Algorithm runs an exhaustive search on all possible assignments of variables, using unit propagation and literal elimination (discussed above) to reduce search space on different iterations. While doing the exhaustive search, the algorithm gives partial assignments to variables and determines whether taking that path will lead to satisfiability, backtracking if it does not.

A CNF formula is determined to be SAT if the variables are assigned without a clause in the CNF formula becoming empty. A CNF formula is determined to be UNSAT if the opposite happens, that is, if the assignment of variables leads to one of the clauses being empty.

## 2.4 DPLL Pseudocode

This is the DPLL pseudocode, adapted from [here](#), with steps for unit-propagation and literal elimination as discussed above:

---

**Algorithm 1** DPLL Pseudocode

---

```
procedure DPLL: INPUT = CNF FORMULA, OUTPUT = TRUE OF FALSE
  If Formula is a consistent set of literals, return True
  If Formula is a contains and empty clause, return False
  For every unit clause u in the Formula:
    Formula  $\leftarrow$  unit-propagate w.r.t. u
  For every pure literal p in the Formula:
    Formula  $\leftarrow$  simplify w.r.t. p
  u  $\leftarrow$  choose a literal from Formula
  return DPLL (Formula AND u) or DPLL (Formula AND NOT u)
```

---

## 3 Implementation of DPLL in Haskell

### 3.1 CNF Files

A CNF file represents literals as numeric values. Negative numerals represent complements, for instance the complement of 1 is  $-1$ . Each line in the file represents one clause, and each of the lines end with a 0, which is a special value and is not treated as a literal. As expected, since it is in Conjunctive Normal Form, literals in the lines are OR-ed, while clauses across the file are AND-ed.

CNF files for testing the sequential and parallel DPLL are easily available online. Haskell has a [DIMACS CNF parser library](#) that converts the CNF formulas into a 2-dimensional Unboxed Array. I processed the clauses into a list of lists, i.e. a list of clauses.

### 3.2 Sequential Implementation

There are several sequential solvers easily available, including such as which uses a [MiniSAT Solver](#) available on Hackage. However, for the purpose of this project, I focused mainly on parallelizing a sequential approach that closely follows the DPLL characteristics discussed in Section 2 above.

### 3.3 Parallel Implementation

By studying the DPLL algorithm, we learn that the sequential approach builds up a tree-like structure for the solution space. I think that the last step is the one is especially important: a literal is chosen from the remaining clauses, and the algorithm branches into two children, with each branch assigned to a different polarity of the literal. Therefore, this part of the DPLL algorithm lends itself neatly to parallelization.

Initially, I had tried introducing parallelism to some of the subroutines in DPLL, such as unit propagation, literal elimination, or simplification of clauses. However, I concluded that these steps were not very interesting since they would not significantly narrow down the time it takes to cover the whole solution space. Unit Propagation and literal elimination, for instance, could have been parallelized using a `ParList` since I parse the boolean formulas into list of lists. However, I think that parallelizing tree structure was more interesting since it would significantly affect the speed at which the solution space is explored.

Below is the Haskell pseudocode for the parallelized section of the DPLL algorithm:

```
1 State = {[clauses] , [variable assignments]}
2 parDPLL: depth State
3   | if formula is empty = return [assignments]
4   | otherwise = do
5     lit <- chooseLiteral formula
6     let positive = parDPLL (i-1) (State simplified w.r.t lit)
7         (vars with lit added)
8     let negative = parDPLL (i-1) (State simplified w.r.t -lit)
9         (vars with -lit added)
10
11   if depth > 0 then
12     runEval $ do
13       x <- rpar falseBranch
14       return (case trueBranch of
15         Nothing -> x
16         Just result -> return result)
17   else
18     case trueBranch of
19       Nothing -> falseBranch
20       Just result -> return result
21
22   where
23     formula = updated clauses after unitpropagation s'
24     vars    = updated list of variables assignments
```

`State` holds the formula and a list of variable assignments. The formula is the CNF-SAT problem presented as a list of lists, where the nested lists are clauses. The formula is created by parsing a `.cnf` file provided as a command-line argument.

`True` and `False` are assigned to a literal obtained by `chooseLiteral`, and branched into two paths: `trueBranch` and `falseBranch`. This is the step that continues in parallel. The `rpar` in this case does not need an accompanying `rseq` since it evaluates fully by the time the opposite branch returns i.e. we have to check whether `positivePath` returns a `result` or `Nothing` first. Here is the `rpar` section of the algorithm which is an Eval computation performed by `runEval`:

```
1 x <- rpar falseBranch
2   return (case trueBranch of
3     Nothing -> x
```

```
4 Just result -> return result)
```

The results of the parallel algorithm against the sequential one are analyzed in more details in the next section.

## 4 Evaluation

### 4.1 Datasets used for testing

I mostly worked with problems that have been determined to be unsatisfiable in order to ensure that the algorithms explore the maximum search space rather than race to find the solution. Therefore, the solution space for Unsatisfiable boolean formulas provides a better heuristic for measuring the performance of parallelism.

The files are downloadable from the [SATLIB Benchmark Problems site](#) .

### 4.2 Results

Below is the threadscope diagram for a parallel DPLL running on 2 cores.

I used several cnf files to test the sequential and parallel algorithms, but the results below are from one Unsatisfiable formula. The sequential algorithm took an average of 36.549 seconds over 5 runs.

**These are the results from the parallel algorithm  
Depth = 8, with -N2 on a 4-thread i5 computer:**

| Parallel Algorithm Data: depth = 8 |              |                  |              |                |              |                            |
|------------------------------------|--------------|------------------|--------------|----------------|--------------|----------------------------|
| Depth                              | Total Sparks | Converted Sparks | GC'ed Sparks | Fizzled Sparks | Elapsed Time | Speedup on Sequential Time |
| 1                                  | 1            | 1                | 0            | 0              | 18.13s       | 1.96                       |
| 2                                  | 3            | 1                | 0            | 2              | 18.77s       | 1.97                       |
| 4                                  | 3            | 1                | 0            | 2              | 18.57s       | 1.98                       |
| 8                                  | 255          | 1                | 0            | 254            | 18.20s       | 1.96                       |
| 16                                 | 65535        | 2                | 122          | 65411          | 18.06s       | 1.96                       |
| 32                                 | 4194355      | 6                | 3989850      | 204499         | 18.19s       | 1.96                       |

**Depth = 8, with -N4 on a 4-thread i5 computer:**

| Parallel Algorithm Data: depth = 8 |              |                  |              |                |              |                            |
|------------------------------------|--------------|------------------|--------------|----------------|--------------|----------------------------|
| Depth                              | Total Sparks | Converted Sparks | GC'ed Sparks | Fizzled Sparks | Elapsed Time | Speedup on Sequential Time |
| 1                                  | 1            | 1                | 0            | 0              | 22.15        | 1.65                       |
| 2                                  | 3            | 3                | 0            | 0              | 16.88s       | 2.16                       |
| 4                                  | 15           | 5                | 0            | 10             | 16.80s       | 2.17                       |
| 8                                  | 255          | 5                | 0            | 250            | 16.87s       | 2.17                       |
| 16                                 | 65535        | 18               | 126          | 65391          | 17.29s       | 2.11                       |
| 32                                 | 4194539      | 73               | 3982895      | 211571         | 17.03s       | 2.15                       |

The results above show that parallelism on a Dual-COre Intel Core i5 Macbook. Using -N2 and -N4 arguments did not make much of a difference in runtime when compared to the sequential program. The comparison ratios between sequential runtime and parallel runtime (on both -N2 and -N4) were consistent on all the CNF Formulas.

Figure 1 below shows the eventlog diagrams obtained using threadscope using -N2 argument, with a parallelism depth of 8. The workloads on the two cores look balanced and consistent, ending at the same time. Evenlog diagrams with similar arguments (depth and number of cores) were the same as Figure 1, therefore load balancing was consistent across all the sample problems I tried.

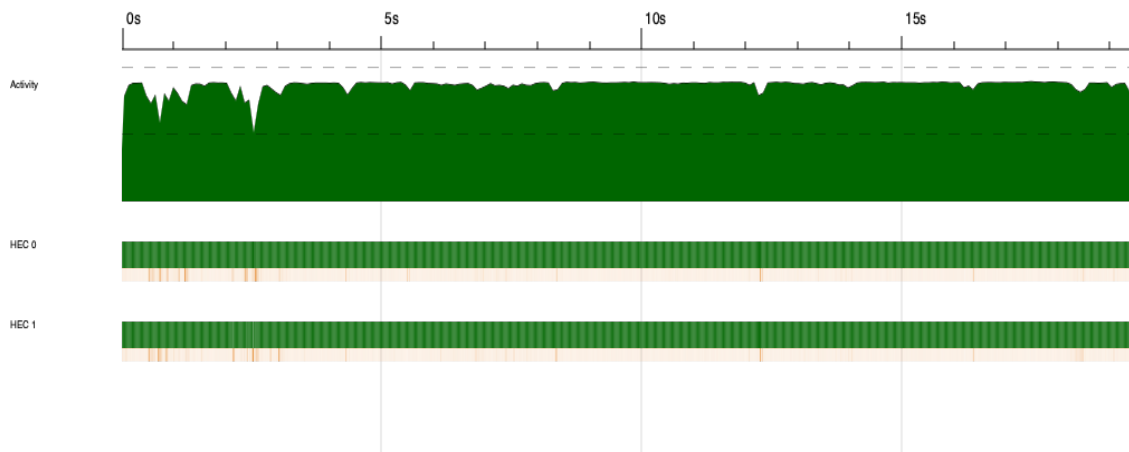


Figure 1: Eventlog for parallel DPLL: depth = 8, -N2

Figure 2 below shows the eventlog diagrams with -N4 argument and depth 8.

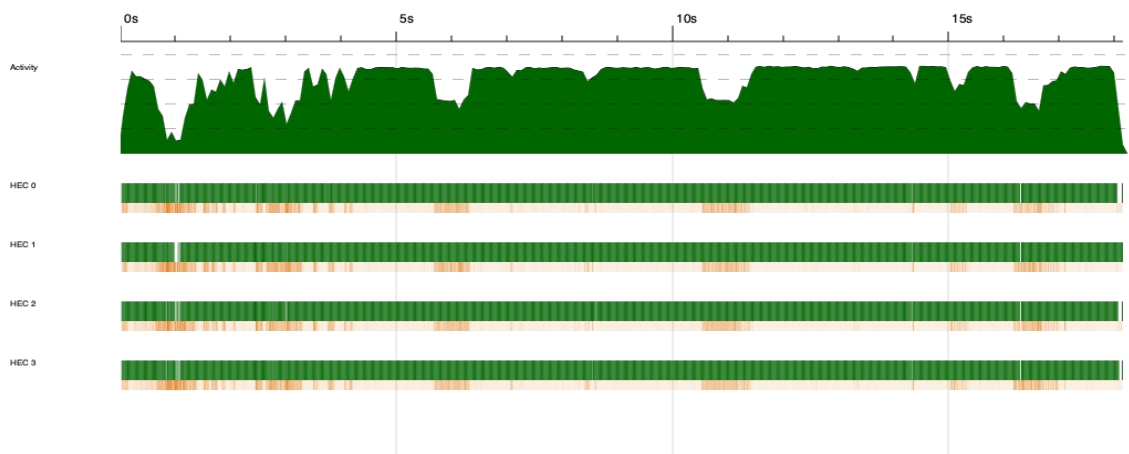


Figure 2: Eventlog for parallel DPLL: depth = 8, -N4

## 5 Lessons Learned

### 5.1 Parallelized DPLL is Faster

The parallel version of the algorithm was consistently about 2 times faster than the sequential version on all the CNF Formulas that I tested on. This is demonstrated by the table under section 4.2. The table below shows the sequential and parallel runtimes on several datasets from the SATLIB Benchmark Problems. I used parallelism depth of 8 and 2 cores (-N2) for the parallel algorithm. The Formulas with suffix *pret* are from Graph-Coloring Problems, and the ones with suffix *hole* are Pigeon-Hole Problems.

| Comparison Between Parallel and Sequential Runtimes |                    |                         |
|---|--------------------|-------------------------|
| CNF Formula   | Sequential Runtime | Parallel Time (depth 8) |
| pret60_25.cnf                                       | 38.574s            | 17.797                  |
| pret60_40.cnf                                       | 40.591s            | 16.884s                 |
| pret60_60.cnf                                       | 39.680s            | 16.85s                  |
| hole7.cnf   | 4.144s             | 1.988s                  |
| hole8.cnf   | 106.588s           | 50.30s                  |

### 5.2 Unsatisfiable CNF Problems Cause A Lot of Fizzling

There was a significant difference on the ratio of sparks converted versus sparks fizzled between Satisfiable CNF Formulas and Unsatisfiable CNF Formulas. As seen in section 4.2, testing on Unsatisfiable CNF Problems result in a lot of sparks getting fizzled especially at depth 8 and beyond. On the other hand, Satisfiable problems consistently resulted in more Converted Sparks than Fizzled Sparks up to depth 10.

According to Marlow (in the class text book), fizzled sparks are caused when the “expression was unevaluated at the time it was sparked but was later evaluated independently by the program.” This may have occurred more in Unsatisfiable CNF Problems since parallel search on the solution space may have resulted in a branch realizing that it has finished doing the work that is still sparked for a separate branch, leading to the other work getting fizzled.

I think this is a part of the parallel algorithm that I can improve in the future.

### 5.3 More Cores Would Likely Increase Speedup Even More

Increasing the depth did not result in a reciprocal increase in speedup that I had expected. I ran the programs on a Dual-Core Intel Core i5 computer, which does not offer processor power to take significant advantage of Haskell parallelism’s potential. I think that more cores would result in incremental speedup (up to some point) as the depth is increased.

My attempts to run on a virtual machine with more power were not successful (yet).

## 6 Code Listing

### 6.1 Lib.hs

```
1 module Lib
2   (readCNF
3   ) where
4
5 import Data.Array.Unboxed
6 import Language.CNF.Parse.ParseDIMACS as DIMACS
7
8 readCNF :: FilePath -> IO (Either [Char] [[Integer]])
9 readCNF filename = do
10   cnfFile <- DIMACS.parseFile filename
11   case cnfFile of
12     Left _ -> return $ Left "Error loading CNF file"
13     Right cnfUarray -> return $ Right $ createFormula cnfUarray
14
15 createFormula :: CNF -> [[Integer]]
16 createFormula cnfUarray = map (map (toInteger) . elems) $ clauses
    cnfUarray
```

### 6.2 BoolSat.hs

```
1 {-# LANGUAGE BangPatterns #-}
2 {--
3
4 ref:
5 https://stackoverflow.com/questions/12547160/how-does-the-dpll-algorithm-work
6 https://en.wikipedia.org/wiki/DPLL\_algorithm
7 https://gist.github.com/adrianlshaw/1807739
8 https://www.cs.cmu.edu/~15414/f17/lectures/10-dpll.pdf
9 --}
10
11 module BoolSat where
12
13 import Data.Maybe
14 import Control.Parallel.Strategies
15
16 type Literal = Integer
17 type Clause = [Literal]
18 type Formula = [Clause]
19 type Record = [Literal]
20
21 {-
22 SolverState holds the formula and the record. The formula is the CNF-SAT
23 problem which is modified as different literals are assigned. The record
24 holds the assignments that are 'true'
25 -}
26 data SolverState = SolverState {formula :: !Formula
27   , record :: !Record
28   } deriving (Show)
29
30 {-Sequential DPLL algorithm that uses unitpropagation and backtracking-}
31 seqDpll :: SolverState -> Maybe Record
32 seqDpll s
33   | null cnf = return rec
34   | otherwise = do
```

```

35     l <- chooseLiteral cnf
36     case seqDpll (SolverState (eliminateLiteral cnf l) (l:rec)) of
37       Just res -> return res
38       Nothing -> seqDpll $ SolverState (eliminateLiteral cnf (-l)) ((-l):
39         rec)
39     where
40       state' = unitpropagate s
41       cnf = formula state'
42       rec = record state'
43
44 {-
45   [DEPRECATED]: first try: parallelize dpll.
46 -}
47 trial_dpllPar :: (Ord a, Num a) => a -> SolverState -> Maybe Record
48 trial_dpllPar _ (SolverState [] rec) = Just rec
49 trial_dpllPar i s
50   | null cnf = return rec
51   | otherwise = do
52     case getUnit cnf of
53       Just u -> trial_dpllPar i $ SolverState (eliminateLiteral cnf u) (u:
54         rec)
55       Nothing ->
56         let
57           dlit = unwrapMaybe $ chooseLiteral cnf
58           trueBranch = trial_dpllPar (i-1)
59             (SolverState (eliminateLiteral cnf dlit) (dlit:rec)
60           )
61           falseBranch = trial_dpllPar (i-1)
62             (SolverState (eliminateLiteral cnf (-dlit)) ((-dlit
63           ):rec))
64         in if i > 0 then
65           runEval $ do
66             x <- rpar falseBranch
67             return (case trueBranch of
68               Nothing -> x
69               Just r -> return r)
70           else
71             case trueBranch of
72               Nothing -> falseBranch
73               Just r -> return r
74     where
75       state' = unitpropagate s
76       cnf = formula state'
77       rec = record state'
78
79 {-
80   Parallel DPLL that does not have any depth. It parallelizes the dpll on
81   the branching part
82   i.e. the last line of the dpll pseudocode. Computes the falseBranch in
83   par with the true one.
84 -}
85 parDpll :: SolverState -> Maybe Record
86 parDpll s
87   | null cnf = return rec
88   | otherwise = do
89     l <- chooseLiteral cnf
90     let trueBranch = parDpll (SolverState (eliminateLiteral cnf l) (l:
91       rec))
92         falseBranch = parDpll (SolverState (eliminateLiteral cnf (-l))
93       ((-l):rec))
94     in runEval $ do

```



```

88     x <- rpar falseBranch
89     return (case trueBranch of
90         Nothing -> x
91         Just r -> return r)
92
93     where
94         state' = unitpropagate s
95         cnf = formula state'
96         rec = record state'
97
98 -----
99 {-
100 Parallel DPLL that has a depth param. It parallelizes the dpll on the
101 branching part
102 i.e. the last line of the dpll pseudocode. Computes the falseBranch in
103 par with the true one.
104 -}
105 parDpll3 :: (Ord t, Num t) => t -> SolverState -> Maybe Record
106 parDpll3 i s
107 | null cnf = return rec
108 | otherwise = do
109     l <- chooseLiteral cnf
110     let trueBranch = parDpll3 (i-1) (SolverState (eliminateLiteral cnf l
111 ) (l:rec))
112         falseBranch = parDpll3 (i-1) (SolverState (eliminateLiteral cnf
113 (-l)) ((-l):rec))
114         if i > 0 then
115             runEval $ do
116                 x <- rpar falseBranch
117                 return (case trueBranch of
118                     Nothing -> x
119                     Just r -> return r)
120         else
121             case trueBranch of
122                 Nothing -> falseBranch
123                 Just r -> return r
124
125     where
126         state' = unitpropagate s
127         cnf = formula state'
128         rec = record state'
129 -----
130 {-More info on unit propagation in the project report-}
131 unitpropagate :: SolverState -> SolverState
132 unitpropagate (SolverState cnf rec) =
133     case getUnit cnf of
134         Nothing -> SolverState cnf rec
135         Just u -> unitpropagate $ SolverState (eliminateLiteral cnf u) (u:rec)
136
137 {-Checks for a literal in the formula and returns it (or Nothing otherwise
138 )-}
139 chooseLiteral :: Formula -> Maybe Literal
140 chooseLiteral cnf = listToMaybe . concat $ cnf
141
142 {-
143 Checks for a clause with only one literal and returns the literal, other
144 wise it returns Nothing
145 -}
146 getUnit :: Formula -> Maybe Literal

```

```

143 getUnit xs = listToMaybe [x | [x] <- xs]
144
145 {-
146   This does what was described under the section of unit propagation in
147   the
148   report. That is, given a literal a to be simplified with, it: " (1)
149   removes
150   instances of other clauses that contain the literal and (2) removes
151   instances of the literal's complement in other clauses.
152   Proof: (a OR _) = a, (a OR (NOT a)) = a
153   -}
154 eliminateLiteral :: Formula -> Literal -> Formula
155 eliminateLiteral cnf l = [simplClause x l | x <- cnf, not (elem l x)]
156   where
157     simplClause c lit = filter (/= -lit) c
158
159 {-solver for sequential; takes cnf formula as input, together with empty
160   list
161   that will hold the output-}
162 seqDpllSolve :: [[Integer]] -> Maybe [Integer]
163 seqDpllSolve = seqDpll . flip SolverState []
164
165 {-solvers for parallel; takes cnf formula as input, together with empty
166   list
167   that will hold the output-}
168 parDpllTrialSolve :: (Ord a, Num a) => a -> [[Integer]] -> Maybe [Integer]
169 parDpllTrialSolve i = trial_dp11Par i . flip SolverState []
170
171 parDpllSolve :: [[Integer]] -> Maybe [Integer]
172 parDpllSolve = parDpll . flip SolverState []
173
174 parDpllSolve3 :: (Ord a, Num a) => a -> [[Integer]] -> Maybe [Integer]
175 parDpllSolve3 i = parDpll3 i . flip SolverState []
176
177 {-obtaining the value in the Maybe type-}
178 unwrapMaybe :: Maybe a -> a
179 unwrapMaybe (Just n) = n
180 unwrapMaybe Nothing = error $ "Nothing is returned here"

```

## 6.3 Main.hs

```

1 module Main where
2
3 import Lib
4 import BoolSat
5 import Control.Monad
6 import System.Environment(getArgs)
7 import System.Exit(die)
8
9 {-
10   Main is set up for 3 types of tests:
11   --}
12
13 -- >>> 1: tests a parallel dp11 that has no max
14 --   parallelization depth
15
16 -- main = do
17 --   args <- getArgs

```

```

18 -- case args of
19 --   [filename] -> do
20 --     cnfFormula <- readCNF filename
21 --     case cnfFormula of
22 --       Left e -> putStrLn e
23 --       Right formula -> putStrLn $
24 --         show $
25 --         parDpllSolve formula
26 --     _ -> putStrLn "Usage: stack run <cnf file>"
27
28
29 -- >>> 2: tests a parallel dpll that has a max
30 --     parallelization depth
31
32 main = do
33   args <- getArgs
34   case args of
35     [filename, depth] -> do
36       cnfFormula <- readCNF filename
37       let d = read $ depth
38       case cnfFormula of
39         Left e -> putStrLn e
40         Right formula -> putStrLn $
41           show $
42           parDpllSolve3 d formula
43       _ -> putStrLn "Usage: stack run <cnf file> <depth of parallelism>"
44
45 -- >>> 3: tests a sequential dpll
46
47 -- main = do
48 --   args <- getArgs
49 --   case args of
50 --     [filename] -> do
51 --       cnfFormula <- readCNF filename
52 --       case cnfFormula of
53 --         Left e -> putStrLn e
54 --         Right formula -> putStrLn $
55 --           show $
56 --           seqDpllSolve formula
57 --       _ -> putStrLn "Usage: ./sequential <cnf file>"

```

## 7 References

Parallel and Concurrent Programming in Haskell by Simon Marlow

Haskell Dimacs CNF parser library on Hackage

DPLL explanation here: <http://www.diag.uniroma1.it/~liberato/ar/dpll/dpll.html>