

CSCE 420 - Spring 2023

Programming Assignment P3

due: **Tues, Apr 25, 2023, 5:00pm**

Overview

The goal of this project is to write a C++ or Python program to do logical inference using **DPLL** (Davis-Putman procedure, i.e. Boolean satisfiability solver). You will use DPLL to answer some questions about the **Australia map-coloring** problem, to generate solutions to the **N-queens problem**, and for **Sammy's Sport Shop**. You will also implement the **unit-clause heuristic**, and evaluate its effect on the efficiency of finding a model.

Implementing DPLL

Command-line interface for the program:

```
usage: DPLL <filename> <literal>* [+UCH]
or:    python DPLL.py <filename> <literal>* [+UCH]
```

After reading in your KB (list of clauses), you will call a function like: `DPLL(clauses, model)`, with `model` initially set to the empty truth assignment over all propositional symbols appearing in the KB. A model can be represented by a hash table that maps propositions (strings) to truth values (integers: 1=true, -1=false, 0=unknown). You may also include extra literals (strings) on the command line, which should be appended to the KB (as unit clauses). See an example of this with the map-color problem below. The optional flag (+UCH) is for turning on the *unit-clause heuristic*, which is turned OFF by default.

`DPLL()` will be implemented as a recursive function, as shown in Fig 7.17 in the textbook. It will determine whether a KB is satisfiable. If a model that is complete and consistent can be found, the print out the final model (truth assignment that satisfies all the clauses); otherwise, print 'unsatisfiable'.

An important part of making DPLL efficient is implementing the **unit-clause heuristic** (UCH), where you give preference to resolving pairs of clauses where one is a unit clause. Remember that you have to re-check clauses as the search progresses, because clauses with more than one literal can become unit-clauses, if all but one of the literals are made false by the model.) Make this heuristic an option on the command line, so you can test whether turning them on increases the efficiency of finding solutions (reduces the number of DPLL calls needed to find a model). By default, your `DPLL()` function should NOT use this heuristic. But the +UCH flag allows the user to selectively turn on this heuristic.

When the program first starts, print out the command line (to save in the output file). Whenever you assign a variable in the model, **you should print out tracing information** that indicates whether it is a **choice-point** (e.g. last 2 lines of the pseudocode), or whether you **forcing** the assignment of a truth value due to the unit-clause heuristic. You should also print out a message whenever **back-tracking** occurs (second line of the pseudocode).

At the end, when you print out the result ('unsatisfiable', or the model), you should also **print out the number of times DPLL() was called** (you can use a simple global variable as a counter). This is effectively equivalent to the number of nodes in the Search Tree (state space) that are searched. See the example transcript at the end of this handout.

Unit Clause Heuristic

While back-tracking makes a big improvement in the efficiency of DPLL (by avoiding large parts of the search space containing models that won't satisfy the clauses), the Unit Clause Heuristic (UCH) can improve the efficiency of finding a satisfying model even further (necessary for scaling DPLL up to larger applications with thousands of variables and clauses).

In the simple version of the DPLL algorithm, with each call, the “next” unassigned proposition (in order) is chosen and tested by assigning it first to T and then F (a choice-point). Instead, UCH scans the list of clauses for a clause that is not satisfied by the current model, and for which all of its literals are made false by the model except one, which is unknown (unassigned) in the model. If such a clause can be found, then DPLL can just bind the truth value of that proposition to its sign as a literal in the clause. This circumvents the choice-point; there is no need to try binding the prop to T (and recursively calling DPLL to see the model can be completed) or F (i.e. “guessing”); the existence of a unit clause determines which truth value can be used.

For example, suppose we have a set of clauses:

1. $a \vee \neg b \vee c \vee \neg d$
2. $d \vee \neg c \vee \neg e$
3. $c \vee \neg f$

Next, consider the model $\{a=F, b=F, c=?, d=F, e=T, f=?\}$. Clause 1 is irrelevant (UCH skips it) because it is already satisfied by $\neg b$. However, clause 2 is effectively considered to be a unit clause, given this model, because d and $\neg e$ are falsified by the model, and $\neg c$ is the only remaining literals which value is unknown in the model. Clause 3 is not unit because it has 2 literals with unknown truth values. Therefore, in this situation, UCH would directly extend the model by adding $c=F$ (since it is a negative literal in the clause), without having to go through the choice-point. Think of writing down the satisfaction of the literals *given the model* as “satisfied” (S) “unsatisfied” (U) or “unknown” (?):

- | | |
|---------------------------------------|------------|
| 1. $a \vee \neg b \vee c \vee \neg d$ | U, S, ?, S |
| 2. $d \vee \neg c \vee \neg e$ | U, ?, U |
| 3. $c \vee \neg f$ | ?, ? |

What makes clause 2 a unit clause is that it has no satisfied literals (S) and exactly 1 unknown literal (U). There is only one possible truth value remaining for that proposition (c) that would make the model satisfy the clause: in this case, c clearly has to be assigned F.

CNF File Format

The input files for DPLL will be KBs consisting of propositional clauses written in the following CNF format. Since a clause is a disjunction of literals (propositions or negated props), we can write a clause as a list of symbols, ignoring the 'or' symbols, and prefixing negated propositions with a minus '-'. For example, *apple v -lime v cherry* would be written as "apple -lime cherry". Symbols can also contain digits and other characters like '_'; let's say it is case-sensitive. Each clause in the CNF file is written on a separate line; additional white-space doesn't matter. Blank lines can be skipped. Comments on any line can start with a '#' and everything after the '#' is ignored. Here is an example of a KB written in CNF format:

```
# clauses (these come from doing implication elimination on rules)
# room2 is accessible only if doors 1 and 2 are unlocked
-door1unlocked -door2unlocked room2accessible
weekend door1unlocked      # (not weekend) -> door1unlocked
-broken_door2 door2unlocked # broken_door2 -> door2unlocked
-wednesday -weekend      # wednesday -> (not weekend)
# facts
wednesday
broken_door2
```

Problems to Solve with DPLL

Australia Map Coloring (example in textbook, Ch. 6)

- write the KB file for map-color in CNF
- use DPLL to generate a model (*please give the answers to all these questions in your write-up, RESULTS.txt, and along with the transcripts of outputs*)
- create a modified KB file to generate a second model, where you force Queensland to be a different color (e.g. add QG to the KB, if QR is true in the first model)
- create a modified KB file by adding another fact that will force an inconsistency. For example, if QG is true in the second model, then it turns out the Victoria will have all have to be green (and thus can't be blue). To make the KB unsatisfiable, add both QG and VB as facts, and use DPLL to show this is *unsatisfiable*.
- In each case, report how many DPLL calls are made on each problem **with versus without** the unit-clause heuristic.

Sammy's Sport Shop

Save the knowledge base you created in Homework #2 and converted to CNF as a .cnf file (i.e. clauses in the ASCII format as described above). The KB should only have the general rules in it, but nothing specific to any particular scenario (i.e. do not include any facts).

Use your DPLL program to generate the solution to the scenario described in Homework #2. Note that you can add the corresponding facts (labels and observations) to the command-line.

Scenario A (same as in HW#2):

- facts: O1Y, O2W, O3Y, L1W, L2Y, L3B
- command: DPLL sammy.cnf O1Y O2W O3Y L1W L2Y L3B
- expected result: check that C2W is in the final model

Now try solving a different version of the problem, where a yellow ball is drawn from box 2 and white from boxes 1 and 3, and box 1 is labeled 'white' (incorrectly) and boxes 2 and 3 are labeled 'both'.

Scenario B (new):

- facts: O1W O2Y O3W L1W L2B L3B
- command: DPLL sammy.cnf O1W L2B L3B O2Y O3W L1W
- result: are the contents of the boxes different than in scenario A?

| | | | |
|---------------------|-------|--------|-------|
| observed: | white | yellow | white |
| boxes: | 1 | 2 | 3 |
| labels (incorrect): | white | both | both |

N-queens

Use your DPLL program to solve the **6-queens problem**. As described in the textbook, the N-queens problem refers to how to place N queens on an NxN chess board such that no queens can attack each other (either vertically, horizontally, or diagonally). Your goal is to encode this as a Boolean satisfiability problem and then use DPLL to generate a model.

To represent the problem, we can use propositional symbols like “Qcr” where c is the column and r is the row. For example, for 4-queens, the layout looks like this:

```
Q11 Q21 Q31 Q41
Q12 Q22 Q32 Q42
Q13 Q23 Q33 Q43
Q14 Q24 Q34 Q44
```

You will need to write a CNF KB with all the clauses necessary for the N-queens problem (where N is an input argument), including clauses *precluding* two queens being in the same row, column, and diagonal. For example, to say the queens in column 1 and 2 cannot both be in row 4, you could say “-Q14 -Q24”. Don’t forget to include sentences saying there has to be at least 1 queen in each column and/or row. There will be a separate KB for 3-queens, 4-queens, 5-queens, and 6-queens. Your KBs will have lots of clauses. Many of the clauses are repetitive and follow a pattern, so you might want to write a *script* to generate them, for each version of the problem (N=3..6).

A transcript of my solution to the 4-queens problem is shown at the end of this handout.

(note: there is no solution for 3-queens, so if you generate the KB and DPLL, it should fail to find a model and say ‘unsatisfiable’)

Your goal is to find a solution to the 4-queens, 5-queens, and 6-queens problem (using the unit-clause heuristic), and show that **3-queens** is *unsatisfiable*.

What to Turn In

Create a subdirectory in your (private) Github project for the class (on github.com) called 'P3'. In P3/, you should have at least the following files:

- **README** – should explain how your program works (e.g. command-line arguments), and any important parameters or constraints
- **makefile** – we must be able to compile your C++ code on compute.cs.tamu.edu by simply calling 'make', or be able to run it using python3
- **DPLL.cpp or DPLL.py**: contains your implementation of DPLL, conforming to the command-line usage above
- **KB files**: .cnf files with clauses written in same format described above
 - **mapcolor.cnf**
 - **sammy.cnf**
 - **3queens.cnf, 4queens.cnf, 5queens.cnf, and 6queens.cnf**
- **transcripts**: – show the outputs
 - **results_mapcolor.txt, results_mapcolor2.txt** (adding a fact on the command line to change the model)
 - **results_mapcolor.UCH.txt, results_mapcolor2.UCH.txt**
 - **results_sammy_scenarioA.txt, results_sammy_scenarioB.txt**
 - **results_3queens.txt, results_4queens.txt, results_5queens.txt, results_6queens.txt**
 - **results_3queens.UCH.txt, results_4queens.UCH.txt, results_5queens.UCH.txt, results_6queens.UCH.txt**
- **RESULTS.txt** –
 - a text file with a table summarizing the results of all your runs:
 - command line inputs
 - model output (just the positive literals) (or 'unsatisfiable')
 - number of DPLL calls without UCH
 - number of DPLL calls with UCH

The date you commit your files and push them to the Github server will be used to determine when it is turned in and whether any late penalties will be applied.

Grading

The materials you turn in will be graded according to the following criteria:

- 20% - does it compile and run without problems on compute.cs.tamu.edu?
- 20% - does the implementation look correct? (DPLL)
- 20% - do the knowledge base files look correct? (clauses)
- 20% - does it run correctly on test cases?
- 20% - do the RESULTS look correct? (transcripts, tracing info, #DPLL calls with and without unit-clause heuristic)

Examples

```

> DPLL mapcolor.cnf ...OR... > python DPLL.py mapcolor.cnf
command: python DPLL.py mapcolor.cnf
model: {'NSWR': 0, 'QR': 0, 'TR': 0, 'VG': 0, 'WAG': 0, 'SAB': 0, 'NTR': 0,
'SAG': 0, 'VB': 0, 'WAB': 0, 'VR': 0, 'TG': 0, 'NSWB': 0, 'QB': 0, 'QG': 0,
'NSWG': 0, 'SAR': 0, 'NTB': 0, 'TB': 0, 'WAR': 0, 'NTG': 0}
trying NSWB=T
model: {'VB': 0, 'VG': 0, 'WAG': 0, 'SAG': 0, 'WAB': 0, 'VR': 0, 'SAR': 0,
'WAR': 0, 'NSWR': 0, 'QR': 0, 'SAB': 0, 'NTR': 0, 'TR': 0, 'NSWB': 1, 'QB':
0, 'QG': 0, 'NSWG': 0, 'TG': 0, 'NTB': 0, 'TB': 0, 'NTG': 0}
trying NSWG=T
model: {'VB': 0, 'VG': 0, 'WAG': 0, 'TR': 0, 'WAB': 0, 'VR': 0, 'TG': 0,
'WAR': 0, 'NSWR': 0, 'QR': 0, 'SAB': 0, 'NTR': 0, 'SAG': 0, 'NSWB': 1, 'QB':
0, 'QG': 0, 'NSWG': 1, 'SAR': 0, 'NTB': 0, 'TB': 0, 'NTG': 0}
backtracking
trying NSWG=F
model: {'VB': 0, 'VG': 0, 'WAG': 0, 'TR': 0, 'WAB': 0, 'VR': 0, 'TG': 0,
'WAR': 0, 'NSWR': 0, 'QR': 0, 'SAB': 0, 'NTR': 0, 'SAG': 0, 'NSWB': 1, 'QB':
0, 'QG': 0, 'NSWG': -1, 'SAR': 0, 'NTB': 0, 'TB': 0, 'NTG': 0}
...
solution:
NSWB: 1
NSWG: -1
NSWR: -1
NTB: 1
NTG: -1
NTR: -1
QB: -1
QG: 1
QR: -1
SAB: -1
SAG: -1
SAR: 1
TB: 1
TG: -1
TR: -1
VB: -1
VG: 1
VR: -1
WAB: -1
WAG: 1
WAR: -1
just the Satisfied (true) propositions:
VG WAG SAR NSWB QG NTB TB
total DPLL calls: 36
UCH=False

```

Re-run it with Unit Clause Heuristic to see if DPLL calls decreases:

```
> DPLL mapcolor.cnf +UCH
...
just the Satisfied (true) propositions:
VG WAG SAR NSWB QG NTB TB
total DPLL calls: 22
UCH=True
```

Since Victoria is green in the model above, re-run with -VG to get a different solution:

```
> DPLL mapcolor.cnf -VG
...
just the Satisfied (true) propositions:
SAG VR WAR QR NSWB NTB TB
total DPLL calls: 74
UCH=False
```

> python Nqueens.py 4 > 4queens.cnf // generate clauses

```
> DPLL 4queens.cnf +UCH
command: DPLL 4queens.cnf +UCH
model: {'Q33': 0, 'Q44': 0, 'Q42': 0, 'Q43': 0, 'Q21': 0, 'Q34': 0, 'Q31': 0,
'Q32': 0, 'Q14': 0, 'Q22': 0, 'Q23': 0, 'Q11': 0, 'Q13': 0, 'Q12': 0, 'Q24':
0, 'Q41': 0}
trying Q11=T
model: {'Q12': 0, 'Q44': 0, 'Q42': 0, 'Q43': 0, 'Q32': 0, 'Q41': 0, 'Q22': 0,
'Q33': 0, 'Q21': 0, 'Q31': 0, 'Q23': 0, 'Q11': 1, 'Q13': 0, 'Q34': 0, 'Q24':
0, 'Q14': 0}
forcing Q12=-1 by UCH
model: {'Q34': 0, 'Q44': 0, 'Q42': 0, 'Q43': 0, 'Q41': 0, 'Q21': 0, 'Q31': 0,
'Q33': 0, 'Q32': 0, 'Q22': 0, 'Q23': 0, 'Q11': 1, 'Q13': 0, 'Q12': -1, 'Q24':
0, 'Q14': 0}
...
solution (model):
Q11: -1
Q12: 1
Q13: -1
Q14: -1
Q21: -1
Q22: -1
Q23: -1
Q24: 1
Q31: 1
Q32: -1
Q33: -1
Q34: -1
Q41: -1
Q42: -1
Q43: 1
```



```

Q44: -1
just the Satisfied (true) propositions:
Q43 Q31 Q12 Q24
total DPLL calls: 34
UCH=True

```

If we were to visualize this solution, it would like this:

```

. Q . .
. . . Q
Q . . .
. . Q .

```

Note, there are also other solutions. If you add “-Q43” on the command-line, it will show you a different solution with no queen in this position:

```

> DPLL 4queens.cnf -Q43 +UCH
...
just the Satisfied (true) propositions:
Q34 Q42 Q21 Q13
total DPLL calls: 44
UCH=True, PSH=False

```

If we were to visualize this alternative solution, it would like this, which is symmetric to the solution above:

```

. . Q .
Q . . .
. . . Q
. Q . .

```