

Homework 1
CSE4051 Program Verification, Fall 2025
Woosuk Lee
due: 10/7(Wed), 23:59

- You must use Python 3 for this assignment.
- Submit one python file per problem via the submission system in the course website. Make sure that your files run without errors.
- You are not allowed to use any external libraries except for Z3 for Python.
- All problems can be solved by encoding them as SAT or MaxSAT problems (reasoning over values other than Boolean values (e.g., integers) is not allowed).

Exercise 1 DIMACS format is a standard format for specifying SAT problems. A DIMACS file consists of a preamble line followed by lines representing clauses. The preamble line starts with the letter ‘p’, followed by the problem type (usually ‘cnf’ for conjunctive normal form), the number of variables, and the number of clauses.

Each clause is represented on a separate line, with positive integers representing variables and negative integers representing negated variables.

Each clause ends with a ‘0’. For example, the clause $x_1 \vee \neg x_2 \vee x_3$ would be represented as "1 -2 3 0".

Here is an example of a DIMACS file representing a simple SAT problem:

Example input:

```
p cnf 3 2
1 -2 0
-1 2 3 0
```

This file specifies a SAT problem with 3 variables (namely, x_1 , x_2 , and x_3) and 2 clauses. The first clause is $x_1 \vee \neg x_2$ and the second clause is $\neg x_1 \vee x_2 \vee x_3$.

Write a python program that reads a DIMACS file and outputs the number of all possible assignments of variables that satisfy all the clauses in the file. You can assume that the input file is named “input.dimacs” and is located in the same directory as your program.

Example output: 5

Here is the starter code:

```

1 from z3 import *
2 def read_dimacs(file_path):
3     with open(file_path, 'r') as f:
4         lines = f.readlines()
5
6     # fill here
7
8     return num_vars, clauses
9
10 def count_satisfying_assignments(num_vars, clauses):
11     # fill here
12     return count
13
14 if __name__ == "__main__":
15     num_vars, clauses = read_dimacs("input.dimacs")
16     result = count_satisfying_assignments(num_vars, clauses)
17     print(result)

```

Exercise 2 There are n switches controlling a single light. The light is on if (1) an odd number of switches are on, (2) two adjacent switches are not both on, and (3) the first or the last switch is on. Write a program that takes n as input and outputs the number of all possible assignments of the switches that turn on the light. The input as a single integer n will be given via standard input.

HINT: You may use Xor operator in Z3 to express the parity constraint (1).

Example input:

5

Example output: 3

The valid switch configurations are 10000, 00001, and 10101, where 1 means on and 0 means off.

Here is the starter code:

```

1 from z3 import *
2 def count_light_on_configurations(n):
3     # fill here
4     return count
5
6 if __name__ == "__main__":
7     n = int(input().strip())
8     result = count_light_on_configurations(n)
9     print(result)

```

Exercise 3 Recall the seat assignment problem discussed in the class. In the class, we considered the case where three people and three seats are given. Now, consider the general case where n people and n seats are given. There are three kinds of constraints: (1) person i cannot sit in seat j , (2) person i must sit in seat j , and (3) person i and person j cannot sit next to each other. The indices of people and seats are 0-based (i.e., from 0 to $n - 1$). Write a program that takes n and a set of constraints as input and outputs the number of all possible assignments of people to seats that satisfy all the constraints. The input is of the following format:

```
[the number of people/seats n]
1 i j (person i cannot sit in seat j)
...
2 i j (person i must sit in seat j)
...
3 i j (person i and person j cannot sit next to each other)
...
```

For example, the following input specifies the case where there are 4 people and 4 seats, person 0 cannot sit in seat 1, person 1 must sit in seat 2, and person 3 cannot sit next to each other.

Example input:

```
4
1 0 1
2 1 2
3 2 3
```

Example output: 2

The two valid orders of people from left to right are (0, 2, 1, 3) and (0, 3, 1, 2).

You can assume that the input will be given as a file named "input.txt" in the same directory as your program.
Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     with open(file_path, 'r') as f:
4         lines = f.readlines()
5         # fill here
6         return n, constraints
7
8 def count_valid_assignments(n, constraints):
9     # fill here
10    return count
11
12 if __name__ == "__main__":
13     n, constraints = read_input("input.txt")
14     result = count_valid_assignments(n, constraints)
15     print(result)
```

Exercise 4 Suppose you are organizing a tennis league with n matches and $2n$ players, where each player has a skill level represented by an integer. You want to schedule matches between players such that the difference in skill levels between any two players in a match is less than or equal to a given threshold k . Write a program that takes as input the number of matches n , the skill levels of the $2n$ players, and the threshold k , and outputs the number of all possible ways to schedule the matches satisfying the skill level constraint. The input will be given in a file named "input.txt" in the following format:

```
[n]
[skill_level_1] [skill_level_2] ... [skill_level_2n]
[k]
```

For example, the following input specifies the case where there are 2 matches (4 players), the skill levels of the players are 1, 3, 5, and 7, and the threshold is 2.

Example input:

```
2
1 3 5 7
2
```

Example output: 1

The only valid pairing is skill 1 with skill 3, and skill 5 with skill 7. Re-ordering the matches or swapping the two players in a match does not create a different schedule.
Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     with open(file_path, 'r') as f:
4         lines = f.readlines()
5         #fill here
6         return n, skill_levels, k
7
8 def count_match_schedules(n, skill_levels, k):
9     # fill here
10    return count
11
12 if __name__ == "__main__":
13     n, skill_levels, k = read_input("input.txt")
14     result = count_match_schedules(n, skill_levels, k)
15     print(result)
```

Exercise 5 Given a directed graph, a starting node, an ending node, and a set of constraints on the nodes, write a program that finds the shortest path from the starting node to the ending node that satisfies all the constraints. The constraints can include: (1) certain nodes must be included in the path, (2) certain nodes must be excluded from the path, (3) certain node must be visited if another node is visited, and (4) certain node cannot be visited if another node is visited.

The indices of nodes are 0-based (i.e., from 0 to n-1). The starting node is 0 and the ending node is n-1. You may assume that there is no cycle in the graph.

The input will be given in a file named "input.txt" in the following format:

```
[n] [m] (number of nodes and edges)
[u1] [v1] (edge from u1 to v1)
...
[um] [vm] (edge from um to vm)
1 [node1] (node1 must be included)
1 [node2] (node2 must be included)
...
2 [node3] (node3 must be excluded)
2 [node4] (node4 must be excluded)
...
3 [node1] [node2] (if node1 is visited, node2 must be visited)
...
4 [node1] [node2] (if node1 is visited, node2 cannot be visited)
...
```

For example, the following input specifies a graph with 4 nodes and 4 edges, where there is an edge from node 0 to node 1, from node 1 to node 2, from node 0 to node 2, and from node 2 to node 3. The starting node is 0 and the ending node is 3. If node 0 is visited, node 1 must be visited, and if node 1 is visited, node 2 cannot be visited.

Example input:

```
4 4
0 1
1 2
0 2
2 3
3 0 1
4 1 2
```

The output should be a sequence of integers representing the nodes in the shortest path from the starting node to the ending node that satisfies all the constraints. If no such path exists, the program should output "No path".

Example output:

No path

HINT: You may encode the problem as a partial weighted MaxSAT problem.

Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     with open(file_path, 'r') as f:
4         lines = f.readlines()
5         # fill here
6     return n, edges, constraints
7
8 def find_shortest_path(n, edges, constraints):
9     start = 0
10    end = n - 1
11    # fill here
12
13 if __name__ == "__main__":
14     n, edges, constraints = read_input("input.txt")
15     result = find_shortest_path(n, edges, constraints)
16     print(result)
```

Exercise 6 You are planning your course schedule for the next semester. You are given a graph representing course prerequisites, a list of courses you have already taken. You should take at least 1 course in the next semester, and all courses are offered in the next semester. Write a program that computes the number of all possible combinations of courses you can take in the next semester while satisfying the prerequisite requirements.

The indices of courses are 0-based (i.e., from 0 to $n-1$). The input will be given in a file named "input.txt" in the following format:

```
[n] [m] (number of courses and prerequisite relations)
[u1] [v1] (course u1 is a prerequisite for course v1)
...
[um] [vm] (course um is a prerequisite for course vm)
[k] (number of courses already taken)
[c1] [c2] ... [ck] (list of courses already taken)
```

For example, the following input specifies a scenario with 4 courses and 3 prerequisite relations, where course 0 is a prerequisite for course 1, course 1 is a prerequisite for course 2, and course 0 is a prerequisite for course 3. You have already taken two courses: courses 0 and 1.

Example input:

```
4 3
0 1
1 2
0 3
2
0 1
```

Example output: 3

The possible sets of courses for the next semester are $\{2\}$, $\{3\}$, and $\{2, 3\}$.

Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     with open(file_path, 'r') as f:
4         lines = f.readlines()
5         # fill here
6     return n, edges, taken_courses
7
8 def count_course_combinations(n, edges, taken_courses):
9     # fill here
10    return count
11
12 if __name__ == "__main__":
13     n, edges, taken_courses = read_input("input.txt")
14     result = count_course_combinations(n, edges, taken_courses)
15     print(result)
```

Exercise 7 Recall the graph coloring problem discussed in the class. In the class, we considered the case where a graph is colored with three colors. In this exercise, given a graph, compute the minimum number of colors needed to color the graph such that no two adjacent nodes have the same color. The indices of nodes are 0-based (i.e., from 0 to $n-1$). The input will be given in a file named "input.txt" in the following format:

```
[n] (number of nodes)
[u1] [v1] (edge from u1 to v1)
...
[um] [vm] (edge from um to vm)
```

For example, the following input specifies a graph with 4 nodes and 4 edges, where there is an edge from node 0 to node 1, from node 1 to node 2, from node 0 to node 2, and from node 2 to node 3.

Example input:

```
4
0 1
1 2
0 2
2 3
```

Example output: 3

Nodes 0, 1, and 2 form a triangle and need three different colors; node 3 can reuse the color of node 0 or node 1.

Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     with open(file_path, 'r') as f:
4         lines = f.readlines()
5         # fill here
6         return n, edges
7
8 def min_colors(n, edges):
9     # fill here
10
11 if __name__ == "__main__":
12     n, edges = read_input("input.txt")
13     result = min_colors(n, edges)
14     print(result)
```

Exercise 8 Recall the equivalence checking of if-then-else chains discussed in the class. A Boolean decision program consists of m ordered rules and a default return value. Each rule has a guard that is a conjunction of literals and returns either 0 (false) or 1 (true). The program returns the value of the first rule whose guard is true. If no guard is true, it returns the default value.

An optimizer changes only the order of the rules. Write a program that counts all assignments to the n Boolean inputs for which the original and re-ordered programs return different values. Count assignments to all n inputs, including inputs that do not occur in any guard. Output 0 if the programs are equivalent.

Input variables are $x_1, \dots, x_n$. A positive literal i denotes x_i , and a negative literal $\neg i$ denotes $\neg x_i$. Rule indices are 0-based. Assume $n \geq 1$ and $m \geq 1$. Each guard has at least one literal. The input will be given in a file named "input.txt" in the following format:

```
[n] [m]
[k] [literal_1] ... [literal_k] [return_value] (rule 0)
... (m rule lines, in the original order)
[default_value]
[p0] [p1] ... [p(m-1)] (new order: a permutation of 0,...,m-1)
```

Example input:

```
3 2
1 1 0
1 2 1
0
1 0
```

The original program tests x_1 before x_2 ; the reordered program tests x_2 before x_1 . They disagree exactly when x_1 and x_2 are both true. The input x_3 can have either value.

Example output: 2

HINT: Compile each chain into a Boolean formula, then enumerate models in which the two formulas differ. Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     # fill here
4     return n, rules, default_value, order
5
6 def count_reordering_counterexamples(n, rules, default_value, order):
7     # rules contains (literals, return_value) pairs.
8     # fill here
9     return count
10
11 if __name__ == "__main__":
12     data = read_input("input.txt")
13     print(count_reordering_counterexamples(*data))
```

Exercise 9 Recall package management from the class. Repair an installation to satisfy new requests. Package i has an initial state $b_i \in \{0, 1\}$ (1 means installed) and a positive integer cost w_i . Installing an absent package or removing an installed package costs w_i ; keeping its state costs 0.

A dependency $(p, \{q_1, \dots, q_k\})$ requires at least one q_i to be installed whenever p is installed in the final state. All dependency rules must hold, including multiple rules for the same package. Conflicting packages cannot both be installed. Requested packages must be installed; others may be installed or removed.

Write a program that outputs the minimum total change cost, or "Impossible" if no valid final state exists. The initial state need not satisfy the rules. Package indices are 0-based. Assume $n \geq 1$, $d, c, r \geq 0$, and $k \geq 1$. The input will be given in a file named "input.txt" in the following format:

```
[n] [d] [c] [r]
[bi] [wi]           (n lines: packages 0,...,n-1)
[p] [k] [q1] ... [qk] (d dependency lines)
[u] [v]           (c conflict lines, u != v)
[requested_package] (r request lines)
```

A section with zero lines is omitted.

Example input:

```
4 1 1 1
1 4
1 1
0 2
0 3
3 2 0 2
0 3
3
```

Package 3 requires package 0 or 2, but conflicts with package 0. The cheapest repair removes package 0, installs packages 2 and 3, and keeps package 1.

Example output: 9

HINT: Use hard constraints for dependencies, conflicts, and requests. Use a weighted soft constraint for keeping each package in its initial state. Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     # dependencies contains (package, alternatives) pairs.
4     # fill here
5     return n, initial, costs, dependencies, conflicts, requested
6 def min_repair_cost(n, initial, costs, dependencies, conflicts,
7     requested):
8     # fill here
9 if __name__ == "__main__":
10     data = read_input("input.txt")
11     print(min_repair_cost(*data))
```

Exercise 10 You are testing a program with b known bugs. There are t available test cases, and each test case can detect a given set of bugs. Choose as few tests as possible so that every bug can be detected by at least one chosen test.

You are given a table of 0s and 1s. Entry $a_{ij} = 1$ means that test i can detect bug j ; $a_{ij} = 0$ means that it cannot. Write a program that outputs the minimum number of tests to choose. If some bug cannot be detected by any available test, output "Impossible".

Assume $t \geq 0$ and $b \geq 1$. Test and bug indices are 0-based. The input will be given in a file named "input.txt" in the following format:

```
[t] [b]
[a00] [a01] ... [a0(b-1)] (bugs detectable by test 0)
... (t rows of b bits each)
```

Example input:

```
3 4
1 1 0 0
0 1 1 0
0 0 1 1
```

Test 0 detects bugs 0 and 1. Test 1 detects bugs 1 and 2. Test 2 detects bugs 2 and 3. Choosing tests 0 and 2 covers all four bugs. No single test covers all bugs.

Example output: 2

HINT: Introduce one Boolean variable per test to indicate whether it is chosen. Require each bug to be covered, and use soft constraints to minimize the number of chosen tests. Here is the starter code:

```
1 from z3 import *
2 def read_input(file_path):
3     # detects is a list of t rows, each containing b bits.
4     # fill here
5     return t, b, detects
6
7 def min_detection_tests(t, b, detects):
8     # Return an integer, or "Impossible".
9     # fill here
10
11 if __name__ == "__main__":
12     data = read_input("input.txt")
13     print(min_detection_tests(*data))
```