

# Quick introduction into SAT/SMT solvers and symbolic execution

Dennis Yurichev <dennis(a)yurichev.com>

December 2015 – May 2017

## Contents

<b>1 This is a draft!</b>	<b>3</b>
<b>2 Thanks</b>	<b>3</b>
<b>3 Introduction</b>	<b>3</b>
<b>4 Is it a hype? Yet another fad?</b>	<b>3</b>
<b>5 SMT<sup>1</sup>-solvers</b>	<b>4</b>
5.1 School-level system of equations . . . . .	4
5.2 Another school-level system of equations . . . . .	5
5.3 Connection between SAT <sup>2</sup> and SMT solvers . . . . .	5
5.4 Zebra puzzle (AKA <sup>3</sup> Einstein puzzle) . . . . .	5
5.5 Sudoku puzzle . . . . .	8
5.5.1 The first idea . . . . .	9
5.5.2 The second idea . . . . .	12
5.5.3 Conclusion . . . . .	14
5.5.4 Homework . . . . .	14
5.5.5 Further reading . . . . .	14
5.5.6 Sudoku as a SAT problem . . . . .	14
5.6 Solving Problem Euler 31: “Coin sums” . . . . .	14
5.7 Using Z3 theorem prover to prove equivalence of some weird alternative to XOR operation . . . . .	15
5.7.1 In SMT-LIB form . . . . .	16
5.7.2 Using universal quantifier . . . . .	17
5.7.3 How the expression works . . . . .	17
5.8 Dietz’s formula . . . . .	17
5.9 Cracking LCG <sup>4</sup> with Z3 . . . . .	18
5.10 Solving pipe puzzle using Z3 SMT-solver . . . . .	20
5.10.1 Generation . . . . .	21
5.10.2 Solving . . . . .	22
5.11 Cracking Minesweeper with Z3 SMT solver . . . . .	25
5.11.1 The method . . . . .	25
5.11.2 The code . . . . .	26
5.12 Recalculating micro-spreadsheet using Z3Py . . . . .	29
5.12.1 Unsat core . . . . .	30
5.12.2 Stress test . . . . .	30
5.12.3 The files . . . . .	32

<sup>1</sup>Satisfiability modulo theories

<sup>2</sup>Boolean satisfiability problem

<sup>3</sup>Also Known As

<sup>4</sup>Linear congruential generator

<b>6</b>	<b>Program synthesis</b>	<b>32</b>
6.1	Synthesis of simple program using Z3 SMT-solver	32
6.1.1	Few notes	35
6.1.2	The code	35
6.2	Rockey dongle: finding unknown algorithm using only input/output pairs	35
6.2.1	Conclusion	39
6.2.2	The files	40
6.2.3	Further work	40
<b>7</b>	<b>Toy decompiler</b>	<b>40</b>
7.1	Introduction	40
7.2	Data structure	40
7.3	Simple examples	41
7.4	Dealing with compiler optimizations	46
7.4.1	Division using multiplication	51
7.5	Obfuscation/deobfuscation	52
7.6	Tests	56
7.6.1	Evaluating expressions	57
7.6.2	Using Z3 SMT-solver for testing	57
7.7	My other implementations of toy decompiler	59
7.7.1	Even simpler toy decompiler	59
7.8	Difference between toy decompiler and commercial-grade one	60
7.9	Further reading	60
7.10	The files	61
<b>8</b>	<b>Symbolic execution</b>	<b>61</b>
8.1	Symbolic computation	61
8.1.1	Rational data type	62
8.2	Symbolic execution	62
8.2.1	Swapping two values using XOR	62
8.2.2	Change endianness	63
8.2.3	Fast Fourier transform	65
8.2.4	Cyclic redundancy check	67
8.2.5	Linear congruential generator	70
8.2.6	Path constraint	71
8.2.7	Division by zero	73
8.2.8	Merge sort	73
8.2.9	Extending Expr class	75
8.2.10	Conclusion	75
8.3	Further reading	75
<b>9</b>	<b>KLEE</b>	<b>76</b>
9.1	Installation	76
9.2	School-level equation	76
9.3	Zebra puzzle	77
9.4	Sudoku	81
9.5	Unit test: HTML/CSS color	85
9.6	Unit test: strcmp() function	87
9.7	UNIX date/time	89
9.8	Inverse function for base64 decoder	93
9.9	CRC (Cyclic redundancy check)	96
9.9.1	Buffer alteration case #1	96
9.9.2	Buffer alteration case #2	97
9.9.3	Recovering input data for given CRC32 value of it	98
9.9.4	In comparison with other hashing algorithms	99
9.10	LZSS decompressor	99
9.11	strtox() from RetroBSD	102
9.12	Unit testing: simple expression evaluator (calculator)	105
9.13	Regular expressions	110
9.14	Exercise	111

<b>10 (Amateur) cryptography</b>	<b>111</b>
10.1 Serious cryptography	111
10.1.1 Attempts to break “serious” crypto	114
10.2 Amateur cryptography	114
10.2.1 Bugs	116
10.2.2 XOR ciphers	116
10.2.3 Other features	116
10.2.4 Examples	116
10.3 Case study: simple hash function	117
10.3.1 Manual decompiling	117
10.3.2 Now let’s use the Z3	120
<b>11 SAT-solvers</b>	<b>123</b>
11.1 CNF form	123
11.2 Example: 2-bit adder	124
11.2.1 MiniSat	126
11.2.2 CryptoMiniSat	127
11.3 Cracking Minesweeper with SAT solver	128
11.3.1 Simple <i>population count</i> function	128
11.3.2 Minesweeper	131
11.4 Conway’s “Game of Life”	134
11.4.1 Reversing back state of “Game of Life”	134
11.4.2 Finding “still lives”	140
11.4.3 The source code	148
<b>12 Acronyms used</b>	<b>148</b>

## 1 This is a draft!

This is very early draft, but still can be interesting for someone.

Latest version is always available at [http://yurichev.com/writings/SAT\\_SMT\\_draft-EN.pdf](http://yurichev.com/writings/SAT_SMT_draft-EN.pdf). Russian version is at [http://yurichev.com/writings/SAT\\_SMT\\_draft-RU.pdf](http://yurichev.com/writings/SAT_SMT_draft-RU.pdf).

Current text version: May 8, 2017.

For news about updates, you may subscribe my twitter<sup>5</sup>, facebook<sup>6</sup>, or github repo<sup>7</sup>.

## 2 Thanks

Leonardo de Moura and Nikolaj Bjorner, for help.

## 3 Introduction

SAT/SMT solvers can be viewed as solvers of huge systems of equations. The difference is that SMT solvers takes systems in arbitrary format, while SAT solvers are limited to boolean equations in CNF<sup>8</sup> form.

A lot of real world problems can be represented as problems of solving system of equations.

## 4 Is it a hype? Yet another fad?

Some people say, this is just another hype. No, SAT is old enough and fundamental to CS<sup>9</sup>. The reason of increased interest to it is that computers gets faster over the last couple decades, so there are attempts to solve old problems using SAT/SMT, which were inaccessible in past.

<sup>5</sup><https://twitter.com/yurichev>

<sup>6</sup><https://www.facebook.com/dennis.yurichev.5>

<sup>7</sup>[https://github.com/dennis714/SAT\\_SMT\\_article](https://github.com/dennis714/SAT_SMT_article)

<sup>8</sup>Conjunctive normal form

<sup>9</sup>Computer science

## 5 SMT-solvers

### 5.1 School-level system of equations

I've got this school-level system of equations copped from Wikipedia <sup>10</sup>:

$$\begin{aligned}3x + 2y - z &= 1 \\2x - 2y + 4z &= -2 \\-x + \frac{1}{2}y - z &= 0\end{aligned}$$

Will it be possible to solve it using Z3? Here it is:

```
#!/usr/bin/python
from z3 import *

x = Real('x')
y = Real('y')
z = Real('z')
s = Solver()
s.add(3*x + 2*y - z == 1)
s.add(2*x - 2*y + 4*z == -2)
s.add(-x + 0.5*y - z == 0)
print s.check()
print s.model()
```

We see this after run:

```
sat
[z = -2, y = -2, x = 1]
```

If we change any equation in some way so it will have no solution, `s.check()` will return “unsat”.

I've used “Real” *sort* (some kind of data type in SMT-solvers) because the last expression equals to  $\frac{1}{2}$ , which is, of course, a real number. For the integer system of equations, “Int” *sort* would work fine.

Python (and other high-level PL<sup>11</sup>s like C#) interface is highly popular, because it's practical, but in fact, there is a standard language for SMT-solvers called SMT-LIB <sup>12</sup>.

Our example rewritten to it looks like this:

```
(declare-const x Real)
(declare-const y Real)
(declare-const z Real)
(assert (= (- (+ (* 3 x) (* 2 y)) z) 1))
(assert (= (+ (- (* 2 x) (* 2 y)) (* 4 z)) -2))
(assert (= (- (+ (- 0 x) (* 0.5 y)) z) 0))
(check-sat)
(get-model)
```

This language is very close to LISP, but is somewhat hard to read for untrained eyes.

Now we run it:

```
% z3 -smt2 example.smt
sat
(model
 (define-fun z () Real
  (- 2.0))
 (define-fun y () Real
  (- 2.0))
 (define-fun x () Real
  1.0)
)
```

So when you look back to my Python code, you may feel that these 3 expressions could be executed. This is not true: Z3Py API offers overloaded operators, so expressions are constructed and passed into the guts of Z3 without any execution <sup>13</sup>. I would call it “embedded DSL<sup>14</sup>”.

<sup>10</sup>[https://en.wikipedia.org/wiki/System\\_of\\_linear\\_equations](https://en.wikipedia.org/wiki/System_of_linear_equations)

<sup>11</sup>Programming Language

<sup>12</sup><http://smtlib.cs.uiowa.edu/papers/smt-lib-reference-v2.5-r2015-06-28.pdf>

<sup>13</sup><https://github.com/Z3Prover/z3/blob/6e852762baf568af2aad1e35019fdf41189e4e12/src/api/python/z3.py>

<sup>14</sup>Domain-specific language

Same thing for Z3 C++ API, you may find there “operator+” declarations and many more <sup>15</sup>. Z3 API<sup>16</sup>s for Java, ML and .NET are also exist <sup>17</sup>.

Z3Py tutorial: <https://github.com/ericpony/z3py-tutorial>.

Z3 tutorial which uses SMT-LIB language: <http://rise4fun.com/Z3/tutorial/guide>.

## 5.2 Another school-level system of equations

I’ve found this somewhere at Facebook:

$$\begin{aligned} \text{○} + \text{○} &= 10 \\ \text{○} \times \text{□} + \text{□} &= 12 \\ \text{○} \times \text{□} - \text{△} \times \text{○} &= \text{○} \\ \text{△} &= ? \end{aligned}$$

Figure 1: System of equations

It’s that easy to solve it in Z3:

```
#!/usr/bin/python
from z3 import *

circle, square, triangle = Ints('circle square triangle')
s = Solver()
s.add(circle+circle==10)
s.add(circle*square+square==12)
s.add(circle*square-triangle*circle==circle)
print s.check()
print s.model()
```

```
sat
[triangle = 1, square = 2, circle = 5]
```

## 5.3 Connection between SAT and SMT solvers

Early SMT-solvers were frontends to SAT solvers, i.e., they translating input SMT expressions into CNF and feed SAT-solver with it. Translation process is sometimes called “bit blasting”. Some SMT-solvers still works in that way: STP uses MiniSAT or CryptoMiniSAT as backend SAT-solver. Some other SMT-solvers are more advanced (like Z3), so they use something even more complex.

## 5.4 Zebra puzzle (AKA Einstein puzzle)

Zebra puzzle is a popular puzzle, defined as follows:

<sup>15</sup><https://github.com/Z3Prover/z3/blob/6e852762baf568af2aad1e35019fdf41189e4e12/src/api/c%2B%2B/z3%2B%2B.h>

<sup>16</sup>Application programming interface

<sup>17</sup><https://github.com/Z3Prover/z3/tree/6e852762baf568af2aad1e35019fdf41189e4e12/src/api>

1. There are five houses.
2. The Englishman lives in the red house.
3. The Spaniard owns the dog.
4. Coffee is drunk in the green house.
5. The Ukrainian drinks tea.
6. The green house is immediately to the right of the ivory house.
7. The Old Gold smoker owns snails.
8. Kools are smoked in the yellow house.
9. Milk is drunk in the middle house.
10. The Norwegian lives in the first house.
11. The man who smokes Chesterfields lives in the house next to the man with the fox.
12. Kools are smoked in the house next to the house where the horse is kept.
13. The Lucky Strike smoker drinks orange juice.
14. The Japanese smokes Parliaments.
15. The Norwegian lives next to the blue house.

Now, who drinks water? Who owns the zebra?

In the interest of clarity, it must be added that each of the five houses is painted a different color, and their inhabitants are of different national extractions, own different pets, drink different beverages and smoke different brands of American cigarets [sic]. One other thing: in statement 6, right means your right.

( [https://en.wikipedia.org/wiki/Zebra\\_Puzzle](https://en.wikipedia.org/wiki/Zebra_Puzzle) )

It's a very good example of CSP<sup>18</sup>.

We would encode each entity as integer variable, representing number of house.

Then, to define that Englishman lives in red house, we will add this constraint: `Englishman == Red`, meaning that number of a house where Englishmen resides and which is painted in red is the same.

To define that Norwegian lives next to the blue house, we don't really know, if it is at left side of blue house or at right side, but we know that house numbers are different by just 1. So we will define this constraint:

`Norwegian == Blue - 1 OR Norwegian == Blue + 1`.

We will also need to limit all house numbers, so they will be in range of 1..5.

We will also use `Distinct` to show that all various entities of the same type are all has different house numbers.

```
#!/usr/bin/env python
from z3 import *

Yellow, Blue, Red, Ivory, Green=Ints('Yellow Blue Red Ivory Green')
Norwegian, Ukrainian, Englishman, Spaniard, Japanese=Ints('Norwegian Ukrainian Englishman Spaniard Japanese')
Water, Tea, Milk, OrangeJuice, Coffee=Ints('Water Tea Milk OrangeJuice Coffee')
Kools, Chesterfield, OldGold, LuckyStrike, Parliament=Ints('Kools Chesterfield OldGold LuckyStrike Parliament')
Fox, Horse, Snails, Dog, Zebra=Ints('Fox Horse Snails Dog Zebra')

s = Solver()

# colors are distinct for all 5 houses:
s.add(Distinct(Yellow, Blue, Red, Ivory, Green))

# all nationalities are living in different houses:
s.add(Distinct(Norwegian, Ukrainian, Englishman, Spaniard, Japanese))

# so are beverages:
s.add(Distinct(Water, Tea, Milk, OrangeJuice, Coffee))

# so are cigarettes:
s.add(Distinct(Kools, Chesterfield, OldGold, LuckyStrike, Parliament))

# so are pets:
s.add(Distinct(Fox, Horse, Snails, Dog, Zebra))
```

<sup>18</sup>Constraint satisfaction problem

```

# limits.
# adding two constraints at once (separated by comma) is the same
# as adding one And() constraint with two subconstraints
s.add(Yellow>=1, Yellow<=5)
s.add(Blue>=1, Blue<=5)
s.add(Red>=1, Red<=5)
s.add(Ivory>=1, Ivory<=5)
s.add(Green>=1, Green<=5)

s.add(Norwegian>=1, Norwegian<=5)
s.add(Ukrainian>=1, Ukrainian<=5)
s.add(Englishman>=1, Englishman<=5)
s.add(Spaniard>=1, Spaniard<=5)
s.add(Japanese>=1, Japanese<=5)

s.add(Water>=1, Water<=5)
s.add(Tea>=1, Tea<=5)
s.add(Milk>=1, Milk<=5)
s.add(OrangeJuice>=1, OrangeJuice<=5)
s.add(Coffee>=1, Coffee<=5)

s.add(Kools>=1, Kools<=5)
s.add(Chesterfield>=1, Chesterfield<=5)
s.add(OldGold>=1, OldGold<=5)
s.add(LuckyStrike>=1, LuckyStrike<=5)
s.add(Parliament>=1, Parliament<=5)

s.add(Fox>=1, Fox<=5)
s.add(Horse>=1, Horse<=5)
s.add(Snails>=1, Snails<=5)
s.add(Dog>=1, Dog<=5)
s.add(Zebra>=1, Zebra<=5)

# main constraints of the puzzle:

# 2.The Englishman lives in the red house.
s.add(Englishman==Red)

# 3.The Spaniard owns the dog.
s.add(Spaniard==Dog)

# 4.Coffee is drunk in the green house.
s.add(Coffee==Green)

# 5.The Ukrainian drinks tea.
s.add(Ukrainian==Tea)

# 6.The green house is immediately to the right of the ivory house.
s.add(Green==Ivory+1)

# 7.The Old Gold smoker owns snails.
s.add(OldGold==Snails)

# 8.Kools are smoked in the yellow house.
s.add(Kools==Yellow)

# 9.Milk is drunk in the middle house.
s.add(Milk==3) # i.e., 3rd house

# 10.The Norwegian lives in the first house.
s.add(Norwegian==1)

# 11.The man who smokes Chesterfields lives in the house next to the man with the fox.
s.add(Or(Chesterfield==Fox+1, Chesterfield==Fox-1)) # left or right

# 12.Kools are smoked in the house next to the house where the horse is kept.
s.add(Or(Kools==Horse+1, Kools==Horse-1)) # left or right

# 13.The Lucky Strike smoker drinks orange juice.
s.add(LuckyStrike==OrangeJuice)

# 14.The Japanese smokes Parliaments.
s.add(Japanese==Parliament)

# 15.The Norwegian lives next to the blue house.

```

```
s.add(Or(Norwegian==Blue+1, Norwegian==Blue-1)) # left or right

r=s.check()
print r
if r==unsat:
    exit(0)
m=s.model()
print(m)
```

When we run it, we got correct result:

```
sat
[Snails = 3,
Blue = 2,
Ivory = 4,
OrangeJuice = 4,
Parliament = 5,
Yellow = 1,
Fox = 1,
Zebra = 5,
Horse = 2,
Dog = 4,
Tea = 2,
Water = 1,
Chesterfield = 2,
Red = 3,
Japanese = 5,
LuckyStrike = 4,
Norwegian = 1,
Milk = 3,
Kools = 1,
OldGold = 3,
Ukrainian = 2,
Coffee = 5,
Green = 5,
Spaniard = 4,
Englishman = 3]
```

## 5.5 Sudoku puzzle

Sudoku puzzle is a 9\*9 grid with some cells filled with values, some are empty:

		5	3					
8							2	
	7			1		5		
4					5	3		
	1			7				6
		3	2				8	
	6		5					9
		4					3	
					9	7		

Unsolved Sudoku

Numbers of each row must be unique, i.e., it must contain all 9 numbers in range of 1..9 without repetition. Same story for each column and also for each 3\*3 square.

This puzzle is good candidate to try [SMT](#) solver on, because it's essentially an unsolved system of equations.



### 5.5.1 The first idea

The only thing we must decide is that how to determine in one expression, if the input 9 variables has all 9 unique numbers? They are not ordered or sorted, after all.

From the school-level arithmetics, we can devise this idea:

$$\underbrace{10^{i_1} + 10^{i_2} + \dots + 10^{i_9}}_9 = 1111111110 \quad (1)$$

Take each input variable, calculate  $10^{i_i}$  and sum them all. If all input values are unique, each will be settled at its own place. Even more than that: there will be no holes, i.e., no skipped values. So, in case of Sudoku, 1111111110 number will be final result, indicating that all 9 input values are unique, in range of 1..9.

Exponentiation is heavy operation, can we use binary operations? Yes, just replace 10 with 2:

$$\underbrace{2^{i_1} + 2^{i_2} + \dots + 2^{i_9}}_9 = 1111111110_2 \quad (2)$$

The effect is just the same, but the final value is in base 2 instead of 10.

Now a working example:

```
import sys
from z3 import *

"""
coordinates:
-----
00 01 02 | 03 04 05 | 06 07 08
10 11 12 | 13 14 15 | 16 17 18
20 21 22 | 23 24 25 | 26 27 28
-----
30 31 32 | 33 34 35 | 36 37 38
40 41 42 | 43 44 45 | 46 47 48
50 51 52 | 53 54 55 | 56 57 58
-----
60 61 62 | 63 64 65 | 66 67 68
70 71 72 | 73 74 75 | 76 77 78
80 81 82 | 83 84 85 | 86 87 88
-----
"""

s=Solver()

# using Python list comprehension, construct array of arrays of BitVec instances:
cells=[[BitVec('cell%d%d' % (r, c), 16) for c in range(9)] for r in range(9)]

# http://www.norvig.com/sudoku.html
# http://www.mirror.co.uk/news/weird-news/worlds-hardest-sudoku-can-you-242294
puzzle="..53.....8.....2..7..1.5..4.....53...1..7...6...32...8..6.5.....9..4....3.....97..."

# process text line:
current_column=0
current_row=0
for i in puzzle:
    if i!='.':
        s.add(cells[current_row][current_column]==BitVecVal(int(i),16))
        current_column=current_column+1
    if current_column==9:
        current_column=0
        current_row=current_row+1

one=BitVecVal(1,16)
mask=BitVecVal(0b111111110,16)

# for all 9 rows
for r in range(9):
    s.add(((one<<cells[r][0]) +
           (one<<cells[r][1]) +
           (one<<cells[r][2]) +
           (one<<cells[r][3]) +
           (one<<cells[r][4]) +
           (one<<cells[r][5]) +
           (one<<cells[r][6]) +
           (one<<cells[r][7]) +
```

```

        (one<<cells[r][8])==mask)
# for all 9 columns
for c in range(9):
    s.add(((one<<cells[0][c]) +
           (one<<cells[1][c]) +
           (one<<cells[2][c]) +
           (one<<cells[3][c]) +
           (one<<cells[4][c]) +
           (one<<cells[5][c]) +
           (one<<cells[6][c]) +
           (one<<cells[7][c]) +
           (one<<cells[8][c]))==mask)
# enumerate all 9 squares
for r in range(0, 9, 3):
    for c in range(0, 9, 3):
        # add constraints for each 3*3 square:
        s.add((one<<cells[r+0][c+0]) +
              (one<<cells[r+0][c+1]) +
              (one<<cells[r+0][c+2]) +
              (one<<cells[r+1][c+0]) +
              (one<<cells[r+1][c+1]) +
              (one<<cells[r+1][c+2]) +
              (one<<cells[r+2][c+0]) +
              (one<<cells[r+2][c+1]) +
              (one<<cells[r+2][c+2]))==mask)
# print s.check()
s.check()
# print s.model()
m=s.model()
for r in range(9):
    for c in range(9):
        sys.stdout.write (str(m[cells[r][c]])+" ")
    print ""

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SMT/sudoku\\_plus.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SMT/sudoku_plus.py) )

```

% time python sudoku_plus.py
1 4 5 3 2 7 6 9 8
8 3 9 6 5 4 1 2 7
6 7 2 9 1 8 5 4 3
4 9 6 1 8 5 3 7 2
2 1 8 4 7 3 9 5 6
7 5 3 2 9 6 4 8 1
3 6 7 5 4 2 8 1 9
9 8 4 7 6 1 2 3 5
5 2 1 8 3 9 7 6 4

real    0m11.717s
user    0m10.896s
sys     0m0.068s

```

Even more, we can replace summing operation to logical OR:

$$\underbrace{2^{i_1} \vee 2^{i_2} \vee \dots \vee 2^{i_9}}_9 = 1111111110_2 \quad (3)$$

```

import sys
from z3 import *

"""
coordinates:
-----
00 01 02 | 03 04 05 | 06 07 08
10 11 12 | 13 14 15 | 16 17 18
20 21 22 | 23 24 25 | 26 27 28
-----
30 31 32 | 33 34 35 | 36 37 38
40 41 42 | 43 44 45 | 46 47 48
50 51 52 | 53 54 55 | 56 57 58
-----
60 61 62 | 63 64 65 | 66 67 68

```

```

70 71 72 | 73 74 75 | 76 77 78
80 81 82 | 83 84 85 | 86 87 88
-----
"""

s=Solver()

# using Python list comprehension, construct array of arrays of BitVec instances:
cells=[[BitVec('cell%d%d' % (r, c), 16) for c in range(9)] for r in range(9)]

# http://www.norvig.com/sudoku.html
# http://www.mirror.co.uk/news/weird-news/worlds-hardest-sudoku-can-you-242294
puzzle="..53.....8.....2..7..1.5..4.....53...1..7...6..32...8..6.5.....9..4....3.....97.."

# process text line:
current_column=0
current_row=0
for i in puzzle:
    if i!='.':
        s.add(cells[current_row][current_column]==BitVecVal(int(i),16))
        current_column=current_column+1
    if current_column==9:
        current_column=0
        current_row=current_row+1

one=BitVecVal(1,16)
mask=BitVecVal(0b111111110,16)

# for all 9 rows
for r in range(9):
    s.add(((one<<cells[r][0]) |
           (one<<cells[r][1]) |
           (one<<cells[r][2]) |
           (one<<cells[r][3]) |
           (one<<cells[r][4]) |
           (one<<cells[r][5]) |
           (one<<cells[r][6]) |
           (one<<cells[r][7]) |
           (one<<cells[r][8]))==mask)

# for all 9 columns
for c in range(9):
    s.add(((one<<cells[0][c]) |
           (one<<cells[1][c]) |
           (one<<cells[2][c]) |
           (one<<cells[3][c]) |
           (one<<cells[4][c]) |
           (one<<cells[5][c]) |
           (one<<cells[6][c]) |
           (one<<cells[7][c]) |
           (one<<cells[8][c]))==mask)

# enumerate all 9 squares
for r in range(0, 9, 3):
    for c in range(0, 9, 3):
        # add constraints for each 3*3 square:
        s.add(one<<cells[r+0][c+0] |
              one<<cells[r+0][c+1] |
              one<<cells[r+0][c+2] |
              one<<cells[r+1][c+0] |
              one<<cells[r+1][c+1] |
              one<<cells[r+1][c+2] |
              one<<cells[r+2][c+0] |
              one<<cells[r+2][c+1] |
              one<<cells[r+2][c+2]==mask)

#print s.check()
s.check()
#print s.model()
m=s.model()

for r in range(9):
    for c in range(9):
        sys.stdout.write (str(m[cells[r][c]])+" ")
    print ""

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SMT/sudoku\\_or.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SMT/sudoku_or.py) )  
Now it works much faster. Z3 handles OR operation over bit vectors better than addition?

```
% time python sudoku_or.py
1 4 5 3 2 7 6 9 8
8 3 9 6 5 4 1 2 7
6 7 2 9 1 8 5 4 3
4 9 6 1 8 5 3 7 2
2 1 8 4 7 3 9 5 6
7 5 3 2 9 6 4 8 1
3 6 7 5 4 2 8 1 9
9 8 4 7 6 1 2 3 5
5 2 1 8 3 9 7 6 4

real    0m1.429s
user    0m1.393s
sys     0m0.036s
```

The puzzle I used as example is dubbed as one of the hardest known <sup>19</sup> (well, for humans). It took  $\approx 1.4$  seconds on my Intel Core i3-3110M 2.4GHz notebook to solve it.

### 5.5.2 The second idea

My first approach is far from effective, I did what first came to my mind and worked. Another approach is to use `distinct` command from SMTLIB, which tells Z3 that some variables must be distinct (or unique). This command is also available in Z3 Python interface.

I've rewritten my first Sudoku solver, now it operates over *Int sort*, it has `distinct` commands instead of bit operations, and now also other constraint added: each cell value must be in 1..9 range, because, otherwise, Z3 will offer (although correct) solution with too big and/or negative numbers.

```
import sys
from z3 import *

"""
-----
00 01 02 | 03 04 05 | 06 07 08
10 11 12 | 13 14 15 | 16 17 18
20 21 22 | 23 24 25 | 26 27 28
-----
30 31 32 | 33 34 35 | 36 37 38
40 41 42 | 43 44 45 | 46 47 48
50 51 52 | 53 54 55 | 56 57 58
-----
60 61 62 | 63 64 65 | 66 67 68
70 71 72 | 73 74 75 | 76 77 78
80 81 82 | 83 84 85 | 86 87 88
-----
"""

s=Solver()

# using Python list comprehension, construct array of arrays of BitVec instances:
cells=[[Int('cell%d%d' % (r, c)) for c in range(9)] for r in range(9)]

# http://www.norvig.com/sudoku.html
# http://www.mirror.co.uk/news/weird-news/worlds-hardest-sudoku-can-you-242294
puzzle="..53.....8.....2..7..1.5..4....53...1..7...6..32...8..6.5....9..4....3.....97..."

# process text line:
current_column=0
current_row=0
for i in puzzle:
    if i!='.':
        s.add(cells[current_row][current_column]==int(i))
        current_column=current_column+1
    if current_column==9:
        current_column=0
        current_row=current_row+1

# this is important, because otherwise, Z3 will report correct solutions with too big and/or negative numbers in
cells
```

<sup>19</sup><http://www.mirror.co.uk/news/weird-news/worlds-hardest-sudoku-can-you-242294>

```

for r in range(9):
    for c in range(9):
        s.add(cells[r][c]>=1)
        s.add(cells[r][c]<=9)

# for all 9 rows
for r in range(9):
    s.add(Distinct(cells[r][0],
                  cells[r][1],
                  cells[r][2],
                  cells[r][3],
                  cells[r][4],
                  cells[r][5],
                  cells[r][6],
                  cells[r][7],
                  cells[r][8]))

# for all 9 columns
for c in range(9):
    s.add(Distinct(cells[0][c],
                  cells[1][c],
                  cells[2][c],
                  cells[3][c],
                  cells[4][c],
                  cells[5][c],
                  cells[6][c],
                  cells[7][c],
                  cells[8][c]))

# enumerate all 9 squares
for r in range(0, 9, 3):
    for c in range(0, 9, 3):
        # add constraints for each 3*3 square:
        s.add(Distinct(cells[r+0][c+0],
                      cells[r+0][c+1],
                      cells[r+0][c+2],
                      cells[r+1][c+0],
                      cells[r+1][c+1],
                      cells[r+1][c+2],
                      cells[r+2][c+0],
                      cells[r+2][c+1],
                      cells[r+2][c+2]))

#print s.check()
s.check()
#print s.model()
m=s.model()

for r in range(9):
    for c in range(9):
        sys.stdout.write (str(m[cells[r][c]])+" ")
    print ""

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SMT/sudoku2.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SMT/sudoku2.py) )

```

% time python sudoku2.py
1 4 5 3 2 7 6 9 8
8 3 9 6 5 4 1 2 7
6 7 2 9 1 8 5 4 3
4 9 6 1 8 5 3 7 2
2 1 8 4 7 3 9 5 6
7 5 3 2 9 6 4 8 1
3 6 7 5 4 2 8 1 9
9 8 4 7 6 1 2 3 5
5 2 1 8 3 9 7 6 4

real    0m0.382s
user    0m0.346s
sys     0m0.036s

```

That's much faster.

### 5.5.3 Conclusion

SMT-solvers are so helpful, is that our Sudoku solver has nothing else, we have just defined relationships between variables (cells).

### 5.5.4 Homework

As it seems, true Sudoku puzzle is the one which has only one solution. The piece of code I've included here shows only the first one. Using the method described earlier (5.6, also called "model counting"), try to find more solutions, or prove that the solution you have just found is the only one possible.

### 5.5.5 Further reading

<http://www.norvig.com/sudoku.html>

### 5.5.6 Sudoku as a SAT problem

It's also possible to represent Sudoku puzzle as a huge CNF equation and use SAT-solver to find solution, but it's just trickier.

Some articles about it: *Building a Sudoku Solver with SAT*<sup>20</sup>, Tjark Weber, *A SAT-based Sudoku Solver*<sup>21</sup>, Ines Lynce, Joel Ouaknine, *Sudoku as a SAT Problem*<sup>22</sup>, Gihwon Kwon, Himanshu Jain, *Optimized CNF Encoding for Sudoku Puzzles*<sup>23</sup>.

SMT-solver can also use SAT-solver in its core, so it does all mundane translating work. As a "compiler", it may not do this in the most efficient way, though.

## 5.6 Solving Problem Euler 31: "Coin sums"

(This text was first published in my blog<sup>24</sup> at 10-May-2013.)

In England the currency is made up of pound, £, and pence, p, and there are eight coins in general circulation:

1p, 2p, 5p, 10p, 20p, 50p, £1 (100p) and £2 (200p). It is possible to make £2 in the following way:

1£1 + 150p + 220p + 15p + 12p + 31p How many different ways can £2 be made using any number of coins?

( Problem Euler 31 — Coin sums )

Using Z3 for solving this is overkill, and also slow, but nevertheless, it works, showing all possible solutions as well. The piece of code for blocking already found solution and search for next, and thus, counting all solutions, was taken from Stack Overflow answer<sup>25</sup>. This is also called "model counting". Constraints like "a>=0" must be present, because Z3 solver will find solutions with negative numbers.

```
#!/usr/bin/python
from z3 import *
a,b,c,d,e,f,g,h = Ints('a b c d e f g h')
s = Solver()
s.add(1*a + 2*b + 5*c + 10*d + 20*e + 50*f + 100*g + 200*h == 200,
      a>=0, b>=0, c>=0, d>=0, e>=0, f>=0, g>=0, h>=0)
result=[]
```

<sup>20</sup>[http://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-005-elements-of-software-construction-fall-assignments/MIT6\\_005F11\\_ps4.pdf](http://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-005-elements-of-software-construction-fall-assignments/MIT6_005F11_ps4.pdf)

<sup>21</sup><https://www.lri.fr/~conchon/mpri/weber.pdf>

<sup>22</sup><http://sat.inesc-id.pt/~ines/publications/aimath06.pdf>

<sup>23</sup><http://www.cs.cmu.edu/~hjain/papers/sudoku-as-SAT.pdf>

<sup>24</sup><http://dennisyurichev.blogspot.de/2013/05/in-england-currency-is-made-up-of-pound.html>

<sup>25</sup><http://stackoverflow.com/questions/11867611/z3py-checking-all-solutions-for-equation>, another question: <http://stackoverflow.com/questions/13395391/z3-finding-all-satisfying-models>

```

while True:
    if s.check() == sat:
        m = s.model()
        print m
        result.append(m)
        # Create a new constraint the blocks the current model
        block = []
        for d in m:
            # d is a declaration
            if d.arity() > 0:
                raise Z3Exception("uninterpreted functions are not supported")
            # create a constant from declaration
            c=d()
            #print c, m[d]
            if is_array(c) or c.sort().kind() == Z3_UNINTERPRETED_SORT:
                raise Z3Exception("arrays and uninterpreted sorts are not supported")
            block.append(c != m[d])
        #print "new constraint:",block
        s.add(Or(block))
    else:
        print len(result)
        break

```

Works very slow, and this is what it produces:

```

[h = 0, g = 0, f = 0, e = 0, d = 0, c = 0, b = 0, a = 200]
[f = 1, b = 5, a = 0, d = 1, g = 1, h = 0, c = 2, e = 1]
[f = 0, b = 1, a = 153, d = 0, g = 0, h = 0, c = 1, e = 2]
...
[f = 0, b = 31, a = 33, d = 2, g = 0, h = 0, c = 17, e = 0]
[f = 0, b = 30, a = 35, d = 2, g = 0, h = 0, c = 17, e = 0]
[f = 0, b = 5, a = 50, d = 2, g = 0, h = 0, c = 24, e = 0]

```

73682 results in total.

## 5.7 Using Z3 theorem prover to prove equivalence of some weird alternative to XOR operation

(The test was first published in my blog at April 2015: <http://blog.yurichev.com/node/86>).

There is a “A Hacker’s Assistant” program<sup>26</sup> (*Aha!*) written by Henry Warren, who is also the author of the great “Hacker’s Delight” book.

The *Aha!* program is essentially *superoptimizer*<sup>27</sup>, which blindly brute-force a list of some generic RISC CPU instructions to achieve shortest possible (and jumpless or branch-free) CPU code sequence for desired operation. For example, *Aha!* can find jumpless version of `abs()` function easily.

Compiler developers use superoptimization to find shortest possible (and/or jumpless) code, but I tried to do otherwise—to find longest code for some primitive operation. I tried *Aha!* to find equivalent of basic XOR operation without usage of the actual XOR instruction, and the most bizarre example *Aha!* gave is:

```

Found a 4-operation program:
add  r1,ry,rx
and  r2,ry,rx
mul  r3,r2,-2
add  r4,r3,r1
Expr: (((y & x)*-2) + (y + x))

```

And it’s hard to say, why/where we can use it, maybe for obfuscation, I’m not sure. I would call this *suboptimization* (as opposed to *superoptimization*). Or maybe *superdeoptimization*.

But my another question was also, is it possible to prove that this is correct formula at all? The *Aha!* checking some input/output values against XOR operation, but of course, not all the possible values. It is 32-bit code, so it may take very long time to try all possible 32-bit inputs to test it.

We can try Z3 theorem prover for the job. It’s called *prover*, after all.

So I wrote this:

```

#!/usr/bin/python
from z3 import *

x = BitVec('x', 32)

```

<sup>26</sup><http://www.hackersdelight.org/>

<sup>27</sup><http://en.wikipedia.org/wiki/Superoptimization>

```

y = BitVec('y', 32)
output = BitVec('output', 32)
s = Solver()
s.add(x^y==output)
s.add(((y & x)*0xFFFFFFFF) + (y + x)!=output)
print s.check()

```

In plain English language, this means “are there any case for  $x$  and  $y$  where  $x \oplus y$  doesn't equals to  $((y \& x) * -2) + (y + x)$ ?” ...and Z3 prints “unsat”, meaning, it can't find any counterexample to the equation. So this *Aha!* result is proved to be working just like XOR operation.

Oh, I also tried to extend the formula to 64 bit:

```

#!/usr/bin/python
from z3 import *

x = BitVec('x', 64)
y = BitVec('y', 64)
output = BitVec('output', 64)
s = Solver()
s.add(x^y==output)
s.add(((y & x)*0xFFFFFFFF) + (y + x)!=output)
print s.check()

```

Nope, now it says “sat”, meaning, Z3 found at least one counterexample. Oops, it's because I forgot to extend -2 number to 64-bit value:

```

#!/usr/bin/python
from z3 import *

x = BitVec('x', 64)
y = BitVec('y', 64)
output = BitVec('output', 64)
s = Solver()
s.add(x^y==output)
s.add(((y & x)*0xFFFFFFFFFFFFFFFF) + (y + x)!=output)
print s.check()

```

Now it says “unsat”, so the formula given by *Aha!* works for 64-bit code as well.

### 5.7.1 In SMT-LIB form

Now we can rephrase our expression to more suitable form:  $(x + y - ((x \& y) \ll 1))$ . It also works well in Z3Py:

```

#!/usr/bin/python
from z3 import *

x = BitVec('x', 64)
y = BitVec('y', 64)
output = BitVec('output', 64)
s = Solver()
s.add(x^y==output)
s.add((x + y - ((x & y) << 1)) != output)
print s.check()

```

Here is how to define it in SMT-LIB way:

```

(declare-const x (_ BitVec 64))
(declare-const y (_ BitVec 64))
(assert
  (not
    (=
      (bvsub
        (bvadd x y)
        (bvshl (bvand x y) (_ bv1 64)))
      (bvxor x y)
    )
  )
)
(check-sat)

```



## 5.7.2 Using universal quantifier

Z3 supports universal quantifier `exists`, which is true if at least one set of variables satisfied underlying condition:

```
(declare-const x (_ BitVec 64))
(declare-const y (_ BitVec 64))
(assert
  (exists ((x (_ BitVec 64)) (y (_ BitVec 64)))
    (not (=
      (bvsb
        (bvadd x y)
        (bvshl (bvand x y) (_ bv1 64))
      )
      (bvxor x y)
    ))
  )
)
(check-sat)
```

It returns “unsat”, meaning, Z3 couldn’t find any counterexample of the equation, i.e., it’s not exist.

This is also known as  $\exists$  in mathematical logic lingo.

Z3 also supports universal quantifier `forall`, which is true if the equation is true for all possible values. So we can rewrite our SMT-LIB example as:

```
(declare-const x (_ BitVec 64))
(declare-const y (_ BitVec 64))
(assert
  (forall ((x (_ BitVec 64)) (y (_ BitVec 64)))
    (=
      (bvsb
        (bvadd x y)
        (bvshl (bvand x y) (_ bv1 64))
      )
      (bvxor x y)
    )
  )
)
(check-sat)
```

It returns “sat”, meaning, the equation is correct for all possible 64-bit `x` and `y` values, like them all were checked.

Mathematically speaking:  $\forall n \in \mathbb{N} (x \oplus y = (x + y - ((x \& y) \ll 1)))$  <sup>28</sup>

## 5.7.3 How the expression works

First of all, binary addition can be viewed as binary XORing with carrying (11.2). Here is an example: let’s add 2 (10b) and 2 (10b). XORing these two values resulting 0, but there is a carry generated during addition of two second bits. That carry bit is propagated further and settles at the place of the 3rd bit: 100b. 4 (100b) is hence a final result of addition.

If the carry bits are not generated during addition, the addition operation is merely XORing. For example, let’s add 1 (1b) and 2 (10b).  $1 + 2$  equals to 3, but  $1 \oplus 2$  is also 3.

If the addition is XORing plus carry generation and application, we should eliminate effect of carrying somehow here. The first part of the expression  $(x + y)$  is addition, the second  $((x \& y) \ll 1)$  is just calculation of every carry bit which was used during addition. If to subtract carry bits from the result of addition, the only XOR effect is left then.

It’s hard to say how Z3 proves this: maybe it just simplifies the equation down to single XOR using simple boolean algebra rewriting rules?

## 5.8 Dietz’s formula

One of the impressive examples of *Aha!* work is finding of Dietz’s formula<sup>29</sup>, which is the code of computing average number of two numbers without overflow (which is important if you want to find average number of

<sup>28</sup>  $\forall$  means equation must be true for all possible values, which are chosen from natural numbers ( $\mathbb{N}$ ).

<sup>29</sup> <http://aggregate.org/MAGIC/#Average%20of%20Integers>

numbers like 0xFFFFFFFF0 and so on, using 32-bit registers).

Taking this in input:

```
int userfun(int x, int y) { // To find Dietz's formula for
                          // the floor-average of two
                          // unsigned integers.
    return ((unsigned long long)x + (unsigned long long)y) >> 1;
}
```

...the *Aha!* gives this:

```
Found a 4-operation program:
and  r1,ry,rx
xor  r2,ry,rx
shrs r3,r2,1
add  r4,r3,r1
Expr: ((y ^ x) >>s 1) + (y & x)
```

And it works correctly<sup>30</sup>. But how to prove it?

We will place Dietz's formula on the left side of equation and  $x + y/2$  (or  $x + y >> 1$ ) on the right side:

$$\forall n \in 0..2^{64} - 1. (x \& y) + (x \oplus y) >> 1 = x + y >> 1$$

One important thing is that we can't operate on 64-bit values on right side, because result will overflow. So we will zero extend inputs on right side by 1 bit (in other words, we will just 1 zero bit before each value). The result of Dietz's formula will also be extended by 1 bit. Hence, both sides of the equation will have a width of 65 bits:

```
(declare-const x (_ BitVec 64))
(declare-const y (_ BitVec 64))
(assert
  (forall ((x (_ BitVec 64)) (y (_ BitVec 64)))
    (=
      (( _ zero_extend 1)
       (bvadd
         (bvand x y)
         (bvlsht (bv XOR x y) (_ bv1 64))
       ))
      (bvlsht
        (bvadd (( _ zero_extend 1) x) (( _ zero_extend 1) y))
        (_ bv1 65)
      ))
    )
  )
)
(check-sat)
```

Z3 says "sat".

65 bits are enough, because the result of addition of two biggest 64-bit values has width of 65 bits:

0xFF...FF + 0xFF...FF = 0x1FF...FE.

As in previous example about XOR equivalent, `(not (= ... ))` and `exists` can also be used here instead of `forall`.

## 5.9 Cracking LCG with Z3

(This text is first appeared in my blog in June 2015 at <http://yurichev.com/blog/modulo/>.)

There are well-known weaknesses of LCG<sup>31</sup>, but let's see, if it would be possible to crack it straightforwardly, without any special knowledge. We will define all relations between LCG states in terms of Z3. Here is a test program:

<sup>30</sup>For those who interesting how it works, its mechanics is closely related to the weird XOR alternative we just saw. That's why I placed these two pieces of text one after another.

<sup>31</sup>[http://en.wikipedia.org/wiki/Linear\\_congruential\\_generator#Advantages\\_and\\_disadvantages\\_of\\_LCGs](http://en.wikipedia.org/wiki/Linear_congruential_generator#Advantages_and_disadvantages_of_LCGs), <http://www.reteam.org/papers/e59.pdf>, <http://stackoverflow.com/questions/8569113/why-1103515245-is-used-in-rand/8574774#8574774>

```

#include <stdlib.h>
#include <stdio.h>
#include <time.h>

int main()
{
    int i;

    srand(time(NULL));

    for (i=0; i<10; i++)
        printf ("%d\n", rand()%100);
};

```

It is printing 10 pseudorandom numbers in 0..99 range:

```

37
29
74
95
98
40
23
58
61
17

```

Let's say we are observing only 8 of these numbers (from 29 to 61) and we need to predict next one (17) and/or previous one (37).

The program is compiled using MSVC 2013 (I choose it because its LCG is simpler than that in Glib):

```

.text:0040112E rand      proc near
.text:0040112E          call     __getptd
.text:00401133          imul   ecx, [eax+0x14], 214013
.text:0040113A          add    ecx, 2531011
.text:00401140          mov    [eax+14h], ecx
.text:00401143          shr   ecx, 16
.text:00401146          and   ecx, 7FFFh
.text:0040114C          mov   eax, ecx
.text:0040114E          retn
.text:0040114E rand      endp

```

Let's define **LCG** in Z3Py:

```

#!/usr/bin/python
from z3 import *

output_prev = BitVec('output_prev', 32)
state1 = BitVec('state1', 32)
state2 = BitVec('state2', 32)
state3 = BitVec('state3', 32)
state4 = BitVec('state4', 32)
state5 = BitVec('state5', 32)
state6 = BitVec('state6', 32)
state7 = BitVec('state7', 32)
state8 = BitVec('state8', 32)
state9 = BitVec('state9', 32)
state10 = BitVec('state10', 32)
output_next = BitVec('output_next', 32)

s = Solver()

s.add(state2 == state1*214013+2531011)
s.add(state3 == state2*214013+2531011)
s.add(state4 == state3*214013+2531011)
s.add(state5 == state4*214013+2531011)
s.add(state6 == state5*214013+2531011)
s.add(state7 == state6*214013+2531011)
s.add(state8 == state7*214013+2531011)
s.add(state9 == state8*214013+2531011)
s.add(state10 == state9*214013+2531011)

s.add(output_prev==URem((state1>>16)&0x7FFF,100))
s.add(URem((state2>>16)&0x7FFF,100)==29)
s.add(URem((state3>>16)&0x7FFF,100)==74)

```

```

s.add(URem((state4>>16)&0x7FFF,100)==95)
s.add(URem((state5>>16)&0x7FFF,100)==98)
s.add(URem((state6>>16)&0x7FFF,100)==40)
s.add(URem((state7>>16)&0x7FFF,100)==23)
s.add(URem((state8>>16)&0x7FFF,100)==58)
s.add(URem((state9>>16)&0x7FFF,100)==61)
s.add(output_next==URem((state10>>16)&0x7FFF,100))

print(s.check())
print(s.model())

```

*URem* states for *unsigned remainder*. It works for some time and gave us correct result!

```

sat
[state3 = 2276903645,
state4 = 1467740716,
state5 = 3163191359,
state7 = 4108542129,
state8 = 2839445680,
state2 = 998088354,
state6 = 4214551046,
state1 = 1791599627,
state9 = 548002995,
output_next = 17,
output_prev = 37,
state10 = 1390515370]

```

I added  $\approx 10$  states to be sure result will be correct. It may be not in case of smaller set of information.

That is the reason why [LCG](#) is not suitable for any security-related task. This is why cryptographically secure pseudorandom number generators exist: they are designed to be protected against such simple attack. Well, at least if [NSA](#)<sup>32</sup> don't get involved <sup>33</sup>.

Security tokens like "RSA SecurID" can be viewed just as [CPRNG](#)<sup>34</sup> with a secret seed. It shows new pseudorandom number each minute, and the server can predict it, because it knows the seed. Imagine if such token would implement [LCG](#)—it would be much easier to break!

## 5.10 Solving pipe puzzle using Z3 SMT-solver

"Pipe puzzle" is a popular puzzle (just google "pipe puzzle" and look at images).

This is how shuffled puzzle looks like:




Figure 2: Shuffled puzzle

...and solved:

<sup>32</sup>National Security Agency

<sup>33</sup>[https://en.wikipedia.org/wiki/Dual\\_EC\\_DRBG](https://en.wikipedia.org/wiki/Dual_EC_DRBG)

<sup>34</sup>Cryptographically Secure Pseudorandom Number Generator





Figure 3: Solved puzzle

Let's try to find a way to solve it.

### 5.10.1 Generation

First, we need to generate it. Here is my quick idea on it. Take 8\*16 array of cells. Each cell may contain some type of block. There are joints between cells:



Blue lines are horizontal joints, red lines are vertical joints. We just set each joint to 0/false (absent) or 1/true (present), randomly.

Once set, it's now easy to find type for each cell. There are:

joints	our internal name	angle	symbol
0	type 0	0°	(space)
2	type 2a	0°	
2	type 2a	90°	-
2	type 2b	0°	┌
2	type 2b	90°	┐
2	type 2b	180°	└
2	type 2b	270°	┘

3	type 3	0°	┆
3	type 3	90°	┆
3	type 3	180°	┆
3	type 3	270°	┆
4	type 4	0°	┆

*Dangling* joints can be preset at a first stage (i.e., cell with only one joint), but they are removed recursively, these cells are transforming into empty cells. Hence, at the end, all cells has at least two joints, and the whole plumbing system has no connections with outer world—I hope this would make things clearer.


The C source code of generator is here: [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/SMT/pipe/generator](https://github.com/dennis714/SAT_SMT_article/tree/master/SMT/pipe/generator). All horizontal joints are stored in the global array *hjoints[]* and vertical in *vjoints[]*.

The C program generates ANSI-colored output like it has been showed above (5.10, 5.10) plus an array of types, with no angle information about each cell:

```
[
["0", "0", "2b", "3", "2a", "2a", "2a", "3", "3", "2a", "3", "2b", "2b", "2b", "0", "0"],
["2b", "2b", "3", "2b", "0", "0", "2b", "3", "3", "3", "3", "3", "4", "2b", "0", "0"],
["3", "4", "2b", "0", "0", "0", "3", "2b", "2b", "4", "2b", "3", "4", "2b", "2b", "2b"],
["2b", "4", "3", "2a", "3", "3", "3", "2b", "2b", "3", "3", "3", "2a", "2b", "4", "3"],
["0", "2b", "3", "2b", "3", "4", "2b", "3", "3", "2b", "3", "3", "3", "0", "2a", "2a"],
["0", "0", "2b", "2b", "0", "3", "3", "4", "3", "4", "3", "3", "3", "2b", "3", "3"],
["0", "2b", "3", "2b", "0", "3", "3", "4", "3", "4", "4", "3", "0", "3", "4", "3"],
["0", "2b", "3", "3", "2a", "3", "2b", "2b", "3", "3", "3", "3", "2a", "3", "3", "2b"],
]
```

### 5.10.2 Solving

First of all, we would think about 8\*16 array of cells, where each has four bits: “T” (top), “B” (bottom), “L” (left), “R” (right). Each bit represents half of joint.



Now we define arrays of each of four half-joints + angle information:

```
HEIGHT=8
WIDTH=16

# if T/B/R/L is Bool instead of Int, Z3 solver will work faster
T=[[Bool('cell_%d_%d_top' % (r, c)) for c in range(WIDTH)] for r in range(HEIGHT)]
B=[[Bool('cell_%d_%d_bottom' % (r, c)) for c in range(WIDTH)] for r in range(HEIGHT)]
R=[[Bool('cell_%d_%d_right' % (r, c)) for c in range(WIDTH)] for r in range(HEIGHT)]
L=[[Bool('cell_%d_%d_left' % (r, c)) for c in range(WIDTH)] for r in range(HEIGHT)]
A=[[Int('cell_%d_%d_angle' % (r, c)) for c in range(WIDTH)] for r in range(HEIGHT)]
```

We know that if each of half-joints is present, corresponding half-joint must be also present, and vice versa. We define this using these constraints:

```

# shorthand variables for True and False:
t=True
f=False

# "top" of each cell must be equal to "bottom" of the cell above
# "bottom" of each cell must be equal to "top" of the cell below
# "left" of each cell must be equal to "right" of the cell at left
# "right" of each cell must be equal to "left" of the cell at right
for r in range(HEIGHT):
    for c in range(WIDTH):
        if r!=0:
            s.add(T[r][c]==B[r-1][c])
        if r!=HEIGHT-1:
            s.add(B[r][c]==T[r+1][c])
        if c!=0:
            s.add(L[r][c]==R[r][c-1])
        if c!=WIDTH-1:
            s.add(R[r][c]==L[r][c+1])

# "left" of each cell of first column shouldn't have any connection
# so is "right" of each cell of the last column
for r in range(HEIGHT):
    s.add(L[r][0]==f)
    s.add(R[r][WIDTH-1]==f)

# "top" of each cell of the first row shouldn't have any connection
# so is "bottom" of each cell of the last row
for c in range(WIDTH):
    s.add(T[0][c]==f)
    s.add(B[HEIGHT-1][c]==f)

```

Now we'll enumerate all cells in the initial array (5.10.1). First two cells are empty there. And the third one has type "2b". This is "r" and it can be oriented in 4 possible ways. And if it has angle 0°, bottom and right half-joints are present, others are absent. If it has angle 90°, it looks like "┐", and bottom and left half-joints are present, others are absent.

In plain English: "if cell of this type has angle 0°, these half-joints must be present **OR** if it has angle 90°, these half-joints must be present, **OR**, etc, etc."

Likewise, we define all these rules for all types and all possible angles:

```

for r in range(HEIGHT):
    for c in range(WIDTH):
        ty=cells_type[r][c]

        if ty=="0":
            s.add(A[r][c]==f)
            s.add(T[r][c]==f, B[r][c]==f, L[r][c]==f, R[r][c]==f)

        if ty=="2a":
            s.add(Or(And(A[r][c]==0, L[r][c]==f, R[r][c]==f, T[r][c]==t, B[r][c]==t), # |
                    And(A[r][c]==90, L[r][c]==t, R[r][c]==t, T[r][c]==f, B[r][c]==f))) # -

        if ty=="2b":
            s.add(Or(And(A[r][c]==0, L[r][c]==f, R[r][c]==t, T[r][c]==f, B[r][c]==t), # r
                    And(A[r][c]==90, L[r][c]==t, R[r][c]==f, T[r][c]==f, B[r][c]==t), # ]
                    And(A[r][c]==180, L[r][c]==t, R[r][c]==f, T[r][c]==t, B[r][c]==f), # ]
                    And(A[r][c]==270, L[r][c]==f, R[r][c]==t, T[r][c]==t, B[r][c]==f))) # L

        if ty=="3":
            s.add(Or(And(A[r][c]==0, L[r][c]==f, R[r][c]==t, T[r][c]==t, B[r][c]==t), # T
                    And(A[r][c]==90, L[r][c]==t, R[r][c]==t, T[r][c]==f, B[r][c]==t), # T
                    And(A[r][c]==180, L[r][c]==t, R[r][c]==f, T[r][c]==t, B[r][c]==t), # T
                    And(A[r][c]==270, L[r][c]==t, R[r][c]==t, T[r][c]==t, B[r][c]==f))) # L

        if ty=="4":
            s.add(A[r][c]==0)
            s.add(T[r][c]==t, B[r][c]==t, L[r][c]==t, R[r][c]==t) # †

```

Full source code is here: [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SMT/pipe/solver/solve\\_pipe\\_puzzle1.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SMT/pipe/solver/solve_pipe_puzzle1.py).

It produces this result (prints angle for each cell and (pseudo)graphical representation):

```

sat
  0  0  0  90  90  90  90  90  90  90  90  90  0  90  0  0
  0  90  0 180  0  0  0 180  0  90 180  0  0 180  0  0
  0  0 180  0  0  0  0 180 270  0 180  0  0  90  0  90
270  0  90  90  90  90 270  90  0 270  90 180  0 270  0 180
  0 270 180  0 270  0  90  0 180  0 270 270 180  0  0  0
  0  0 270 180  0  0 270  0 270  0  90  90 270  90  0 180
  0  0  90  90  0  0  90  0  90  0  0 180  0  0  0 180
  0 270 270 270  90 270 180 270 270 270 270 270  90 270 270 180
  r T - - - T T - T r r r
r r t J      r t t T t t t J
t t J      t J L t J t t r r r
L t T - T T L r r L T t | L t t
L t r L t r t t r L L t | |
  L J      t L t L t T T L r t t
r T r      t T t T t t t t t t t
L L L - L J L L L L L - L L J

```

Figure 4: Solver script output

It worked  $\approx 4$  seconds on my old and slow Intel Atom N455 1.66GHz. Is it fast? I don't know, but again, what is really cool, we do not know about any mathematical background of all this, we just defined cells, (half-)joints and defined relations between them.

Now the next question is, how many solutions are possible? Using method described earlier (5.6), I've altered solver script<sup>35</sup> and solver said two solutions are possible.

Let's compare these two solutions using gvimdiff:

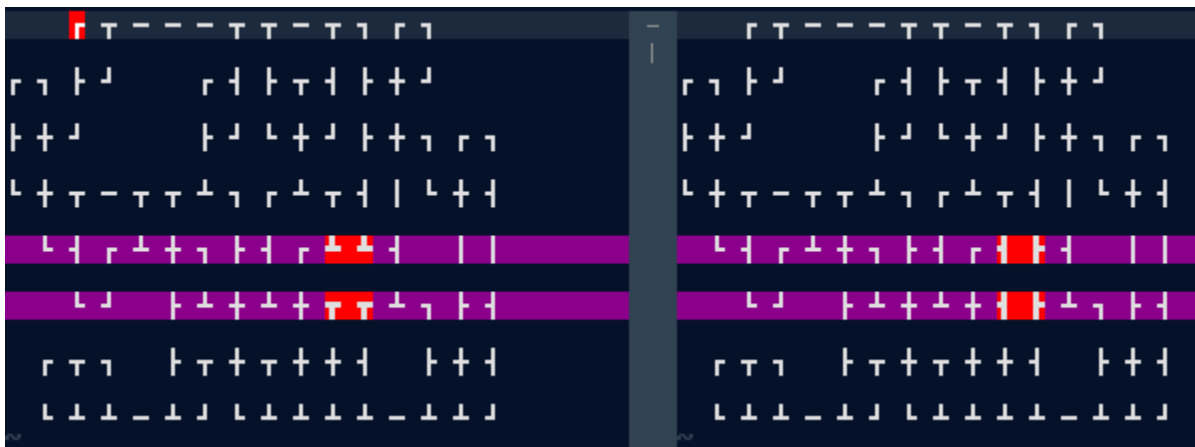


Figure 5: gvimdiff output (pardon my red cursor at left pane at left-bottom corner)

4 cells in the middle can be orientated differently. Perhaps, other puzzles may produce different results.

P.S. *Half-joint* is defined as boolean type. But in fact, the first version of the solver has been written using integer type for half-joints, and 0 was used for False and 1 for True. I did it so because I wanted to make source code tidier and narrower without using long words like "False" and "True". And it worked, but slower. Perhaps, Z3 handles boolean data types faster? Better? Anyway, I writing this to note that integer type can also be used instead of boolean, if needed.

<sup>35</sup>[https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SMT/pipe/solver/solve\\_pipe\\_puzzle2.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SMT/pipe/solver/solve_pipe_puzzle2.py)



## 5.11 Cracking Minesweeper with Z3 SMT solver

For those who are not very good at playing Minesweeper (like me), it's possible to predict bombs' placement without touching debugger.

Here is a clicked somewhere and I see revealed empty cells and cells with known number of "neighbours":



What we have here, actually? Hidden cells, empty cells (where bombs are not present), and empty cells with numbers, which shows how many bombs are placed nearby.

### 5.11.1 The method

Here is what we can do: we will try to place a bomb to all possible hidden cells and ask Z3 SMT solver, if it can disprove the very fact that the bomb can be placed there.

Take a look at this fragment. "?" mark is for hidden cell, "." is for empty cell, number is a number of neighbours.

	C1	C2	C3
R1	?	?	?
R2	?	3	.
R3	?	1	.

So there are 5 hidden cells. We will check each hidden cell by placing a bomb there. Let's first pick top/left cell:

	C1	C2	C3
R1	*	?	?
R2	?	3	.
R3	?	1	.

Then we will try to solve the following system of equations ( $RrCc$  is cell of row  $r$  and column  $c$ ):

- $R1C2+R2C1+R2C2=1$  (because we placed bomb at  $R1C1$ )
- $R2C1+R2C2+R3C1=1$  (because we have "1" at  $R3C2$ )
- $R1C1+R1C2+R1C3+R2C1+R2C2+R2C3+R3C1+R3C2+R3C3=3$  (because we have "3" at  $R2C2$ )
- $R1C2+R1C3+R2C2+R2C3+R3C2+R3C3=0$  (because we have "." at  $R2C3$ )
- $R2C2+R2C3+R3C2+R3C3=0$  (because we have "." at  $R3C3$ )

As it turns out, this system of equations is satisfiable, so there could be a bomb at this cell. But this information is not interesting to us, since we want to find cells we can freely click on. And we will try another one. And if the equation will be unsatisfiable, that would imply that a bomb cannot be there and we can click on it.

### 5.11.2 The code

```
#!/usr/bin/python

known=[
"01?10001?",
"01?100011",
"011100000",
"000000000",
"111110011",
"?????1001?",
"?????3101?",
"?????211?",
"?????????"]

from z3 import *
import sys

WIDTH=len(known[0])
HEIGHT=len(known)

print "WIDTH=", WIDTH, "HEIGHT=", HEIGHT

def chk_bomb(row, col):

    s=Solver()

    cells=[[Int('cell_r=%d_c=%d' % (r,c)) for c in range(WIDTH+2)] for r in range(HEIGHT+2)]

    # make border
    for c in range(WIDTH+2):
        s.add(cells[0][c]==0)
        s.add(cells[HEIGHT+1][c]==0)
    for r in range(HEIGHT+2):
        s.add(cells[r][0]==0)
        s.add(cells[r][WIDTH+1]==0)

    for r in range(1,HEIGHT+1):
        for c in range(1,WIDTH+1):

            t=known[r-1][c-1]
            if t in "012345678":
                s.add(cells[r][c]==0)
                # we need empty border so the following expression would be able to work for all possible cases:
                s.add(cells[r-1][c-1] + cells[r-1][c] + cells[r-1][c+1] + cells[r][c-1] + cells[r][c+1] + cells[
r+1][c-1] + cells[r+1][c] + cells[r+1][c+1]==int(t))

    # place bomb:
    s.add(cells[row][col]==1)

    result=str(s.check())
    if result=="unsat":
        print "row=%d col=%d, unsat!" % (row, col)

# enumerate all hidden cells:
for r in range(1,HEIGHT+1):
    for c in range(1,WIDTH+1):
        if known[r-1][c-1]=="?":
            chk_bomb(r, c)
```

The code is almost self-explanatory. We need border for the same reason, why Conway's "Game of Life" implementations also has border (to make calculation function simpler). Whenever we know that the cell is free of bomb, we put zero there. Whenever we know number of neighbours, we add a constraint, again, just like in "Game of Life": number of neighbours must be equal to the number we have seen in the Minesweeper. Then we place bomb somewhere and check.

Let's run:

```
row=1 col=3, unsat!
row=6 col=2, unsat!
row=6 col=3, unsat!
row=7 col=4, unsat!
row=7 col=9, unsat!
row=8 col=9, unsat!
```

These are cells where I can click safely, so I did:



Now we have more information, so we update input:

```
known=[
"01110001?",
"01?100011",
"011100000",
"000000000",
"111110011",
"?11?1001?",
"???331011",
"?????2110",
"???????10"]
```

I run it again:

```
row=7 col=1, unsat!
row=7 col=2, unsat!
row=7 col=3, unsat!
row=8 col=3, unsat!
row=9 col=5, unsat!
row=9 col=6, unsat!
```

I click on these cells again:



I update it again:

```
known=[
"01110001?",
"01?100011",
"011100000",
```

```
"000000000",
"111110011",
"?11?1001?",
"222331011",
"?22??2110",
"????22?10"]
```

```
row=8 col=2, unsat!
row=9 col=4, unsat!
```



This is last update:

```
known=[
"01110001?",
"01?100011",
"011100000",
"000000000",
"111110011",
"?11?1001?",
"222331011",
"?22??2110",
"???322?10"]
```

...last result:

```
row=9 col=1, unsat!
row=9 col=2, unsat!
```

Voila!



The source code: [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SMT/minesweeper/minesweeper\\_solver.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SMT/minesweeper/minesweeper_solver.py).

Some discussion on HN: <https://news.ycombinator.com/item?id=13797375>.  
See also: cracking Minesweeper using SAT solver: 11.3.

## 5.12 Recalculating micro-spreadsheet using Z3Py

There is a nice exercise<sup>36</sup>: write a program to recalculate micro-spreadsheet, like this one:

1	0	B0+B2	A0*B0*C0
123	10	12	11
667	A0+B1	(C1*A0)*122	A3+C2

As it turns out, though overkill, this can be solved using Z3 with little effort:

```
#!/usr/bin/python

from z3 import *
import sys, re

# MS Excel or LibreOffice style.
# first top-left cell is A0, not A1
def coord_to_name(R, C):
    return "ABCDEFGHIJKLMNPOQRSTUVWXYZ"[R]+str(C)

# open file and parse it as list of lists:
f=open(sys.argv[1],"r")
# filter(None, ...) to remove empty sublists:
ar=filter(None, [item.rstrip().split() for item in f.readlines()])
f.close()

WIDTH=len(ar[0])
HEIGHT=len(ar)

# cells{} is a dictionary with keys like "A0", "B9", etc:
cells={}
for R in range(HEIGHT):
    for C in range(WIDTH):
        name=coord_to_name(R, C)
        cells[name]=Int(name)

s=Solver()

cur_R=0
cur_C=0

for row in ar:
    for c in row:
        # string like "A0+B2" becomes "cells["A0"]+cells["B2"]":
        c=re.sub(r'([A-Z]{1}[0-9]+)', r'cells["\1"]', c)
        st="cells[\"%s\"]==%s" % (coord_to_name(cur_R, cur_C), c)
        # evaluate string. Z3Py expression is constructed at this step:
        e=eval(st)
        # add constraint:
        s.add (e)
        cur_C=cur_C+1
    cur_R=cur_R+1
    cur_C=0

result=str(s.check())
print result
if result=="sat":
    m=s.model()
    for r in range(HEIGHT):
        for c in range(WIDTH):
            sys.stdout.write (str(m[cells[coord_to_name(r, c)]])+"\t")
            sys.stdout.write ("\n")
```

( <https://github.com/dennis714/yurichev.com/blob/master/blog/spreadsheet/1.py> )

All we do is just creating pack of variables for each cell, named A0, B1, etc, of integer type. All of them are stored in *cells[]* dictionary. Key is a string. Then we parse all the strings from cells, and add to list of constraints  $A0=123$  (in case of number in cell) or  $A0=B1+C2$  (in case of expression in cell). There is a slight preparation: string like  $A0+B2$  becomes  $cells["A0"]+cells["B2"]$ .

<sup>36</sup>Blog post in Russian: <http://thesz.livejournal.com/280784.html>

Then the string is evaluated using Python `eval()` method, which is highly dangerous<sup>37</sup>: imagine if end-user could add a string to cell other than expression? Nevertheless, it serves our purposes well, because this is a simplest way to pass a string with expression into Z3.

Z3 do the job with little effort:

```
% python 1.py test1
sat
1      0      135      82041
123    10     12       11
667    11     1342     83383
```

### 5.12.1 Unsat core

Now the problem: what if there is circular dependency? Like:

```
1      0      B0+B2      A0*B0
123    10     12       11
C1+123 C0*123 A0*122     A3+C2
```

Two first cells of the last row (C0 and C1) are linked to each other. Our program will just tells “unsat”, meaning, it couldn’t satisfy all constraints together. We can’t use this as error message reported to end-user, because it’s highly unfriendly.

However, we can fetch *unsat core*, i.e., list of variables which Z3 finds conflicting.

```
...
s=Solver()
s.set(unsat_core=True)
...
# add constraint:
s.assert_and_track(e, coord_to_name(cur_R, cur_C))
...
if result=="sat":
...
else:
    print s.unsat_core()
```

( <https://github.com/dennis714/yurichev.com/blob/master/blog/spreadsheet/2.py> )

We should explicitly turn on *unsat core* support and use `assert_and_track()` instead of `add()` method, because this feature slows down the whole process, and is turned off by default. That works:

```
% python 2.py test_circular
unsat
[C0, C1]
```

Perhaps, these variables could be removed from the 2D array, marked as *unresolved* and the whole spreadsheet could be recalculated again.

### 5.12.2 Stress test

How to generate large random spreadsheet? What we can do. First, create random DAG<sup>38</sup>, like this one:

<sup>37</sup><http://stackoverflow.com/questions/1832940/is-using-eval-in-python-a-bad-practice>

<sup>38</sup>Directed acyclic graph




Figure 6: Random DAG

Arrows will represent information flow. So a vertex (node) which has no incoming arrows to it (indegree=0), can be set to a random number. Then we use topological sort to find dependencies between vertices. Then we assign spreadsheet cell names to each vertex. Then we generate random expression with random operations/numbers/cells to each cell, with the use of information from topological sorted graph.

Wolfram Mathematica:

```
(* Utility functions *)
In[1]:= findSublistBeforeElementByValue[lst_,element_]:=lst[[ 1;;Position[lst, element][[1]][[1]]-1]]

(* Input in ∞1.. range. 1->A0, 2->A1, etc *)
In[2]:= vertexToName[x_,width_]:=StringJoin[FromCharacterCode[ToCharacterCode["A"]][[1]]+Floor[(x-1)/width],
ToString[Mod[(x-1),width]]]

In[3]:= randomNumberAsString[]:=ToString[RandomInteger[{1,1000}]]

In[4]:= interleaveListWithRandomNumbersAsStrings[lst_]:=Riffle[lst,Table[randomNumberAsString[],Length[lst]-1]]

(* We omit division operation because micro-spreadsheet evaluator can't handle division by zero *)
In[5]:= interleaveListWithRandomOperationsAsStrings[lst_]:=Riffle[lst,Table[RandomChoice[{"+","-","*"}],Length[
lst]-1]]
```

```

In[6]:= randomNonNumberExpression[g_, vertex_] := StringJoin[interleaveListWithRandomOperationsAsStrings[
interleaveListWithRandomNumbersAsStrings[Map[vertexToName[#, WIDTH]&, pickRandomNonDependentVertices[g, vertex
]]]]]

In[7]:= pickRandomNonDependentVertices[g_, vertex_] := DeleteDuplicates[RandomChoice[
findSublistBeforeElementByValue[TopologicalSort[g], vertex], RandomInteger[{1, 5}]]]

In[8]:= assignNumberOrExpr[g_, vertex_] := If[VertexInDegree[g, vertex] == 0, randomNumberAsString[],
randomNonNumberExpression[g, vertex]]

(* Main part *)
(* Create random graph *)
In[21]:= WIDTH=7; HEIGHT=8; TOTAL=WIDTH*HEIGHT
Out[21]= 56

In[24]:= g=DirectedGraph[RandomGraph[BernoulliGraphDistribution[TOTAL, 0.05]], "Acyclic"];

...

(* Generate random expressions and numbers *)
In[26]:= expressions=Map[assignNumberOrExpr[g, #]&, VertexList[g]];

(* Make 2D table of it *)
In[27]:= t=Partition[expressions, WIDTH];

(* Export as tab-separated values *)
In[28]:= Export["/home/dennis/1.txt", t, "TSV"]
Out[28]= /home/dennis/1.txt

In[29]:= Grid[t, Frame->All, Alignment->Left]

```

Here is an output from `Grid[]`:

846	499	A3*913-H4	808	278	303
B4*860-D2	999	59	442	425	A5*163+B2+127*C2*927*D3*213+C1
G6*379-C3-436-C4-289+H6	972	804	D2	G5+108-F1*413-D3	B5
F2	E0	B6-731-D3+791+B4*92+C1	551	F4*922*C2+760*A6-992+B4-184-A4	B1-624-E3
519	G1*402+D1*107*G3-458*A1	D3	B4	B3*811-D3*345+E0	B5
F5-531+B5-222*E4	9	B5+106*B6+600-B1	E3	A5+866*F6+695-A3*226+C6	F4*102*E4*998-H0
C3-956*A5	G4*408-D3*290*B6-899*G5+400+F1	B2-701+H6	A3+782*A5+46-B3-731+C1	42	287
B4-792*H4*407+F6-425-E1	D2	D3	F2-327*G4*35*E1	E1+376*A6-606*F6*554+C5	E3

Using this script, I can generate random spreadsheet of  $26 \cdot 500 = 13000$  cells, which seems to be processed in couple of seconds.

### 5.12.3 The files

The files, including Mathematica notebook: <https://github.com/dennis714/yurichev.com/tree/master/blog/spreadsheet>.

## 6 Program synthesis

Program synthesis is a process of automatic program generation, in accordance with some specific goals.

### 6.1 Synthesis of simple program using Z3 SMT-solver

Sometimes, multiplication operation can be replaced with a several operations of shifting/addition/subtraction. Compilers do so, because pack of instructions can be executed faster.

For example, multiplication by 19 is replaced by GCC 5.4 with pair of instructions: `lea edx, [eax+eax*8]` and

`lea eax, [eax+edx*2]`. This is sometimes also called “superoptimization”.

Let’s see if we can find a shortest possible instructions pack for some specified multiplier.

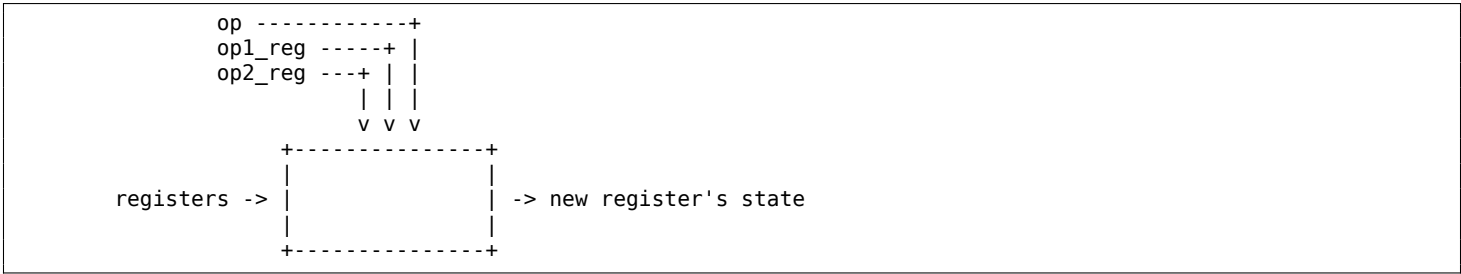
As I’ve already wrote once, SMT-solver can be seen as a solver of huge systems of equations. The task is to construct such system of equations, which, when solved, could produce a short program. I will use electronics analogy here, it can make things a little simpler.

First of all, what our program will be consting of? There will be 3 operations allowed: ADD/SUB/SHL. Only registers allowed as operands, except for the second operand of SHL (which could be in 1..31 range). Each register will be assigned only once (as in SSA<sup>39</sup>).

<sup>39</sup>Static single assignment form



And there will be some “magic block”, which takes all previous register states, it also takes operation type, operands and produces a value of next register’s state.



Now let’s take a look on our schematics on top level:

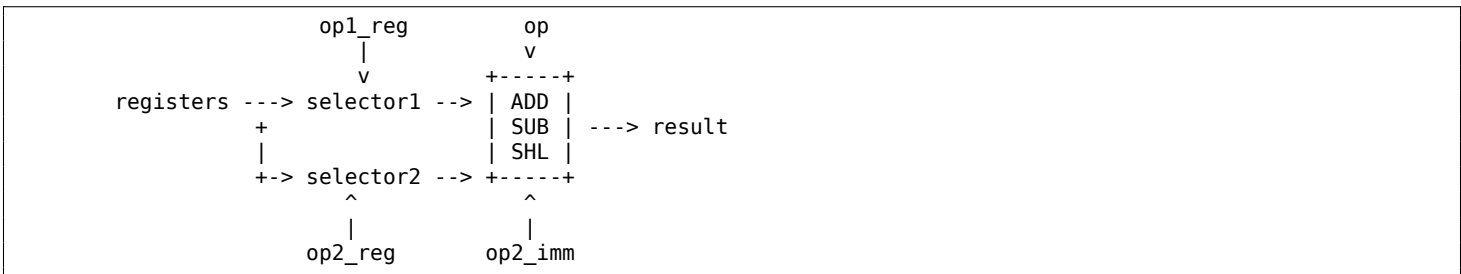
```
0 -> blk -> blk -> blk .. -> blk -> 0
1 -> blk -> blk -> blk .. -> blk -> multiplier
```

Each block takes previous state of registers and produces new states. There are two chains. First chain takes 0 as state of R0 at the very beginning, and the chain is supposed to produce 0 at the end (since zero multiplied by any value is still zero). The second chain takes 1 and must produce multiplier as the state of very last register (since 1 multiplied by multiplier must equal to multiplier).

Each block is “controlled” by operation type, operands, etc. For each column, there is each own set.

Now you can view these two chains as two equations. The ultimate goal is to find such state of all operation types and operands, so the first chain will equal to 0, and the second to multiplier.

Let’s also take a look into “magic block” inside:



Each selector can be viewed as a simple multipositional switch. If operation is SHL, a value in range of 1..31 is used as second operand.

So you can imagine this electric circuit and your goal is to turn all switches in such a state, so two chains will have 0 and multiplier on output. This sounds like logic puzzle in some way. Now we will try to use Z3 to solve this puzzle.

First, we define all variables:

```
R=[[BitVec('S_s%d_c%d' % (s, c), 32) for s in range(MAX_STEPS)] for c in range (CHAINS)]
op=[Int('op_s%d' % s) for s in range(MAX_STEPS)]
op1_reg=[Int('op1_reg_s%d' % s) for s in range(MAX_STEPS)]
op2_reg=[Int('op2_reg_s%d' % s) for s in range(MAX_STEPS)]
op2_imm=[BitVec('op2_imm_s%d' % s, 32) for s in range(MAX_STEPS)]
```

R[][] is registers state for each chain and each step.

On contrary, `op / op1_reg / op2_reg / op2_imm` variables are defined for each step, but for both chains, since both chains at each column has the same operation/operands.

Now we must limit count of operations, and also, register’s number for each step must not be bigger than step number, in other words, instruction at each step is allowed to access only registers which were already set before:

```
for s in range(1, STEPS):
    # for each step
    sl.add(And(op[s]>=0, op[s]<=2))
    sl.add(And(op1_reg[s]>=0, op1_reg[s]<s))
    sl.add(And(op2_reg[s]>=0, op2_reg[s]<s))
    sl.add(And(op2_imm[s]>=1, op2_imm[s]<=31))
```

Fix register of first step for both chains:

```
for c in range(CHAINS):
    # for each chain:
```

```
sl.add(R[c][0]==chain_inputs[c])
sl.add(R[c][STEPS-1]==chain_inputs[c]*multiplier)
```

Now let's add "magic blocks":

```
for s in range(1, STEPS):
    sl.add(R[c][s]==simulate_op(R,c, op[s], op1_reg[s], op2_reg[s], op2_imm[s]))
```

Now how "magic block" is defined?

```
def selector(R, c, s):
    # for all MAX_STEPS:
    return If(s==0, R[c][0],
            If(s==1, R[c][1],
            If(s==2, R[c][2],
            If(s==3, R[c][3],
            If(s==4, R[c][4],
            If(s==5, R[c][5],
            If(s==6, R[c][6],
            If(s==7, R[c][7],
            If(s==8, R[c][8],
            If(s==9, R[c][9],
            0)))))))) # default

def simulate_op(R, c, op, op1_reg, op2_reg, op2_imm):
    op1_val=selector(R,c,op1_reg)
    return If(op==0, op1_val + selector(R, c, op2_reg),
            If(op==1, op1_val - selector(R, c, op2_reg),
            If(op==2, op1_val << op2_imm,
            0))) # default
```

This is very important to understand: if the operation is ADD/SUB, `op2_imm`'s value is just ignored. Otherwise, if operation is SHL, value of `op2_reg` is ignored. Just like in case of digital circuit.

The code: [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/pgm\\_synth/mult.py](https://github.com/dennis714/SAT_SMT_article/blob/master/pgm_synth/mult.py).

Now let's see how it works:

```
% ./mult.py 12
multiplier= 12
attempt, STEPS= 2
unsat
attempt, STEPS= 3
unsat
attempt, STEPS= 4
sat!
r1=SHL r0, 2
r2=SHL r1, 1
r3=ADD r1, r2
tests are OK
```

The first step is always a step containing 0/1, or, r0. So when our solver reporting about 4 steps, this means 3 instructions.

Something harder:

```
% ./mult.py 123
multiplier= 123
attempt, STEPS= 2
unsat
attempt, STEPS= 3
unsat
attempt, STEPS= 4
unsat
attempt, STEPS= 5
sat!
r1=SHL r0, 2
r2=SHL r1, 5
r3=SUB r2, r1
r4=SUB r3, r0
tests are OK
```

Now the code multiplying by 1234:

```
r1=SHL r0, 6
r2=ADD r0, r1
r3=ADD r2, r1
```

```
r4=SHL r2, 4
r5=ADD r2, r3
r6=ADD r5, r4
```

Looks great, but it took  $\approx 23$  seconds to find it on my Intel Xeon CPU E31220 @ 3.10GHz. I agree, this is far from practical usage. Also, I'm not quite sure that this piece of code will work faster than a single multiplication instruction. But anyway, it's a good demonstration of SMT solvers capabilities.

The code multiplying by 12345 ( $\approx 150$  seconds):

```
r1=SHL r0, 5
r2=SHL r0, 3
r3=SUB r2, r1
r4=SUB r1, r3
r5=SHL r3, 9
r6=SUB r4, r5
r7=ADD r0, r6
```

Multiplication by 123456 ( $\approx 8$  minutes!):

```
r1=SHL r0, 9
r2=SHL r0, 13
r3=SHL r0, 2
r4=SUB r1, r2
r5=SUB r3, r4
r6=SHL r5, 4
r7=ADD r1, r6
```

### 6.1.1 Few notes

I've removed SHR instruction support, simply because the code multiplying by a constant makes no use of it. Even more: it's not a problem to add support of constants as second operand for all instructions, but again, you wouldn't find a piece of code which does this job and uses some additional constants. Or maybe I wrong?

Of course, for another job you'll need to add support of constants and other operations. But at the same time, it will work slower and slower. So I had to keep [ISA](#)<sup>40</sup> of this toy [CPU](#)<sup>41</sup> as compact as possible.

### 6.1.2 The code

[https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/pgm\\_synt/mult.py](https://github.com/dennis714/SAT_SMT_article/blob/master/pgm_synt/mult.py).

## 6.2 Rockey dongle: finding unknown algorithm using only input/output pairs

(This text was first published in August 2012 in my blog: <http://blog.yurichev.com/node/71>.)

Some smartcards can execute Java or .NET code - that's the way to hide your sensitive algorithm into chip that is very hard to break (decapsulate). For example, one may encrypt/decrypt data files by hidden crypto algorithm rendering software piracy of such software close to impossible — an encrypted data file created on software with connected smartcard would be impossible to decrypt on cracked version of the same software. (This leads to many nuisances, though.)

That's what called *black box*.

Some software protection dongles offers this functionality too. One example is [Rockey 4](#)<sup>42</sup>.

---

<sup>40</sup>Instruction Set Architecture

<sup>41</sup>Central processing unit

<sup>42</sup><http://www.rockey.nl/en/rockey.html>




Figure 7: Rockey 4 dongle

This is a small dongle connected via USB. It contains some user-defined memory but also memory for user algorithms.

The virtual (toy) CPU for these algorithms is very simple: it offers only 8 16-bit registers (however, only 4 can be set and read) and 8 operations (addition, subtraction, cyclic left shifting, multiplication, OR, XOR, AND, negation).

Second instruction argument can be a constant (from 0 to 63) instead of register.

Each algorithm is described by string like

`A=A+B, B=C*13, D=D^A, C=B*55, C=C&A, D=D|A, A=A*9, A=A&B.`

There are no memory, stack, conditional/unconditional jumps, etc.

Each algorithm, obviously, can't have side effects, so they are actually *pure functions* and their results can be *memoized*.

By the way, as it has been mentioned in Rockey 4 manual, first and last instruction cannot have constants. Maybe that's because these fields used for some internal data: each algorithm start and end should be marked somehow internally anyway.

Would it be possible to reveal hidden impossible-to-read algorithm only by recording input/output dongle traffic? Common sense tells us "no". But we can try anyway.

Since, my goal wasn't to break into some Rockey-protected software, I was interested only in limits (which algorithms could we find), so I make some things simpler: we will work with only 4 16-bit registers, and there will be only 6 operations (add, subtract, multiply, OR, XOR, AND).

Let's first calculate, how much information will be used in brute-force case.

There are 384 of all possible instructions in `reg=reg,op,reg` format for 4 registers and 6 operations, and also 6144 instructions in `reg=reg,op,constant` format. Remember that constant limited to 63 as maximal value? That helps us a little.

So, there are 6528 of all possible instructions. This means, there are  $\approx 1.1 \cdot 10^{19}$  5-instruction algorithms. That's too much.

How can we express each instruction as system of equations? While remembering some school mathematics, I wrote this:

```
Function one_step()=
# Each Bx is integer, but may be only 0 or 1.
# only one of B1..B4 and B5..B9 can be set
reg1=B1*A + B2*B + B3*C + B4*D
reg_or_constant2=B5*A + B6*B + B7*C + B8*D + B9*constant
reg1 should not be equal to reg_or_constant2
# Only one of B10..B15 can be set
result=result+B10*(reg1*reg2)
result=result+B11*(reg1^reg2)
result=result+B12*(reg1+reg2)
result=result+B13*(reg1-reg2)
result=result+B14*(reg1|reg2)
result=result+B15*(reg1&reg2)
B16 - true if register isn't updated in this part
B17 - true if register is updated in this part
(B16 cannot be equal to B17)
A=B16*A + B17*result
B=B18*A + B19*result
C=B20*A + B21*result
```

```
D=B22*A + B23*result
```

That's how we can express each instruction in algorithm.

5-instructions algorithm can be expressed like this:

```
one_step (one_step (one_step (one_step (one_step (input_registers)))))).
```

Let's also add five known input/output pairs and we'll get system of equations like this:

```
one_step (one_step (one_step (one_step (one_step (input_1))))==output_1
one_step (one_step (one_step (one_step (one_step (input_2))))==output_2
one_step (one_step (one_step (one_step (one_step (input_3))))==output_3
one_step (one_step (one_step (one_step (one_step (input_4))))==output_4
.. etc
```

So the question now is to find  $5 \cdot 2^3$  boolean values satisfying known input/output pairs.

I wrote small utility to probe Rockey 4 algorithm with random numbers, it produce results in form:

```
RY_CALCULATE1: (input) p1=30760 p2=18484 p3=41200 p4=61741 (output) p1=49244 p2=11312 p3=27587 p4=12657
RY_CALCULATE1: (input) p1=51139 p2=7852 p3=53038 p4=49378 (output) p1=58991 p2=34134 p3=40662 p4=9869
RY_CALCULATE1: (input) p1=60086 p2=52001 p3=13352 p4=45313 (output) p1=46551 p2=42504 p3=61472 p4=1238
RY_CALCULATE1: (input) p1=48318 p2=6531 p3=51997 p4=30907 (output) p1=54849 p2=20601 p3=31271 p4=44794
```

p1/p2/p3/p4 are just another names for A/B/C/D registers.

Now let's start with Z3. We will need to express Rockey 4 toy CPU in Z3Py (Z3 Python [API](#)) terms.

It can be said, my Python script is divided into two parts:

- constraint definitions (like, *output\_1 should be n for input\_1=m, constant cannot be greater than 63, etc*);
- functions constructing system of equations.

This piece of code define some kind of *structure* consisting of 4 named 16-bit variables, each represent register in our toy CPU.

```
Registers_State=Datatype ('Registers_State')
Registers_State.declare('cons', ('A', BitVecSort(16)), ('B', BitVecSort(16)), ('C', BitVecSort(16)), ('D',
    BitVecSort(16)))
Registers_State=Registers_State.create()
```

These enumerations define two new types (or *sorts* in Z3's terminology):

```
Operation, (OP_MULT, OP_MINUS, OP_PLUS, OP_XOR, OP_OR, OP_AND) = EnumSort('Operation', ('OP_MULT', 'OP_MINUS', '
    OP_PLUS', 'OP_XOR', 'OP_OR', 'OP_AND'))
Register, (A, B, C, D) = EnumSort('Register', ('A', 'B', 'C', 'D'))
```

This part is very important, it defines all variables in our system of equations. `op_step` is type of operation in instruction. `reg_or_constant` is selector between register and constant in second argument — *False* if it's a register and *True* if it's a constant. `reg_step` is a destination register of this instruction. `reg1_step` and `reg2_step` are just registers at arg1 and arg2. `constant_step` is constant (in case it's used in instruction instead of arg2).

```
op_step=[Const('op_step%s' % i, Operation) for i in range(STEPS)]
reg_or_constant_step=[Bool('reg_or_constant_step%s' % i) for i in range(STEPS)]
reg_step=[Const('reg_step%s' % i, Register) for i in range(STEPS)]
reg1_step=[Const('reg1_step%s' % i, Register) for i in range(STEPS)]
reg2_step=[Const('reg2_step%s' % i, Register) for i in range(STEPS)]
constant_step = [BitVec('constant_step%s' % i, 16) for i in range(STEPS)]
```

Adding constraints is very simple. Remember, I wrote that each constant cannot be larger than 63?

```
# according to Rockey 4 dongle manual, arg2 in first and last instructions cannot be a constant
s.add (reg_or_constant_step[0]==False)
s.add (reg_or_constant_step[STEPS-1]==False)

...

for x in range(STEPS):
    s.add (constant_step[x]>=0, constant_step[x]<=63)
```

Known input/output values are added as constraints too.  
 Now let's see how to construct our system of equations:

```
# Register, Registers_State -> int
def register_selector (register, input_registers):
    return If(register==A, Registers_State.A(input_registers),
              If(register==B, Registers_State.B(input_registers),
                If(register==C, Registers_State.C(input_registers),
                  If(register==D, Registers_State.D(input_registers),
                    0)))) # default
```

This function returning corresponding register value from *structure*. Needless to say, the code above is not executed. `If()` is Z3Py function. The code only declares the function, which will be used in another. Expression declaration resembling LISP PL in some way.

Here is another function where `register_selector()` is used:

```
# Bool, Register, Registers_State, int -> int
def register_or_constant_selector (register_or_constant, register, input_registers, constant):
    return If(register_or_constant==False, register_selector(register, input_registers), constant)
```

The code here is never executed too. It only constructs one small piece of very big expression. But for the sake of simplicity, one can think all these functions will be called during bruteforce search, many times, at fastest possible speed.

```
# Operation, Bool, Register, Register, Int, Registers_State -> int
def one_op (op, register_or_constant, reg1, reg2, constant, input_registers):
    arg1=register_selector(reg1, input_registers)
    arg2=register_or_constant_selector (register_or_constant, reg2, input_registers, constant)
    return If(op==OP_MULT, arg1*arg2,
              If(op==OP_MINUS, arg1-arg2,
                If(op==OP_PLUS, arg1+arg2,
                  If(op==OP_XOR, arg1^arg2,
                    If(op==OP_OR, arg1|arg2,
                      If(op==OP_AND, arg1&arg2,
                        0)))))) # default
```

Here is the expression describing each instruction. `new_val` will be assigned to destination register, while all other registers' values are copied from input registers' state:

```
# Bool, Register, Operation, Register, Register, Int, Registers_State -> Registers_State
def one_step (register_or_constant, register_assigned_in_this_step, op, reg1, reg2, constant, input_registers):
    new_val=one_op(op, register_or_constant, reg1, reg2, constant, input_registers)
    return If (register_assigned_in_this_step==A, Registers_State.cons (new_val,
                                                                        Registers_State.B(input_registers),
                                                                        Registers_State.C(input_registers),
                                                                        Registers_State.D(input_registers)),
              If (register_assigned_in_this_step==B, Registers_State.cons (Registers_State.A(input_registers),
                                                                        new_val,
                                                                        Registers_State.C(input_registers),
                                                                        Registers_State.D(input_registers)),
              If (register_assigned_in_this_step==C, Registers_State.cons (Registers_State.A(input_registers),
                                                                        Registers_State.B(input_registers),
                                                                        new_val,
                                                                        Registers_State.D(input_registers)),
              If (register_assigned_in_this_step==D, Registers_State.cons (Registers_State.A(input_registers),
                                                                        Registers_State.B(input_registers),
                                                                        Registers_State.C(input_registers),
                                                                        new_val),
              Registers_State.cons(0,0,0,0)))) # default
```

This is the last function describing a whole *n*-step program:

```
def program(input_registers, STEPS):
    cur_input=input_registers
    for x in range(STEPS):
        cur_input=one_step (reg_or_constant_step[x], reg_step[x], op_step[x], reg1_step[x], reg2_step[x],
                           constant_step[x], cur_input)
    return cur_input
```

Again, for the sake of simplicity, it can be said, now Z3 will try each possible registers/operations/constants against this expression to find such combination which satisfy all input/output pairs. Sounds absurdic, but this is close to reality. SAT/SMT-solvers indeed tries them all. But the trick is to prune search tree as early as possible, so it will work for some reasonable time. And this is hardest problem for solvers.

Now let's start with very simple 3-step algorithm:  $B=A^D$ ,  $C=D*D$ ,  $D=A*C$ . Please note: register **A** left unchanged. I programmed Rockey 4 dongle with the algorithm, and recorded algorithm outputs are:

```
RY_CALCULATE1: (input) p1=8803 p2=59946 p3=36002 p4=44743 (output) p1=8803 p2=36004 p3=7857 p4=24691
RY_CALCULATE1: (input) p1=5814 p2=55512 p3=52155 p4=55813 (output) p1=5814 p2=52403 p3=33817 p4=4038
RY_CALCULATE1: (input) p1=25206 p2=2097 p3=55906 p4=22705 (output) p1=25206 p2=15047 p3=10849 p4=43702
RY_CALCULATE1: (input) p1=10044 p2=14647 p3=27923 p4=7325 (output) p1=10044 p2=15265 p3=47177 p4=20508
RY_CALCULATE1: (input) p1=15267 p2=2690 p3=47355 p4=56073 (output) p1=15267 p2=57514 p3=26193 p4=53395
```

It took about one second and only 5 pairs above to find algorithm (on my quad-core Xeon E3-1220 3.1GHz, however, Z3 solver working in single-thread mode):

```
B = A ^ D
C = D * D
D = C * A
```

Note the last instruction: **C** and **A** registers are swapped comparing to version I wrote by hand. But of course, this instruction is working in the same way, because multiplication is commutative operation.

Now if I try to find 4-step program satisfying to these values, my script will offer this:

```
B = A ^ D
C = D * D
D = A * C
A = A | A
```

...and that's really fun, because the last instruction do nothing with value in register **A**, it's like **NOP**<sup>43</sup>—but still, algorithm is correct for all values given.

Here is another 5-step algorithm:  $B=B^D$ ,  $C=A*22$ ,  $A=B*19$ ,  $A=A\&42$ ,  $D=B\&C$  and values:

```
RY_CALCULATE1: (input) p1=61876 p2=28737 p3=28636 p4=50362 (output) p1=32 p2=46331 p3=50552 p4=33912
RY_CALCULATE1: (input) p1=46843 p2=43355 p3=39078 p4=24552 (output) p1=8 p2=63155 p3=47506 p4=45202
RY_CALCULATE1: (input) p1=22425 p2=51432 p3=40836 p4=14260 (output) p1=0 p2=65372 p3=34598 p4=34564
RY_CALCULATE1: (input) p1=44214 p2=45766 p3=19778 p4=59924 (output) p1=2 p2=22738 p3=55204 p4=20608
RY_CALCULATE1: (input) p1=27348 p2=49060 p3=31736 p4=59576 (output) p1=0 p2=22300 p3=11832 p4=1560
```

It took 37 seconds and we've got:

```
B = D ^ B
C = A * 22
A = B * 19
A = A & 42
D = C & B
```

$A=A\&42$  was correctly deduced (look at these five p1's at output (assigned to output **A** register): 32,8,0,2,0)

6-step algorithm  $A=A+B$ ,  $B=C*13$ ,  $D=D^A$ ,  $C=C\&A$ ,  $D=D|B$ ,  $A=A\&B$  and values:

```
RY_CALCULATE1: (input) p1=4110 p2=35411 p3=54308 p4=47077 (output) p1=32832 p2=50644 p3=36896 p4=60884
RY_CALCULATE1: (input) p1=12038 p2=7312 p3=39626 p4=47017 (output) p1=18434 p2=56386 p3=2690 p4=64639
RY_CALCULATE1: (input) p1=48763 p2=27663 p3=12485 p4=20563 (output) p1=10752 p2=31233 p3=8320 p4=31449
RY_CALCULATE1: (input) p1=33174 p2=38937 p3=54005 p4=38871 (output) p1=4129 p2=46705 p3=4261 p4=48761
RY_CALCULATE1: (input) p1=46587 p2=36275 p3=6090 p4=63976 (output) p1=258 p2=13634 p3=906 p4=48966
```

90 seconds and we've got:

```
A = A + B
B = C * 13
D = D ^ A
D = B | D
C = C & A
A = B & A
```

But that was simple, however. Some 6-step algorithms are not possible to find, for example:

$A=A^B$ ,  $A=A*9$ ,  $A=A^C$ ,  $A=A*19$ ,  $A=A^D$ ,  $A=A\&B$ . Solver was working too long (up to several hours), so I didn't even know is it possible to find it anyway.

## 6.2.1 Conclusion

This is in fact an exercise in program synthesis.

Some short algorithms for tiny CPUs are really possible to find using so small set set of data. Of course it's still not possible to reveal some complex algorithm, but this method definitely should not be ignored.

<sup>43</sup>No Operation

## 6.2.2 The files

Rockey 4 dongle programmer and reader, Rockey 4 manual, Z3Py script for finding algorithms, input/output pairs: [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/pgm\\_synth/rokey\\_files](https://github.com/dennis714/SAT_SMT_article/tree/master/pgm_synth/rokey_files).

## 6.2.3 Further work

Perhaps, constructing LISP-like S-expression can be better than a program for toy-level CPU.

It's also possible to start with smaller constants and then proceed to bigger. This is somewhat similar to increasing password length in password brute-force cracking.

# 7 Toy decompiler

## 7.1 Introduction

A modern-day compiler is a product of hundreds of developer/year. At the same time, toy compiler can be an exercise for a student for a week (or even weekend).

Likewise, commercial decompiler like Hex-Rays can be extremely complex, while toy decompiler like this one, can be easy to understand and remake.

The following decompiler written in Python, supports only short basic blocks, with no jumps. Memory is also not supported.


## 7.2 Data structure

Our toy decompiler will use just one single data structure, representing expression tree.

Many programming textbooks has an example of conversion from Fahrenheit temperature to Celsius, using the following formula:

$$celsius = (fahrenheit - 32) \cdot \frac{5}{9}$$

This expression can be represented as a tree:



How to store it in memory? We see here 3 types of nodes: 1) numbers (or values); 2) arithmetical operations; 3) symbols (like "INPUT").

Many developers with OOP<sup>44</sup> in their mind will create some kind of class. Other developer maybe will use "variant type".

I'll use simplest possible way of representing this structure: a Python tuple. First element of tuple can be a string: either "EXPR\_OP" for operation, "EXPR\_SYMBOL" for symbol or "EXPR\_VALUE" for value. In case of symbol or value, it follows the string. In case of operation, the string followed by another tuples.

Node type and operation type are stored as plain strings—to make debugging output easier to read.

There are *constructors* in our code, in OOP sense:

<sup>44</sup>Object-oriented programming



```

def create_val_expr (val):
    return ("EXPR_VALUE", val)

def create_symbol_expr (val):
    return ("EXPR_SYMBOL", val)

def create_binary_expr (op, op1, op2):
    return ("EXPR_OP", op, op1, op2)

```

There are also *accessors*:

```

def get_expr_type(e):
    return e[0]

def get_symbol (e):
    assert get_expr_type(e)=="EXPR_SYMBOL"
    return e[1]

def get_val (e):
    assert get_expr_type(e)=="EXPR_VALUE"
    return e[1]

def is_expr_op(e):
    return get_expr_type(e)=="EXPR_OP"

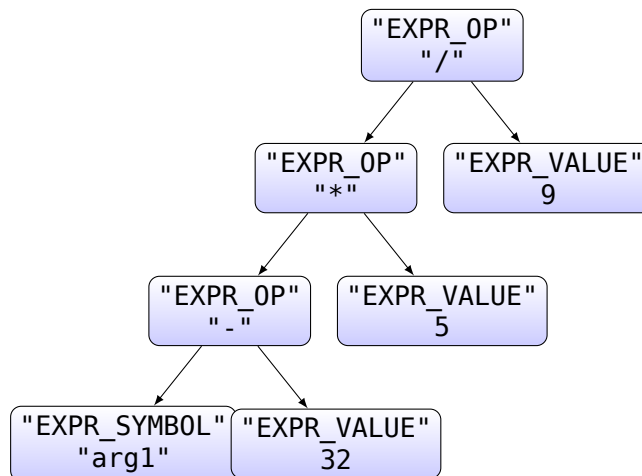
def get_op (e):
    assert is_expr_op(e)
    return e[1]

def get_op1 (e):
    assert is_expr_op(e)
    return e[2]

def get_op2 (e):
    assert is_expr_op(e)
    return e[3]

```

The temperature conversion expression we just saw will be represented as:



...or as Python expression:

```

('EXPR_OP', '/',
 ('EXPR_OP', '*',
 ('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32)),
 ('EXPR_VALUE', 5)),
 ('EXPR_VALUE', 9))

```

In fact, this is [AST<sup>45</sup>](#) in its simplest form. [ASTs](#) are used heavily in compilers.

## 7.3 Simple examples

Let's start with simplest example:

<sup>45</sup>Abstract syntax tree

```
mov    rax, rdi
imul   rax, rsi
```

At start, these symbols are assigned to registers: RAX=initial\_RAX, RBX=initial\_RBX, RDI=arg1, RSI=arg2, RDX=arg3, RCX=arg4.

When we handle MOV instruction, we just copy expression from RDI to RAX. When we handle IMUL instruction, we create a new expression, adding together expressions from RAX and RSI and putting result into RAX again.

I can feed this to decompiler and we will see how register's state is changed through processing:

```
python td.py --show-registers --python-expr tests/mul.s
...
line=[mov    rax, rdi]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_SYMBOL', 'initial_RBX')
rdx=('EXPR_SYMBOL', 'arg3')
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_SYMBOL', 'arg1')

line=[imul   rax, rsi]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_SYMBOL', 'initial_RBX')
rdx=('EXPR_SYMBOL', 'arg3')
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_SYMBOL', 'arg2'))
...
result=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_SYMBOL', 'arg2'))
```

IMUL instruction is mapped to "\*" string, and then new expression is constructed in `handle_binary_op()`, which puts result into RAX.

In this output, the data structures are dumped using Python `str()` function, which does mostly the same, as `print()`.

Output is bulky, and we can turn off Python expressions output, and see how this internal data structure can be rendered neatly using our internal `expr_to_string()` function:

```
python td.py --show-registers tests/mul.s
...
line=[mov    rax, rdi]
rcx=arg4
rsi=arg2
rbx=initial_RBX
rdx=arg3
rdi=arg1
rax=arg1

line=[imul   rax, rsi]
rcx=arg4
rsi=arg2
rbx=initial_RBX
rdx=arg3
rdi=arg1
rax=(arg1 * arg2)
...
result=(arg1 * arg2)
```

Slightly advanced example:

```
imul   rdi, rsi
leaq   rax, [rdi+rdx]
```

LEA instruction is treated just as ADD.

```
python td.py --show-registers --python-expr tests/mul_add.s
```

```
...  
line=[imul      rdi, rsi]  
rcx=('EXPR_SYMBOL', 'arg4')  
rsi=('EXPR_SYMBOL', 'arg2')  
rbx=('EXPR_SYMBOL', 'initial_RBX')  
rdx=('EXPR_SYMBOL', 'arg3')  
rdi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_SYMBOL', 'arg2'))  
rax=('EXPR_SYMBOL', 'initial_RAX')  
  
line=[lea      rax, [rdi+rdx]]  
rcx=('EXPR_SYMBOL', 'arg4')  
rsi=('EXPR_SYMBOL', 'arg2')  
rbx=('EXPR_SYMBOL', 'initial_RBX')  
rdx=('EXPR_SYMBOL', 'arg3')  
rdi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_SYMBOL', 'arg2'))  
rax=('EXPR_OP', '+', ('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_SYMBOL', 'arg2')), ('EXPR_SYMBOL', 'arg3'))  
)  
  
...  
result=('EXPR_OP', '+', ('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_SYMBOL', 'arg2')), ('EXPR_SYMBOL', 'arg3'))
```

And again, let's see this expression dumped neatly:

```
python td.py --show-registers tests/mul_add.s
```

```
...  
result=((arg1 * arg2) + arg3)
```

Now another example, where we use 2 input arguments:

```
imul    rdi, rdi, 1234  
imul    rsi, rsi, 5678  
lea     rax, [rdi+rsi]
```

```
python td.py --show-registers --python-expr tests/mul_add3.s
```

```
...  
line=[imul      rdi, rdi, 1234]  
rcx=('EXPR_SYMBOL', 'arg4')  
rsi=('EXPR_SYMBOL', 'arg2')  
rbx=('EXPR_SYMBOL', 'initial_RBX')  
rdx=('EXPR_SYMBOL', 'arg3')  
rdi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 1234))  
rax=('EXPR_SYMBOL', 'initial_RAX')  
  
line=[imul      rsi, rsi, 5678]  
rcx=('EXPR_SYMBOL', 'arg4')  
rsi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg2'), ('EXPR_VALUE', 5678))  
rbx=('EXPR_SYMBOL', 'initial_RBX')  
rdx=('EXPR_SYMBOL', 'arg3')  
rdi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 1234))  
rax=('EXPR_SYMBOL', 'initial_RAX')  
  
line=[lea      rax, [rdi+rsi]]  
rcx=('EXPR_SYMBOL', 'arg4')  
rsi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg2'), ('EXPR_VALUE', 5678))  
rbx=('EXPR_SYMBOL', 'initial_RBX')  
rdx=('EXPR_SYMBOL', 'arg3')  
rdi=('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 1234))  
rax=('EXPR_OP', '+', ('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 1234)), ('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg2'), ('EXPR_VALUE', 5678)))  
)  
  
...  
result=('EXPR_OP', '+', ('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 1234)), ('EXPR_OP', '*', ('EXPR_SYMBOL', 'arg2'), ('EXPR_VALUE', 5678)))
```

...and now neat output:

```
python td.py --show-registers tests/mul_add3.s
```

```
...
```

```
result=((arg1 * 1234) + (arg2 * 5678))
```

Now conversion program:

```
mov    rax, rdi
sub    rax, 32
imul   rax, 5
mov    rbx, 9
idiv   rbx
```

You can see, how register's state is changed over execution (or parsing).

Raw:

```
python td.py --show-registers --python-expr tests/fahr_to_celsius.s
```

```
...
```

```
line=[mov    rax, rdi]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_SYMBOL', 'initial_RBX')
rdx=('EXPR_SYMBOL', 'arg3')
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_SYMBOL', 'arg1')
```

```
line=[sub    rax, 32]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_SYMBOL', 'initial_RBX')
rdx=('EXPR_SYMBOL', 'arg3')
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32))
```

```
line=[imul   rax, 5]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_SYMBOL', 'initial_RBX')
rdx=('EXPR_SYMBOL', 'arg3')
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_OP', '*', ('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32)), ('EXPR_VALUE', 5))
```

```
line=[mov    rbx, 9]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_VALUE', 9)
rdx=('EXPR_SYMBOL', 'arg3')
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_OP', '*', ('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32)), ('EXPR_VALUE', 5))
```

```
line=[idiv   rbx]
rcx=('EXPR_SYMBOL', 'arg4')
rsi=('EXPR_SYMBOL', 'arg2')
rbx=('EXPR_VALUE', 9)
rdx=('EXPR_OP', '%', ('EXPR_OP', '*', ('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32)), ('EXPR_VALUE', 5)), ('EXPR_VALUE', 9))
rdi=('EXPR_SYMBOL', 'arg1')
rax=('EXPR_OP', '/', ('EXPR_OP', '*', ('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32)), ('EXPR_VALUE', 5)), ('EXPR_VALUE', 9))
```

```
...
```

```
result=('EXPR_OP', '/', ('EXPR_OP', '*', ('EXPR_OP', '-', ('EXPR_SYMBOL', 'arg1'), ('EXPR_VALUE', 32)), ('EXPR_VALUE', 5)), ('EXPR_VALUE', 9))
```

Neat:

```
python td.py --show-registers tests/fahr_to_celsius.s
```

```
...
```

```
line=[mov    rax, rdi]
```

```

rcx=arg4
rsi=arg2
rbx=initial_RBX
rdx=arg3
rdi=arg1
rax=arg1

line=[sub      rax, 32]
rcx=arg4
rsi=arg2
rbx=initial_RBX
rdx=arg3
rdi=arg1
rax=(arg1 - 32)

line=[imul     rax, 5]
rcx=arg4
rsi=arg2
rbx=initial_RBX
rdx=arg3
rdi=arg1
rax=((arg1 - 32) * 5)

line=[mov      rbx, 9]
rcx=arg4
rsi=arg2
rbx=9
rdx=arg3
rdi=arg1
rax=((arg1 - 32) * 5)

line=[idiv     rbx]
rcx=arg4
rsi=arg2
rbx=9
rdx=(((arg1 - 32) * 5) % 9)
rdi=arg1
rax=(((arg1 - 32) * 5) / 9)

...

result=(((arg1 - 32) * 5) / 9)

```

It is interesting to note that IDIV instruction also calculates remainder of division, and it is placed into RDX register. It's not used, but is available for use.

This is how quotient and remainder are stored in registers:

```

def handle_unary_DIV_IDIV (registers, op1):
    op1_expr=register_or_number_in_string_to_expr (registers, op1)
    current_RAX=registers["rax"]
    registers["rax"]=create_binary_expr ("/", current_RAX, op1_expr)
    registers["rdx"]=create_binary_expr ("%", current_RAX, op1_expr)

```

Now this is `align2grain()` function<sup>46</sup>:

```

; uint64_t align2grain (uint64_t i, uint64_t grain)
;   return ((i + grain-1) & ~(grain-1));

; rdi=i
; rsi=grain

sub    rsi, 1
add    rdi, rsi
not    rsi
and    rdi, rsi
mov    rax, rdi

```

```

...

line=[sub      rsi, 1]
rcx=arg4
rsi=(arg2 - 1)

```

<sup>46</sup>Taken from <https://docs.oracle.com/javase/specs/jvms/se6/html/Compiling.doc.html>

```

rbx=initial_RBX
rdx=arg3
rdi=arg1
rax=initial_RAX

line=[add      rdi, rsi]
rcx=arg4
rsi=(arg2 - 1)
rbx=initial_RBX
rdx=arg3
rdi=(arg1 + (arg2 - 1))
rax=initial_RAX

line=[not      rsi]
rcx=arg4
rsi=(~(arg2 - 1))
rbx=initial_RBX
rdx=arg3
rdi=(arg1 + (arg2 - 1))
rax=initial_RAX

line=[and      rdi, rsi]
rcx=arg4
rsi=(~(arg2 - 1))
rbx=initial_RBX
rdx=arg3
rdi=((arg1 + (arg2 - 1)) & ~(arg2 - 1))
rax=initial_RAX

line=[mov      rax, rdi]
rcx=arg4
rsi=(~(arg2 - 1))
rbx=initial_RBX
rdx=arg3
rdi=((arg1 + (arg2 - 1)) & ~(arg2 - 1))
rax=((arg1 + (arg2 - 1)) & ~(arg2 - 1))

...

result=((arg1 + (arg2 - 1)) & ~(arg2 - 1))

```

## 7.4 Dealing with compiler optimizations

The following piece of code ...

```

mov    rax, rdi
add    rax, rax

```

...will be transformed into  $(arg1 + arg1)$  expression. It can be reduced to  $(arg1 * 2)$ . Our toy decompiler can identify patterns like such and rewrite them.

```

# X+X -> X*2
def reduce_ADD1 (expr):
    if is_expr_op(expr) and get_op (expr)=="+" and get_op1 (expr)==get_op2 (expr):
        return dbg_print_reduced_expr ("reduce_ADD1", expr, create_binary_expr ("*", get_op1 (expr),
            create_val_expr (2)))

    return expr # no match

```

This function will just test, if the current node has `EXPR_OP` type, operation is "+" and both children are equal to each other. By the way, since our data structure is just tuple of tuples, Python can compare them using plain "==" operation. If the testing is finished successfully, current node is then replaced with a new expression: we take one of children, we construct a node of `EXPR_VALUE` type with "2" number in it, and then we construct a node of `EXPR_OP` type with "\*".

`dbg_print_reduced_expr()` serving solely debugging purposes—it just prints the old and the new (reduced) expressions.

Decompiler is then traverse expression tree recursively in *deep-first search* fashion.

```

def reduce_step (e):
    if is_expr_op (e)==False:
        return e # expr isn't EXPR_OP, nothing to reduce (we don't reduce EXPR_SYMBOL and EXPR_VAL)

```

```

if is_unary_op(get_op(e)):
    # recreate expr with reduced operand:
    return reducers(create_unary_expr (get_op(e), reduce_step (get_op1 (e))))
else:
    # recreate expr with both reduced operands:
    return reducers(create_binary_expr (get_op(e), reduce_step (get_op1 (e)), reduce_step (get_op2 (e))))
...

# same as "return ... (reduce_MUL1 (reduce_ADD1 (reduce_ADD2 (... expr))))"
reducers=compose([
    ...
    reduce_ADD1, ...
    ...])

def reduce (e):
    print "going to reduce " + expr_to_string (e)
    new_expr=reduce_step(e)
    if new_expr==e:
        return new_expr # we are done here, expression can't be reduced further
    else:
        return reduce(new_expr) # reduced expr has been changed, so try to reduce it again

```

Reduction functions called again and again, as long, as expression changes.  
 Now we run it:

```

python td.py tests/add1.s
...
going to reduce (arg1 + arg1)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
going to reduce (arg1 * 2)
result=(arg1 * 2)

```

So far so good, now what if we would try this piece of code?

```

mov    rax, rdi
add    rax, rax
add    rax, rax
add    rax, rax

```

```

python td.py tests/add2.s
...
working out tests/add2.s
going to reduce (((arg1 + arg1) + (arg1 + arg1)) + ((arg1 + arg1) + (arg1 + arg1)))
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() ((arg1 * 2) + (arg1 * 2)) -> ((arg1 * 2) * 2)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() ((arg1 * 2) + (arg1 * 2)) -> ((arg1 * 2) * 2)
reduction in reduce_ADD1() (((arg1 * 2) * 2) + ((arg1 * 2) * 2)) -> (((arg1 * 2) * 2) * 2)
going to reduce (((arg1 * 2) * 2) * 2)
result=(((arg1 * 2) * 2) * 2)

```

This is correct, but too verbose.

We would like to rewrite  $(X^n)*m$  expression to  $X*(n*m)$ , where  $n$  and  $m$  are numbers. We can do this by adding another function like `reduce_ADD1()`, but there is much better option: we can make matcher for tree. You can think about it as regular expression matcher, but over trees.

```

def bind_expr (key):
    return ("EXPR_WILDCARD", key)

def bind_value (key):
    return ("EXPR_WILDCARD_VALUE", key)

def match_EXPR_WILDCARD (expr, pattern):
    return {pattern[1] : expr} # return {key : expr}

```

```

def match_EXPR_WILDCARD_VALUE (expr, pattern):
    if get_expr_type (expr)!="EXPR_VALUE":
        return None
    return {pattern[1] : get_val(expr)} # return {key : expr}

def is_commutative (op):
    return op in ["+", "*", "&", "|", "^"]

def match_two_ops (op1_expr, op1_pattern, op2_expr, op2_pattern):
    m1=match (op1_expr, op1_pattern)
    m2=match (op2_expr, op2_pattern)
    if m1==None or m2==None:
        return None # one of match for operands returned False, so we do the same
    # join two dicts from both operands:
    rt={}
    rt.update(m1)
    rt.update(m2)
    return rt

def match_EXPR_OP (expr, pattern):
    if get_expr_type(expr)!=get_expr_type(pattern): # be sure, both EXPR_OP.
        return None
    if get_op (expr)!=get_op (pattern): # be sure, ops type are the same.
        return None

    if (is_unary_op(get_op(expr))):
        # match unary expression.
        return match (get_op1 (expr), get_op1 (pattern))
    else:
        # match binary expression.

        # first try match operands as is.
        m=match_two_ops (get_op1 (expr), get_op1 (pattern), get_op2 (expr), get_op2 (pattern))
        if m!=None:
            return m
        # if matching unsuccessful, AND operation is commutative, try also swapped operands.
        if is_commutative (get_op (expr))==False:
            return None
        return match_two_ops (get_op1 (expr), get_op2 (pattern), get_op2 (expr), get_op1 (pattern))

# returns dict in case of success, or None
def match (expr, pattern):
    t=get_expr_type(pattern)
    if t=="EXPR_WILDCARD":
        return match_EXPR_WILDCARD (expr, pattern)
    elif t=="EXPR_WILDCARD_VALUE":
        return match_EXPR_WILDCARD_VALUE (expr, pattern)
    elif t=="EXPR_SYMBOL":
        if expr==pattern:
            return {}
        else:
            return None
    elif t=="EXPR_VALUE":
        if expr==pattern:
            return {}
        else:
            return None
    elif t=="EXPR_OP":
        return match_EXPR_OP (expr, pattern)
    else:
        raise AssertionError

```

Now how we will use it:

```

# (X*A)*B -> X*(A*B)
def reduce_MUL1 (expr):
    m=match (expr, create_binary_expr ("*", (create_binary_expr ("*", bind_expr ("X"), bind_value ("A")),
    bind_value ("B"))))
    if m==None:
        return expr # no match

    return dbg_print_reduced_expr ("reduce_MUL1", expr, create_binary_expr ("*",
    m["X"], # new op1
    create_val_expr (m["A"] * m["B"]))) # new op2

```



We take input expression, and we also construct pattern to be matched. Matcher works recursively over both expressions synchronously. Pattern is also expression, but can use two additional node types: *EXPR\_WILDCARD* and *EXPR\_WILDCARD\_VALUE*. These nodes are supplied with keys (stored as strings). When matcher encounters *EXPR\_WILDCARD* in pattern, it just stashes current expression and will return it. If matcher encounters *EXPR\_WILDCARD\_VALUE*, it does the same, but only in case the current node has *EXPR\_VALUE* type.

`bind_expr()` and `bind_value()` are functions which create nodes with the types we have seen.

All this means, `reduce_MUL1()` function will search for the expression in form  $(X*A)*B$ , where  $A$  and  $B$  are numbers. In other cases, matcher will return input expression untouched, so these reducing function can be chained.

Now when `reduce_MUL1()` encounters (sub)expression we are interesting in, it will return dictionary with keys and expressions. Let's add `print m` call somewhere before return and rerun:

```
python td.py tests/add2.s
...
going to reduce (((arg1 + arg1) + (arg1 + arg1)) + ((arg1 + arg1) + (arg1 + arg1)))
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() ((arg1 * 2) + (arg1 * 2)) -> ((arg1 * 2) * 2)
{'A': 2, 'X': ('EXPR_SYMBOL', 'arg1'), 'B': 2}
reduction in reduce_MUL1() ((arg1 * 2) * 2) -> (arg1 * 4)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() (arg1 + arg1) -> (arg1 * 2)
reduction in reduce_ADD1() ((arg1 * 2) + (arg1 * 2)) -> ((arg1 * 2) * 2)
{'A': 2, 'X': ('EXPR_SYMBOL', 'arg1'), 'B': 2}
reduction in reduce_MUL1() ((arg1 * 2) * 2) -> (arg1 * 4)
reduction in reduce_ADD1() ((arg1 * 4) + (arg1 * 4)) -> ((arg1 * 4) * 2)
{'A': 4, 'X': ('EXPR_SYMBOL', 'arg1'), 'B': 2}
reduction in reduce_MUL1() ((arg1 * 4) * 2) -> (arg1 * 8)
going to reduce (arg1 * 8)
...
result=(arg1 * 8)
```

The dictionary has keys we supplied plus expressions matcher found. We then can use them to create new expression and return it. Numbers are just summed while forming second operand to "\*" operation.

Now a real-world optimization technique—optimizing GCC replaced multiplication by 31 by shifting and subtraction operations:

```
mov    rax, rdi
sal    rax, 5
sub    rax, rdi
```

Without reduction functions, our decompiler will translate this into  $((arg1 \ll 5) - arg1)$ . We can replace shifting left by multiplication:

```
# X<<n -> X*(2^n)
def reduce_SHL1 (expr):
    m=match (expr, create_binary_expr("<<", bind_expr("X"), bind_value("Y")))
    if m==None:
        return expr # no match

    return dbg_print_reduced_expr ("reduce_SHL1", expr, create_binary_expr ("*", m["X"], create_val_expr (1<<m["Y"])))
```

Now we getting  $((arg1 * 32) - arg1)$ . We can add another reduction function:

```
# (X*n)-X -> X*(n-1)
def reduce_SUB3 (expr):
    m=match (expr, create_binary_expr("-",
        create_binary_expr ("*", bind_expr("X1"), bind_value("N")),
        bind_expr("X2")))

    if m!=None and match (m["X1"], m["X2"])!=None:
        return dbg_print_reduced_expr ("reduce_SUB3", expr, create_binary_expr ("*", m["X1"], create_val_expr (m["N"]-1)))
    else:
        return expr # no match
```

Matcher will return two X's, and we must be assured that they are equal. In fact, in previous versions of this toy decompiler, I did comparison with plain "==" , and it worked. But we can reuse `match()` function for the same purpose, because it will process commutative operations better. For example, if X1 is "Q+1" and X2 is "1+Q", expressions are equal, but plain "==" will not work. On the other side, `match()` function, when encounter "+" operation (or another commutative operation), and it fails with comparison, it will also try swapped operand and will try to compare again.

However, to understand it easier, for a moment, you can imagine there is "==" instead of the second `match()` .

Anyway, here is what we've got:

```
working out tests/mul31_GCC.s
going to reduce ((arg1 << 5) - arg1)
reduction in reduce_SHL1() (arg1 << 5) -> (arg1 * 32)
reduction in reduce_SUB3() ((arg1 * 32) - arg1) -> (arg1 * 31)
going to reduce (arg1 * 31)
...
result=(arg1 * 31)
```

Another optimization technique is often seen in ARM thumb code: AND-ing a value with a value like 0xFFFFF0, is implemented using shifts:

```
mov rax, rdi
shr rax, 4
shl rax, 4
```

This code is quite common in ARM thumb code, because it's a headache to encode 32-bit constants using couple of 16-bit thumb instructions, while single 16-bit instruction can shift by 4 bits left or right.

Also, the expression  $(x \gg 4) \ll 4$  can be jokingly called as "twitching operator": I've heard the "--i++" expression was called like this in Russian-speaking social networks, it was some kind of meme ("operator podergivaniya").

Anyway, these reduction functions will be used:

```
# X>>n -> X / (2^n)
...
def reduce_SHR2 (expr):
    m=match(expr, create_binary_expr(">>", bind_expr("X"), bind_value("Y")))
    if m==None or m["Y"]>=64:
        return expr # no match

    return dbg_print_reduced_expr ("reduce_SHR2", expr, create_binary_expr ("/", m["X"],
        create_val_expr (1<<m["Y"])))

...

# X<<n -> X*(2^n)
def reduce_SHL1 (expr):
    m=match (expr, create_binary_expr("<<", bind_expr ("X"), bind_value ("Y")))
    if m==None:
        return expr # no match

    return dbg_print_reduced_expr ("reduce_SHL1", expr, create_binary_expr ("*", m["X"], create_val_expr (1<<m["Y"])))

...

# FIXME: slow
# returns True if n=2^x or popcnt(n)=1
def is_2n(n):
    return bin(n).count("1")==1

# AND operation using DIV/MUL or SHL/SHR
# (X / (2^n)) * (2^n) -> X&(~((2^n)-1))
def reduce_MUL2 (expr):
    m=match(expr, create_binary_expr ("*", create_binary_expr ("/", bind_expr("X"), bind_value("N1")),
        bind_value("N2")))
    if m==None or m["N1"]!=m["N2"] or is_2n(m["N1"])==False: # short-circuit expression
        return expr # no match

    return dbg_print_reduced_expr("reduce_MUL2", expr, create_binary_expr("&", m["X"],
        create_val_expr(~(m["N1"]-1)&0xffffffff)))
```

Now the result:

```
working out tests/AND_by_shifts2.s
going to reduce ((arg1 >> 4) << 4)
reduction in reduce_SHR2() (arg1 >> 4) -> (arg1 / 16)
reduction in reduce_SHL1() ((arg1 / 16) << 4) -> ((arg1 / 16) * 16)
reduction in reduce_MUL2() ((arg1 / 16) * 16) -> (arg1 & 0xfffffffffffff0)
going to reduce (arg1 & 0xfffffffffffff0)
...
result=(arg1 & 0xfffffffffffff0)
```

### 7.4.1 Division using multiplication

Division is often replaced by multiplication for performance reasons.

From school-level arithmetics, we can remember that division by 3 can be replaced by multiplication by  $\frac{1}{3}$ . In fact, sometimes compilers do so for floating-point arithmetics, for example, FDIV instruction in x86 code can be replaced by FMUL. At least MSVC 6.0 will replace division by 3 by multiplication by  $\frac{1}{3}$  and sometimes it's hard to be sure, what operation was in original source code.

But when we operate over integer values and CPU registers, we can't use fractions. However, we can rework fraction:

$$result = \frac{x}{3} = x \cdot \frac{1}{3} = x \cdot \frac{1 \cdot MagicNumber}{3 \cdot MagicNumber}$$

Given the fact that division by  $2^n$  is very fast, we now should find that *MagicNumber*, for which the following equation will be true:  $2^n = 3 \cdot MagicNumber$ .

This code performing division by 10:

```
mov    rax, rdi
movabs rdx, 0cccccccccccccdh
mul    rdx
shr    rdx, 3
mov    rax, rdx
```

Division by  $2^{64}$  is somewhat hidden: lower 64-bit of product in RAX is not used (dropped), only higher 64-bit of product (in RDX) is used and then shifted by additional 3 bits.

RDX register is set during processing of MUL/IMUL like this:

```
def handle_unary_MUL_IMUL (registers, op1):
    op1_expr=register_or_number_in_string_to_expr (registers, op1)
    result=create_binary_expr ("*", registers["rax"], op1_expr)
    registers["rax"]=result
    registers["rdx"]=create_binary_expr (">>", result, create_val_expr(64))
```

In other words, the assembly code we have just seen multiplies by  $\frac{0cccccccccccccdh}{2^{64+3}}$ , or divides by  $\frac{2^{64+3}}{0cccccccccccccdh}$ . To find divisor we just have to divide numerator by denominator.

```
# n = magic number
# m = shifting coefficient
# return = 1 / (n / 2^m) = 2^m / n
def get_divisor (n, m):
    return (2**float(m))/float(n)

# (X*n)>>m, where m>=64 -> X/...
def reduce_div_by_MUL (expr):
    m=match (expr, create_binary_expr(">>", create_binary_expr ("*", bind_expr("X"), bind_value("N")), bind_value("M")))
    if m==None:
        return expr # no match

    divisor=get_divisor(m["N"], m["M"])
    return dbg_print_reduced_expr ("reduce_div_by_MUL", expr, create_binary_expr ("/", m["X"], create_val_expr (int(divisor))))
```

This works, but we have a problem: this rule takes  $(arg1 * 0cccccccccccccd) \gg 64$  expression first and finds divisor to be equal to 1.25. This is correct: result is shifted by 3 bits after (or divided by 8), and  $1.25 \cdot 8 = 10$ . But our toy decompiler doesn't support real numbers.

We can solve this problem in the following way: if divisor has fractional part, we postpone reducing, with a hope, that two subsequent right shift operations will be reduced into single one:

```
# (X*n)>>m, where m>=64 -> X/...
def reduce_div_by_MUL (expr):
    m=match (expr, create_binary_expr(">>", create_binary_expr ("*", bind_expr("X"), bind_value("N")),
        bind_value("M")))
    if m==None:
        return expr # no match

    divisor=get_divisor(m["N"], m["M"])
    if math.floor(divisor)==divisor:
        return dbg_print_reduced_expr ("reduce_div_by_MUL", expr, create_binary_expr ("/", m["X"],
            create_val_expr (int(divisor))))
    else:
        print "reduce_div_by_MUL(): postponing reduction, because divisor=", divisor
        return expr
```

That works:

```
working out tests/div_by_mult10_unsigned.s
going to reduce (((arg1 * 0xffffffff) >> 64) >> 3)
reduce_div_by_MUL(): postponing reduction, because divisor= 1.25
reduction in reduce_SHR1() (((arg1 * 0xffffffff) >> 64) >> 3) -> ((arg1 * 0xffffffff) >> 67)
going to reduce ((arg1 * 0xffffffff) >> 67)
reduction in reduce_div_by_MUL() ((arg1 * 0xffffffff) >> 67) -> (arg1 / 10)
going to reduce (arg1 / 10)
result=(arg1 / 10)
```

I don't know if this is best solution. In early version of this decompiler, it processed input expression in two passes: first pass for everything except division by multiplication, and the second pass for the latter. I don't know which way is better. Or maybe we could support real numbers in expressions?

Couple of words about better understanding division by multiplication. Many people miss "hidden" division by  $2^{32}$  or  $2^{64}$ , when lower 32-bit part (or 64-bit part) of product is not used (or just dropped). Also, there is misconception that modulo inverse is used here. This is close, but not the same thing. Extended Euclidean algorithm is usually used to find *magic coefficient*, but in fact, this algorithm is rather used to solve the equation. You can solve it using any other method. Also, needless to mention, the equation is unsolvable for some divisors, because this is diophantine equation (i.e., equation allowing result to be only integer), since we work on integer CPU registers, after all.

## 7.5 Obfuscation/deobfuscation

Despite simplicity of our decompiler, we can see how to deobfuscate (or optimize) using several simple tricks.

For example, this piece of code does nothing:

```
mov rax, rdi
xor rax, 12345678h
xor rax, 0deadbeefh
xor rax, 12345678h
xor rax, 0deadbeefh
```

We would need these rules to tame it:

```
# (X^n)^m -> X^(n^m)
def reduce_XOR4 (expr):
    m=match(expr,
        create_binary_expr("^",
            create_binary_expr ("^", bind_expr("X"), bind_value("N")),
            bind_value("M")))

    if m!=None:
        return dbg_print_reduced_expr ("reduce_XOR4", expr, create_binary_expr ("^", m["X"],
            create_val_expr (m["N"]^m["M"])))
    else:
        return expr # no match

...

# X op 0 -> X, where op is ADD, OR, XOR, SUB
def reduce_op_0 (expr):
    # try each:
    for op in ["+", "|", "^", "-"]:
        m=match(expr, create_binary_expr(op, bind_expr("X"), create_val_expr (0)))
        if m!=None:
```

```

return dbg_print_reduced_expr ("reduce_op_0", expr, m["X"])

# default:
return expr # no match

```

```

working out tests/t9_obf.s
going to reduce (((arg1 ^ 0x12345678) ^ 0xdeadbeef) ^ 0x12345678) ^ 0xdeadbeef)
reduction in reduce_XOR4() ((arg1 ^ 0x12345678) ^ 0xdeadbeef) -> (arg1 ^ 0xcc99e897)
reduction in reduce_XOR4() ((arg1 ^ 0xcc99e897) ^ 0x12345678) -> (arg1 ^ 0xdeadbeef)
reduction in reduce_XOR4() ((arg1 ^ 0xdeadbeef) ^ 0xdeadbeef) -> (arg1 ^ 0x0)
going to reduce (arg1 ^ 0x0)
reduction in reduce_op_0() (arg1 ^ 0x0) -> arg1
going to reduce arg1
result=arg1

```

This piece of code can be deobfuscated (or optimized) as well:

```

; toggle last bit

mov rax, rdi
mov rbx, rax
mov rcx, rbx
mov rsi, rcx
xor rsi, 12345678h
xor rsi, 12345679h
mov rax, rsi

```

```

working out tests/t7_obf.s
going to reduce ((arg1 ^ 0x12345678) ^ 0x12345679)
reduction in reduce_XOR4() ((arg1 ^ 0x12345678) ^ 0x12345679) -> (arg1 ^ 1)
going to reduce (arg1 ^ 1)
result=(arg1 ^ 1)

```

I also used *aha!*<sup>47</sup> superoptimizer to find weird piece of code which does nothing.

*Aha!* is so called superoptimizer, it tries various piece of codes in brute-force manner, in attempt to find shortest possible alternative for some mathematical operation. While sane compiler developers use superoptimizers for this task, I tried it in opposite way, to find oddest pieces of code for some simple operations, including **NOP** operation. In past, I've used it to find weird alternative to XOR operation (5.7).

So here is what *aha!* can find for **NOP**:

```

; do nothing (as found by aha)

mov rax, rdi
and rax, rax
or rax, rax

```

```

# X & X -> X
def reduce_AND3 (expr):
    m=match (expr, create_binary_expr ("&", bind_expr ("X1"), bind_expr ("X2")))
    if m!=None and match (m["X1"], m["X2"])!=None:
        return dbg_print_reduced_expr("reduce_AND3", expr, m["X1"])
    else:
        return expr # no match

...

# X | X -> X
def reduce_OR1 (expr):
    m=match (expr, create_binary_expr ("|", bind_expr ("X1"), bind_expr ("X2")))
    if m!=None and match (m["X1"], m["X2"])!=None:
        return dbg_print_reduced_expr("reduce_OR1", expr, m["X1"])
    else:
        return expr # no match

```

```

working out tests/t11_obf.s
going to reduce ((arg1 & arg1) | (arg1 & arg1))
reduction in reduce_AND3() (arg1 & arg1) -> arg1
reduction in reduce_AND3() (arg1 & arg1) -> arg1
reduction in reduce_OR1() (arg1 | arg1) -> arg1
going to reduce arg1
result=arg1

```

<sup>47</sup><http://www.hackersdelight.org/aha/aha.pdf>

This is weirder:

```
; do nothing (as found by aha)

;Found a 5-operation program:
; neg r1,rx
; neg r2,rx
; neg r3,r1
; or r4,rx,2
; and r5,r4,r3
; Expr: ((x | 2) & -(-(x)))

    mov rax, rdi
    neg rax
    neg rax
    or rdi, 2
    and rax, rdi
```

Rules added (I used "NEG" string to represent sign change and to be different from subtraction operation, which is just minus ("-")):

```
# (op(op X)) -> X, where both ops are NEG or NOT
def reduce_double_NEG_or_NOT (expr):
    # try each:
    for op in ["NEG", "~"]:
        m=match (expr, create_unary_expr (op, create_unary_expr (op, bind_expr("X"))))
        if m!=None:
            return dbg_print_reduced_expr ("reduce_double_NEG_or_NOT", expr, m["X"])

    # default:
    return expr # no match

...

# X & (X | ...) -> X
def reduce_AND2 (expr):
    m=match (expr, create_binary_expr ("&", create_binary_expr ("|", bind_expr ("X1"), bind_expr ("REST")), bind_expr ("X2")))
    if m!=None and match (m["X1"], m["X2"])!=None:
        return dbg_print_reduced_expr("reduce_AND2", expr, m["X1"])
    else:
        return expr # no match
```

```
going to reduce ((-(-arg1)) & (arg1 | 2))
reduction in reduce_double_NEG_or_NOT() (-(-arg1)) -> arg1
reduction in reduce_AND2() (arg1 & (arg1 | 2)) -> arg1
going to reduce arg1
result=arg1
```

I also forced *aha!* to find piece of code which adds 2 with no addition/subtraction operations allowed:

```
; arg1+2, without add/sub allowed, as found by aha:

;Found a 4-operation program:
; not r1,rx
; neg r2,r1
; not r3,r2
; neg r4,r3
; Expr: -((~(-(~(x))))))

    mov rax, rdi
    not rax
    neg rax
    not rax
    neg rax
```

Rule:

```
# (- (~X)) -> X+1
def reduce_NEG_NOT (expr):
    m=match (expr, create_unary_expr ("NEG", create_unary_expr ("~", bind_expr("X"))))
    if m==None:
        return expr # no match

    return dbg_print_reduced_expr ("reduce_NEG_NOT", expr, create_binary_expr ("+", m["X"],create_val_expr(1)))
```

```

working out tests/add_by_not_neg.s
going to reduce (-(~(-(~arg1))))
reduction in reduce_NEG_NOT() (-(~arg1)) -> (arg1 + 1)
reduction in reduce_NEG_NOT() (-(~(arg1 + 1))) -> ((arg1 + 1) + 1)
reduction in reduce_ADD3() ((arg1 + 1) + 1) -> (arg1 + 2)
going to reduce (arg1 + 2)
result=(arg1 + 2)

```

This is artifact of two's complement system of signed numbers representation. Same can be done for subtraction (just swap NEG and NOT operations).

Now let's add some fake luggage to Fahrenheit-to-Celsius example:

```

; celsius = 5 * (fahr-32) / 9
; fake luggage:
mov     rbx, 12345h
mov     rax, rdi
sub     rax, 32
; fake luggage:
add     rbx, rax
imul   rax, 5
mov     rbx, 9
idiv   rbx
; fake luggage:
sub     rdx, rax

```

It's not a problem for our decompiler, because the noise is left in RDX register, and not used at all:

```

working out tests/fahr_to_celsius_obf1.s
line=[mov     rbx, 12345h]
rcx=arg4
rsi=arg2
rbx=0x12345
rdx=arg3
rdi=arg1
rax=initial_RAX

line=[mov     rax, rdi]
rcx=arg4
rsi=arg2
rbx=0x12345
rdx=arg3
rdi=arg1
rax=arg1

line=[sub     rax, 32]
rcx=arg4
rsi=arg2
rbx=0x12345
rdx=arg3
rdi=arg1
rax=(arg1 - 32)

line=[add     rbx, rax]
rcx=arg4
rsi=arg2
rbx=(0x12345 + (arg1 - 32))
rdx=arg3
rdi=arg1
rax=(arg1 - 32)

line=[imul   rax, 5]
rcx=arg4
rsi=arg2
rbx=(0x12345 + (arg1 - 32))
rdx=arg3
rdi=arg1
rax=((arg1 - 32) * 5)

line=[mov     rbx, 9]
rcx=arg4
rsi=arg2
rbx=9
rdx=arg3
rdi=arg1
rax=((arg1 - 32) * 5)

```

```

line=[idiv      rbx]
rcx=arg4
rsi=arg2
rbx=9
rdx=(((arg1 - 32) * 5) % 9)
rdi=arg1
rax=(((arg1 - 32) * 5) / 9)

line=[sub      rdx, rax]
rcx=arg4
rsi=arg2
rbx=9
rdx=(((arg1 - 32) * 5) % 9) - (((arg1 - 32) * 5) / 9)
rdi=arg1
rax=(((arg1 - 32) * 5) / 9)

going to reduce (((arg1 - 32) * 5) / 9)
result=(((arg1 - 32) * 5) / 9)

```

We can try to pretend we affect the result with the noise:

```

; celsius = 5 * (fahr-32) / 9
; fake luggage:
mov     rbx, 12345h
mov     rax, rdi
sub     rax, 32
; fake luggage:
add     rbx, rax
imul   rax, 5
mov     rbx, 9
idiv   rbx
; fake luggage:
sub     rdx, rax
mov     rcx, rax
; OR result with garbage (result of fake luggage):
or      rcx, rdx
; the following instruction shouldn't affect result:
and     rax, rcx

```

...but in fact, it's all reduced by `reduce_AND2()` function we already saw (7.5):

```

working out tests/fahr_to_celsius_obf2.s
going to reduce (((arg1 - 32) * 5) / 9) & (((arg1 - 32) * 5) / 9) | (((arg1 - 32) * 5) % 9) - (((arg1 - 32) * 5) / 9)))
reduction in reduce_AND2() (((arg1 - 32) * 5) / 9) & (((arg1 - 32) * 5) / 9) | (((arg1 - 32) * 5) % 9) - (((arg1 - 32) * 5) / 9))) -> (((arg1 - 32) * 5) / 9)
going to reduce (((arg1 - 32) * 5) / 9)
result=(((arg1 - 32) * 5) / 9)

```

We can see that deobfuscation is in fact the same thing as optimization used in compilers. We can try this function in GCC:

```

int f(int a)
{
    return ~(~a);
};

```

Optimizing GCC 5.4 (x86) generates this:

```

f:
    mov     eax, DWORD PTR [esp+4]
    add     eax, 1
    ret

```

GCC has its own rewriting rules, some of which are, probably, close to what we use here.

## 7.6 Tests

Despite simplicity of the decompiler, it's still error-prone. We need to be sure that original expression and reduced one are equivalent to each other.



## 7.6.1 Evaluating expressions

First of all, we would just evaluate (or *run*, or *execute*) expression with random values as arguments, and then compare results.

Evaluator do arithmetical operations when possible, recursively. When any symbol is encountered, its value (randomly generated before) is taken from a table.

```
un_ops={"NEG":operator.neg,
        "~":operator.invert}

bin_ops={">>":operator.rshift,
         "<<":(lambda x, c: x<<(c&0x3f)), # operator.lshift should be here, but it doesn't handle too big counts
         "&":operator.and_,
         "|":operator.or_,
         "^":operator.xor,
         "+":operator.add,
         "-":operator.sub,
         "*":operator.mul,
         "/":operator.div,
         "%":operator.mod}

def eval_expr(e, symbols):
    t=get_expr_type (e)
    if t=="EXPR_SYMBOL":
        return symbols[get_symbol(e)]
    elif t=="EXPR_VALUE":
        return get_val (e)
    elif t=="EXPR_OP":
        if is_unary_op (get_op (e)):
            return un_ops[get_op(e)](eval_expr(get_op1(e), symbols))
        else:
            return bin_ops[get_op(e)](eval_expr(get_op1(e), symbols), eval_expr(get_op2(e), symbols))
    else:
        raise AssertionError

def do_selftest(old, new):
    for n in range(100):
        symbols={"arg1":random.getrandbits(64),
                "arg2":random.getrandbits(64),
                "arg3":random.getrandbits(64),
                "arg4":random.getrandbits(64)}
        old_result=eval_expr (old, symbols)&0xffffffff # signed->unsigned
        new_result=eval_expr (new, symbols)&0xffffffff # signed->unsigned
        if old_result!=new_result:
            print "self-test failed"
            print "initial expression: "+expr_to_string(old)
            print "reduced expression: "+expr_to_string(new)
            print "initial expression result: ", old_result
            print "reduced expression result: ", new_result
            exit(0)
```

In fact, this is very close to what LISP *EVAL* function does, or even LISP interpreter. However, not all symbols are set. If the expression is using initial values from RAX or RBX (to which symbols "initial\_RAX" and "initial\_RBX" are assigned, decompiler will stop with exception, because no random values assigned to these registers, and these symbols are absent in *symbols* dictionary.

Using this test, I've suddenly found a bug here (despite simplicity of all these reduction rules). Well, no-one protected from eye strain. Nevertheless, the test has a serious problem: some bugs can be revealed only if one of arguments is 0, or 1, or -1. Maybe there are even more special cases exists.

Mentioned above *aha!* superoptimizer tries at least these values as arguments while testing: 1, 0, -1, 0x80000000, 0x7FFFFFFF, 0x80000001, 0x7FFFFFFE, 0x01234567, 0x89ABCDEF, -2, 2, -3, 3, -64, 64, -5, -31415.

Still, you cannot be sure.

## 7.6.2 Using Z3 SMT-solver for testing

So here we will try Z3 SMT-solver. SMT-solver can *prove* that two expressions are equivalent to each other.

For example, with the help of *aha!*, I've found another weird piece of code, which does nothing:

```
; do nothing (obfuscation)
;Found a 5-operation program:
```

```

; neg r1,rx
; neg r2,r1
; sub r3,r1,3
; sub r4,r3,r1
; sub r5,r4,r3
; Expr: (((-x) - 3) - -(x)) - (-(x) - 3))

mov rax, rdi
neg rax
mov rbx, rax
; rbx=-x
mov rcx, rbx
sub rcx, 3
; rcx=-x-3
mov rax, rcx
sub rax, rbx
; rax=(-(x) - 3) - -(x)
sub rax, rcx

```

Using toy decompiler, I've found that this piece is reduced to *arg1* expression:

```

working out tests/t5_obf.s
going to reduce (((-arg1) - 3) - (-arg1)) - ((-arg1) - 3))
reduction in reduce_SUB2() ((-arg1) - 3) -> (-arg1 + 3))
reduction in reduce_SUB5() ((-arg1 + 3)) - (-arg1)) -> ((-arg1 + 3)) + arg1)
reduction in reduce_SUB2() ((-arg1) - 3) -> (-arg1 + 3))
reduction in reduce_ADD_SUB() (((-arg1 + 3)) + arg1) - (-arg1 + 3)) -> arg1
going to reduce arg1
result=arg1

```

But is it correct? I've added a function which can output expression(s) to SMTLIB-format, it's as simple as a function which converts expression to string.

And this is SMTLIB-file for Z3:

```

(assert
  (forall ((arg1 (_ BitVec 64)) (arg2 (_ BitVec 64)) (arg3 (_ BitVec 64)) (arg4 (_ BitVec 64)))
    (=
      (bvsub (bvsub (bvsub (bvneg arg1) #x0000000000000003) (bvneg arg1)) (bvsub (bvneg arg1) #
x0000000000000003))
      arg1
    )
  )
)
(check-sat)

```

In plain English terms, what we asking it to be sure, that *forall* four 64-bit arguments, two expressions are equivalent (second is just *arg1*).

The syntax maybe hard to understand, but in fact, this is very close to LISP, and arithmetical operations are named "bvsb", "bvadd", etc, because "bv" stands for *bit vector*.

While running, Z3 shows "sat", meaning "satisfiable". In other words, Z3 couldn't find counterexample for this expression.

In fact, I can rewrite this expression in the following form: *expr1 != expr2*, and we would ask Z3 to find at least one set of input arguments, for which expressions are not equal to each other:

```

(declare-const arg1 (_ BitVec 64))
(declare-const arg2 (_ BitVec 64))
(declare-const arg3 (_ BitVec 64))
(declare-const arg4 (_ BitVec 64))

(assert
  (not
    (=
      (bvsub (bvsub (bvsub (bvneg arg1) #x0000000000000003) (bvneg arg1)) (bvsub (bvneg arg1) #
x0000000000000003))
      arg1
    )
  )
)
(check-sat)

```

Z3 says "unsat", meaning, it couldn't find any such counterexample. In other words, for all possible input arguments, results of these two expressions are always equal to each other.

Nevertheless, Z3 is not omnipotent. It fails to prove equivalence of the code which performs division by multiplication. First of all, I extended it so both results will have size of 128 bit instead of 64:

```
(declare-const x (_ BitVec 64))
(assert
  (forall ((x (_ BitVec 64)))
    (=
      ((_ zero_extend 64) (bvudiv x (_ bv17 64)))
      (bvlsr (bvmul ((_ zero_extend 64) x) #x0000000000000000f0f0f0f0f0f0f1) (_ bv68 128))
    )
  )
)
(check-sat)
(get-model)
```

(*bv17* is just 64-bit number 17, etc. “bv” stands for “bit vector”, as opposed to integer value.)

Z3 works too long without any answer, and I had to interrupt it.

As Z3 developers mentioned, such expressions are hard for Z3 so far: <https://github.com/Z3Prover/z3/issues/514>.

Still, division by multiplication can be tested using previously described brute-force check.

## 7.7 My other implementations of toy decompiler

When I made attempt to write it in C++, of course, node in expression was represented using class. There is also implementation in pure C<sup>48</sup>, node is represented using structure.

Matchers in both C++ and C versions doesn't return any dictionary, but instead, `bind_value()` functions takes pointer to a variable which will contain value after successful matching. `bind_expr()` takes pointer to a pointer, which will points to the part of expression, again, in case of success. I took this idea from LLVM.

Here are two pieces of code from LLVM source code with couple of reducing rules:

```
// (X >> A) << A -> X
Value *X;
if (match(Op0, m_Exact(m_Shr(m_Value(X), m_Specific(Op1))))))
    return X;
```

( [lib/Analysis/InstructionSimplify.cpp](#) )

```
// (A | B) | C and A | (B | C) -> bswap if possible.
// (A >> B) | (C << D) and (A << B) | (B >> C) -> bswap if possible.
if (match(Op0, m_Or(m_Value(), m_Value())) ||
    match(Op1, m_Or(m_Value(), m_Value())) ||
    (match(Op0, m_LogicalShift(m_Value(), m_Value())) &&
     match(Op1, m_LogicalShift(m_Value(), m_Value())))) {
    if (Instruction *BSwap = MatchBSwap(I))
        return BSwap;
```

( [lib/Transforms/InstCombine/InstCombineAndOrXor.cpp](#) )

As you can see, my matcher tries to mimic LLVM. What I call *reduction* is called *folding* in LLVM. Both terms are popular.

I have also a blog post about LLVM obfuscator, in which LLVM matcher is mentioned: <https://yurichev.com/blog/llvm/>.

Python version of toy decompiler uses strings in place where enumerate data type is used in C version (like `OP_AND`, `OP_MUL`, etc) and symbols used in Racket version<sup>49</sup> (like `'OP_DIV`, etc). This may be seen as inefficient, nevertheless, thanks to strings interning, only address of strings are compared in Python version, not strings as a whole. So strings in Python can be seen as possible replacement for LISP symbols.

### 7.7.1 Even simpler toy decompiler

Knowledge of LISP makes you understand all these things naturally, without significant effort. But when I had no knowledge of it, but still tried to make a simple toy decompiler, I made it using usual text strings which holded expressions for each registers (and even memory).

So when MOV instruction copies value from one register to another, we just copy string. When arithmetical instruction occurred, we do string concatenation:

<sup>48</sup>[https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/toy\\_decompiler/files/C](https://github.com/dennis714/SAT_SMT_article/tree/master/toy_decompiler/files/C)

<sup>49</sup>Racket is Scheme (which is, in turn, LISP dialect) dialect. [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/toy\\_decompiler/files/Racket](https://github.com/dennis714/SAT_SMT_article/tree/master/toy_decompiler/files/Racket)

```

std::string registers[TOTAL];

...

// all 3 arguments are strings
switch (ins, op1, op2)
{
    ...
    case ADD:    registers[op1]="(" + registers[op1] + " + " + registers[op2] + ")";
                break;
    ...
    case MUL:    registers[op1]="(" + registers[op1] + " / " + registers[op2] + ")";
                break;
    ...
}

```

Now you'll have long expressions for each register, represented as strings. For reducing them, you can use plain simple regular expression matcher.

For example, for the rule  $(X*n)+(X*m) \rightarrow X*(n+m)$ , you can match (sub)string using the following regular expression:

`((.*)*(.))*+((.*)*(.))*` <sup>50</sup>. If the string is matched, you're getting 4 groups (or substrings). You then just compare 1st and 3rd using string comparison function, then you check if the 2nd and 4th are numbers, you convert them to numbers, sum them and you make new string, consisting of 1st group and sum, like this: `(" + X + "*" + (int(n) + int(m)) + ")`.

It was naive, clumsy, it was source of great embarrassment, but it worked correctly.

## 7.8 Difference between toy decompiler and commercial-grade one

Perhaps, someone, who currently reading this text, may rush into extending my source code. As an exercise, I would say, that the first step could be support of partial registers: i.e., AL, AX, EAX. This is tricky, but doable.

Another task may be support of FPU<sup>51</sup> x86 instructions (FPU stack modeling isn't a big deal).

The gap between toy decompiler and a commercial decompiler like Hex-Rays is still enormous. Several tricky problems must be solved, at least these:

- C data types: arrays, structures, pointers, etc. This problem is virtually non-existent for JVM<sup>52</sup> (Java, etc) and .NET decompilers, because type information is present in binary files.
- Basic blocks, C/C++ statements. Mike Van Emmerik in his thesis <sup>53</sup> shows how this can be tackled using SSA forms (which are also used heavily in compilers).
- Memory support, including local stack. Keep in mind pointer aliasing problem. Again, decompilers of JVM and .NET files are simpler here.

## 7.9 Further reading

There are several interesting open-source attempts to build decompiler. Both source code and these are interesting study.

- *decomp* by Jim Reuter<sup>54</sup>.
- *DCC* by Cristina Cifuentes<sup>55</sup>.

It is interesting that this decompiler supports only one type (*int*). Maybe this is a reason why DCC decompiler produces source code with *.B* extension? Read more about B typeless language (C predecessor): <https://yurichev.com/blog/typeless/>.

<sup>50</sup>This regular expression string hasn't been properly escaped, for the reason of easier readability and understanding.

<sup>51</sup>Floating-point unit

<sup>52</sup>Java Virtual Machine

<sup>53</sup>[https://yurichev.com/mirrors/vanEmmerik\\_ssa.pdf](https://yurichev.com/mirrors/vanEmmerik_ssa.pdf)

<sup>54</sup><http://www.program-transformation.org/Transform/DecompReadMe>, <http://www.program-transformation.org/Transform/DecompDecompiler>

<sup>55</sup> <http://www.program-transformation.org/Transform/DccDecompiler>, thesis: [https://yurichev.com/mirrors/DCC\\_decompilation\\_thesis.pdf](https://yurichev.com/mirrors/DCC_decompilation_thesis.pdf)

- *Boomerang* by Mike Van Emmerik, Trent Waddington et al<sup>56</sup>.

As I've said, LISP knowledge can help to understand this all much easier. Here is well-known micro-interpreter of LISP by Peter Norvig, also written in Python: <https://web.archive.org/web/20161116133448/http://www.norvig.com/lispy.html>, <https://web.archive.org/web/20160305172301/http://norvig.com/lispy2.html>.

## 7.10 The files

Python version and tests: [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/toy\\_decompiler/files](https://github.com/dennis714/SAT_SMT_article/tree/master/toy_decompiler/files).

There are also C and Racket versions, but outdated.

Keep in mind—this decompiler is still at toy level, and it was tested only on tiny test files supplied.

# 8 Symbolic execution

## 8.1 Symbolic computation

Let's first start with symbolic computation<sup>57</sup>.

Some numbers can only be represented in binary system approximately, like  $\frac{1}{3}$  and  $\pi$ . If we calculate  $\frac{1}{3} \cdot 3$  step-by-step, we may have loss of significance. We also know that  $\sin(\frac{\pi}{2}) = 1$ , but calculating this expression in usual way, we can also have some noise in result. Arbitrary-precision arithmetic<sup>58</sup> is not a solution, because these numbers cannot be stored in memory as a binary number of finite length.

How we could tackle this problem? Humans reduce such expressions using paper and pencil without any calculations. We can mimic human behaviour programmatically if we will store expression as tree and symbols like  $\pi$  will be converted into number at the very last step(s).

This is what Wolfram Mathematica<sup>59</sup> does. Let's start it and try this:

```
In[]:= x + 2*8
Out[]= 16 + x
```

Since Mathematica has no clue what  $x$  is, it's left *as is*, but  $2 \cdot 8$  can be reduced easily, both by Mathematica and by humans, so that is what has done. In some point of time in future, Mathematica's user may assign some number to  $x$  and then, Mathematica will reduce the expression even further.

Mathematica does this because it parses the expression and finds some known patterns. This is also called *term rewriting*<sup>60</sup>. In plain English language it may sounds like this: "if there is a + operator between two known numbers, replace this subexpression by a computed number which is sum of these two numbers, if possible". Just like humans do.

Mathematica also has rules like "replace  $\sin(\pi)$  by 0" and "replace  $\sin(\frac{\pi}{2})$  by 1", but as you can see,  $\pi$  must be preserved as some kind of symbol instead of a number.

So Mathematica left  $x$  as unknown value. This is, in fact, common mistake by Mathematica's users: a small typo in an input expression may lead to a huge irreducible expression with the typo left.

Another example: Mathematica left this deliberately while computing binary logarithm:

```
In[]:= Log[2, 36]
Out[]= Log[36]/Log[2]
```

Because it has a hope that at some point in future, this expression will become a subexpression in another expression and it will be reduced nicely at the very end. But if we really need a numerical answer, we can force Mathematica to calculate it:

```
In[]:= Log[2, 36] // N
Out[]= 5.16993
```

Sometimes unresolved values are desirable:

<sup>56</sup> <http://boomerang.sourceforge.net/>, <http://www.program-transformation.org/Transform/MikeVanEmmerik>, thesis: [https://yurichev.com/mirrors/vanEmmerik\\_ssa.pdf](https://yurichev.com/mirrors/vanEmmerik_ssa.pdf)

<sup>57</sup> [https://en.wikipedia.org/wiki/Symbolic\\_computation](https://en.wikipedia.org/wiki/Symbolic_computation)

<sup>58</sup> [https://en.wikipedia.org/wiki/Arbitrary-precision\\_arithmetic](https://en.wikipedia.org/wiki/Arbitrary-precision_arithmetic)

<sup>59</sup> Another well-known symbolic computation system are [Maxima](#) and [SymPy](#)

<sup>60</sup> <https://en.wikipedia.org/wiki/Rewriting>

```
In[ ]:= Union[{a, b, a, c}, {d, a, e, b}, {c, a}]
Out[ ]= {a, b, c, d, e}
```

Characters in the expression are just unresolved symbols<sup>61</sup> with no connections to numbers or other expressions, so Mathematica left them *as is*.

Another real world example is symbolic integration<sup>62</sup>, i.e., finding formula for integral by rewriting initial expression using some predefined rules. Mathematica also does it:

```
In[ ]:= Integrate[1/(x^5), x]
Out[ ]= -(1/(4 x^4))
```

Benefits of symbolic computation are obvious: it is not prone to loss of significance<sup>63</sup> and round-off errors<sup>64</sup>, but drawbacks are also obvious: you need to store expression in (possible huge) tree and process it many times. Term rewriting is also slow. All these things are extremely clumsy in comparison to a fast FPU.

“Symbolic computation” is opposed to “numerical computation”, the last one is just processing numbers step-by-step, using calculator, CPU or FPU.

Some task can be solved better by the first method, some others – by the second one.

### 8.1.1 Rational data type

Some LISP implementations can store a number as a ratio/fraction<sup>65</sup>, i.e., placing two numbers in a cell (which, in this case, is called *atom* in LISP lingo). For example, you divide 1 by 3, and the interpreter, by understanding that  $\frac{1}{3}$  is an irreducible fraction<sup>66</sup>, creates a cell with 1 and 3 numbers. Some time after, you may multiply this cell by 6, and the multiplication function inside LISP interpreter may return much better result (2 without *noise*).

Printing function in interpreter can also print something like `1 / 3` instead of floating point number.

This is sometimes called “fractional arithmetic” [see Donald E. Knuth, *The Art of Computing Programming*, 3rd ed., (1997), 4.5.1, page 330].

This is not symbolic computation in any way, but this is slightly better than storing ratios/fractions as just floating point numbers.

Drawbacks are clearly visible: you need more memory to store ratio instead of a number; and all arithmetic functions are more complex and slower, because they must handle both numbers and ratios.

Perhaps, because of drawbacks, some programming languages offers separate (*rational*) data type, as language feature, or supported by a library<sup>67</sup>: Haskell, OCaml, Perl, Ruby, Python, Smalltalk, Java, Clojure, C/C++<sup>68</sup>.

## 8.2 Symbolic execution

### 8.2.1 Swapping two values using XOR

There is a well-known (but counterintuitive) algorithm for swapping two values in two variables using XOR operation without use of any additional memory/register:

```
X=X^Y
Y=Y^X
X=X^Y
```

How it works? It would be better to construct an expression at each step of execution.

```
#!/usr/bin/env python
class Expr:
    def __init__(self,s):
        self.s=s

    def __str__(self):
```

<sup>61</sup>Symbol like in LISP

<sup>62</sup>[https://en.wikipedia.org/wiki/Symbolic\\_integration](https://en.wikipedia.org/wiki/Symbolic_integration)

<sup>63</sup>[https://en.wikipedia.org/wiki/Loss\\_of\\_significance](https://en.wikipedia.org/wiki/Loss_of_significance)

<sup>64</sup>[https://en.wikipedia.org/wiki/Round-off\\_error](https://en.wikipedia.org/wiki/Round-off_error)

<sup>65</sup>[https://en.wikipedia.org/wiki/Rational\\_data\\_type](https://en.wikipedia.org/wiki/Rational_data_type)

<sup>66</sup>[https://en.wikipedia.org/wiki/Irreducible\\_fraction](https://en.wikipedia.org/wiki/Irreducible_fraction)

<sup>67</sup>More detailed list: [https://en.wikipedia.org/wiki/Rational\\_data\\_type](https://en.wikipedia.org/wiki/Rational_data_type)

<sup>68</sup>By GNU Multiple Precision Arithmetic Library

```

        return self.s

    def __xor__(self, other):
        return Expr("(" + self.s + "^" + other.s + ")")

def XOR_swap(X, Y):
    X=X^Y
    Y=Y^X
    X=X^Y
    return X, Y

new_X, new_Y=XOR_swap(Expr("X"), Expr("Y"))
print "new_X", new_X
print "new_Y", new_Y

```

It works, because Python is dynamically typed [PL](#), so the function doesn't care what to operate on, numerical values, or on objects of Expr() class.

Here is result:

```

new_X ((X^Y)^(Y^(X^Y)))
new_Y (Y^(X^Y))

```

You can remove double variables in your mind (since XORing by a value twice will result in nothing). At new\_X we can drop two X-es and two Y-es, and single Y will left. At new\_Y we can drop two Y-es, and single X will left.

## 8.2.2 Change endianness

What does this code do?

```

mov     eax, ecx
mov     edx, ecx
shl     edx, 16
and     eax, 0000ff00H
or      eax, edx
mov     edx, ecx
and     edx, 00ff0000H
shr     ecx, 16
or      edx, ecx
shl     eax, 8
shr     edx, 8
or      eax, edx

```

In fact, many reverse engineers play shell game a lot, keeping track of what is stored where, at each point of time.



Figure 8: Hieronymus Bosch - The Conjurer

Again, we can build equivalent function which can take both numerical variables and Expr() objects. We also extend Expr() class to support many arithmetical and boolean operations. Also, Expr() methods would take both Expr() objects on input and integer values.

```
#!/usr/bin/env python
class Expr:
    def __init__(self,s):
        self.s=s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __str__(self):
        return self.s

    def __xor__(self, other):
        return Expr("(" + self.s + "^" + self.convert_to_Expr_if_int(other).s + ")")

    def __and__(self, other):
        return Expr("(" + self.s + "&" + self.convert_to_Expr_if_int(other).s + ")")

    def __or__(self, other):
        return Expr("(" + self.s + "|" + self.convert_to_Expr_if_int(other).s + ")")

    def __lshift__(self, other):
        return Expr("(" + self.s + "<<" + self.convert_to_Expr_if_int(other).s + ")")

    def __rshift__(self, other):
        return Expr("(" + self.s + ">>" + self.convert_to_Expr_if_int(other).s + ")")
```



```

# change endianness
ecx=Expr("initial_ECX") # 1st argument
eax=ecx                # mov    eax, ecx
edx=ecx                # mov    edx, ecx
edx=edx<<16           # shl    edx, 16
eax=eax&0xff00        # and    eax, 0000ff00H
eax=eax|edx           # or     eax, edx
edx=ecx                # mov    edx, ecx
edx=edx&0x00ff0000    # and    edx, 00ff0000H
ecx=ecx>>16           # shr    ecx, 16
edx=edx|ecx           # or     edx, ecx
eax=eax<<8            # shl    eax, 8
edx=edx>>8            # shr    edx, 8
eax=eax|edx           # or     eax, edx

print eax

```

I run it:

```

((((initial_ECX&65280)|(initial_ECX<<16))<<8)|(((initial_ECX&16711680)|(initial_ECX>>16))>>8))

```

Now this is something more readable, however, a bit LISP-y at first sight. In fact, this is a function which change endianness in 32-bit word.

By the way, my Toy Decompiler can do this job as well, but operates on [AST](#) instead of plain strings: [7](#).

### 8.2.3 Fast Fourier transform

I've found one of the smallest possible FFT implementations on [reddit](#):

```

#!/usr/bin/env python
from cmath import exp,pi

def FFT(X):
    n = len(X)
    w = exp(-2*pi*1j/n)
    if n > 1:
        X = FFT(X[:2]) + FFT(X[1:2])
        for k in xrange(n/2):
            xk = X[k]
            X[k] = xk + w**k*X[k+n/2]
            X[k+n/2] = xk - w**k*X[k+n/2]
    return X

print FFT([1,2,3,4,5,6,7,8])

```

Just interesting, what value has each element on output?

```

#!/usr/bin/env python
from cmath import exp,pi

class Expr:
    def __init__(self,s):
        self.s=s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __str__(self):
        return self.s

    def __add__(self, other):
        return Expr("(" + self.s + "+" + self.convert_to_Expr_if_int(other).s + ")")

    def __sub__(self, other):
        return Expr("(" + self.s + "-" + self.convert_to_Expr_if_int(other).s + ")")

    def __mul__(self, other):
        return Expr("(" + self.s + "*" + self.convert_to_Expr_if_int(other).s + ")")

    def __pow__(self, other):

```

```

    return Expr("(" + self.s + "*" + self.convert_to_Expr_if_int(other).s + ")")
def FFT(X):
    n = len(X)
    # cast complex value to string, and then to Expr
    w = Expr(str(exp(-2*pi*1j/n)))
    if n > 1:
        X = FFT(X[:2]) + FFT(X[1:2])
        for k in xrange(n/2):
            xk = X[k]
            X[k] = xk + w**k*X[k+n/2]
            X[k+n/2] = xk - w**k*X[k+n/2]
    return X
input=[Expr("input_%d" % i) for i in range(8)]
output=FFT(input)
for i in range(len(output)):
    print i, ":", output[i]

```

FFT() function left almost intact, the only thing I added: complex value is converted into string and then Expr() object is constructed.

```

0 : (((input_0+((-1-1.22464679915e-16j)**0)*input_4))+((6.12323399574e-17-1j)**0)*(input_2+((-1-1.22464679915e-16j)**0)*input_6)))+((0.707106781187-0.707106781187j)**0)*((input_1+((-1-1.22464679915e-16j)**0)*input_5))+((6.12323399574e-17-1j)**0)*(input_3+((-1-1.22464679915e-16j)**0)*input_7))))
1 : (((input_0-((-1-1.22464679915e-16j)**0)*input_4))+((6.12323399574e-17-1j)**1)*(input_2-((-1-1.22464679915e-16j)**0)*input_6)))+((0.707106781187-0.707106781187j)**1)*((input_1-((-1-1.22464679915e-16j)**0)*input_5))+((6.12323399574e-17-1j)**1)*(input_3-((-1-1.22464679915e-16j)**0)*input_7))))
2 : (((input_0+((-1-1.22464679915e-16j)**0)*input_4))-((6.12323399574e-17-1j)**0)*(input_2+((-1-1.22464679915e-16j)**0)*input_6)))+((0.707106781187-0.707106781187j)**2)*((input_1+((-1-1.22464679915e-16j)**0)*input_5))-((6.12323399574e-17-1j)**0)*(input_3+((-1-1.22464679915e-16j)**0)*input_7))))
3 : (((input_0-((-1-1.22464679915e-16j)**0)*input_4))-((6.12323399574e-17-1j)**1)*(input_2-((-1-1.22464679915e-16j)**0)*input_6)))+((0.707106781187-0.707106781187j)**3)*((input_1-((-1-1.22464679915e-16j)**0)*input_5))-((6.12323399574e-17-1j)**1)*(input_3-((-1-1.22464679915e-16j)**0)*input_7))))
4 : (((input_0+((-1-1.22464679915e-16j)**0)*input_4))+((6.12323399574e-17-1j)**0)*(input_2+((-1-1.22464679915e-16j)**0)*input_6)))-((0.707106781187-0.707106781187j)**0)*((input_1+((-1-1.22464679915e-16j)**0)*input_5))+((6.12323399574e-17-1j)**0)*(input_3+((-1-1.22464679915e-16j)**0)*input_7))))
5 : (((input_0-((-1-1.22464679915e-16j)**0)*input_4))+((6.12323399574e-17-1j)**1)*(input_2-((-1-1.22464679915e-16j)**0)*input_6)))-((0.707106781187-0.707106781187j)**1)*((input_1-((-1-1.22464679915e-16j)**0)*input_5))+((6.12323399574e-17-1j)**1)*(input_3-((-1-1.22464679915e-16j)**0)*input_7))))
6 : (((input_0+((-1-1.22464679915e-16j)**0)*input_4))-((6.12323399574e-17-1j)**0)*(input_2+((-1-1.22464679915e-16j)**0)*input_6)))-((0.707106781187-0.707106781187j)**2)*((input_1+((-1-1.22464679915e-16j)**0)*input_5))-((6.12323399574e-17-1j)**0)*(input_3+((-1-1.22464679915e-16j)**0)*input_7))))
7 : (((input_0-((-1-1.22464679915e-16j)**0)*input_4))-((6.12323399574e-17-1j)**1)*(input_2-((-1-1.22464679915e-16j)**0)*input_6)))-((0.707106781187-0.707106781187j)**3)*((input_1-((-1-1.22464679915e-16j)**0)*input_5))-((6.12323399574e-17-1j)**1)*(input_3-((-1-1.22464679915e-16j)**0)*input_7))))

```

We can see subexpressions in form like  $x^0$  and  $x^1$ . We can eliminate them, since  $x^0 = 1$  and  $x^1 = x$ . Also, we can reduce subexpressions like  $x \cdot 1$  to just  $x$ .

```

def __mul__(self, other):
    op1=self.s
    op2=self.convert_to_Expr_if_int(other).s

    if op1=="1":
        return Expr(op2)
    if op2=="1":
        return Expr(op1)

    return Expr("(" + op1 + "*" + op2 + ")")

def __pow__(self, other):
    op2=self.convert_to_Expr_if_int(other).s
    if op2=="0":
        return Expr("1")
    if op2=="1":
        return Expr(self.s)

    return Expr("(" + self.s + "*" + op2 + ")")

```

```

0 : (((input_0+input_4)+(input_2+input_6))+((input_1+input_5)+(input_3+input_7)))
1 : (((input_0-input_4)+((6.12323399574e-17-1j)*(input_2-input_6)))+(0.707106781187-0.707106781187j)*((input_1-input_5)+((6.12323399574e-17-1j)*(input_3-input_7))))
2 : (((input_0+input_4)-(input_2+input_6))+((0.707106781187-0.707106781187j)**2)*((input_1+input_5)-(input_3+input_7))))

```

```

3 : (((input_0-input_4)-((6.12323399574e-17-1j)*(input_2-input_6)))+(((0.707106781187-0.707106781187j)**3)*((
input_1-input_5)-((6.12323399574e-17-1j)*(input_3-input_7))))))
4 : (((input_0+input_4)+(input_2+input_6))-((input_1+input_5)+(input_3+input_7)))
5 : (((input_0-input_4)+((6.12323399574e-17-1j)*(input_2-input_6)))-((0.707106781187-0.707106781187j)*((input_1-
input_5)+((6.12323399574e-17-1j)*(input_3-input_7))))))
6 : (((input_0+input_4)-(input_2+input_6))-(((0.707106781187-0.707106781187j)**2)*((input_1+input_5)-(input_3+
input_7))))
7 : (((input_0-input_4)-((6.12323399574e-17-1j)*(input_2-input_6)))-(((0.707106781187-0.707106781187j)**3)*((
input_1-input_5)-((6.12323399574e-17-1j)*(input_3-input_7))))))

```

## 8.2.4 Cyclic redundancy check

I've always been wondering, which input bit affects which bit in the final CRC32 value.

From the [CRC<sup>69</sup>](http://web.archive.org/web/20161220015646/http://www.hackersdelight.org/crc.pdf) theory (good and concise introduction: <http://web.archive.org/web/20161220015646/http://www.hackersdelight.org/crc.pdf>) we know that CRC is shifting register with taps.

We will track each bit rather than byte or word, which is highly inefficient, but serves our purpose better:

```

#!/usr/bin/env python
import sys

class Expr:
    def __init__(self,s):
        self.s=s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __str__(self):
        return self.s

    def __xor__(self, other):
        return Expr("(" + self.s + "^" + self.convert_to_Expr_if_int(other).s + ")")

BYTES=1

def crc32(buf):
    #state=[Expr("init %d" % i) for i in range(32)]
    state=[Expr("1") for i in range(32)]
    for byte in buf:
        for n in range(8):
            bit=byte[n]
            to_taps=bit^state[31]
            state[31]=state[30]
            state[30]=state[29]
            state[29]=state[28]
            state[28]=state[27]
            state[27]=state[26]
            state[26]=state[25]^to_taps
            state[25]=state[24]
            state[24]=state[23]
            state[23]=state[22]^to_taps
            state[22]=state[21]^to_taps
            state[21]=state[20]
            state[20]=state[19]
            state[19]=state[18]
            state[18]=state[17]
            state[17]=state[16]
            state[16]=state[15]^to_taps
            state[15]=state[14]
            state[14]=state[13]
            state[13]=state[12]
            state[12]=state[11]^to_taps
            state[11]=state[10]^to_taps
            state[10]=state[9]^to_taps
            state[9]=state[8]
            state[8]=state[7]^to_taps
            state[7]=state[6]^to_taps

```

<sup>69</sup>Cyclic redundancy check



Now we can process this expressions somehow to get a smaller picture on what is affecting what. Let's say, if we can find "in\_2\_3" substring in expression, this means that 3rd bit of 2nd byte of input affects this expression. But even more than that: since this is XOR tree (i.e., expression consisting only of XOR operations), if some input variable is occurring twice, it's *annihilated*, since  $x \oplus x = 0$ . More than that: if a variable occurred even number of times (2, 4, 8, etc), it's annihilated, but left if it's occurred odd number of times (1, 3, 5, etc).

```

for i in range(32):
    #print "state %d=%s" % (i, state[31-i])
    sys.stdout.write ("state %02d: " % i)
    for byte in range(BYTES):
        for bit in range(8):
            s="in_%d_%d" % (byte, bit)
            if str(state[31-i]).count(s) & 1:
                sys.stdout.write ("*")
            else:
                sys.stdout.write (" ")
    sys.stdout.write ("\n")

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/symbolic/4\\_CRC/2.py](https://github.com/dennis714/SAT_SMT_article/blob/master/symbolic/4_CRC/2.py) )

Now this how each bit of 1-byte input buffer affects each bit of the final CRC32 state:

```

state 00:  *
state 01:  * *
state 02:  ** *
state 03:  ** *
state 04:  * ** *
state 05:  * ** *
state 06:    **
state 07:  *   **
state 08:  *   **
state 09:    *
state 10:   *
state 11:   *
state 12:  * *
state 13: ** *
state 14: ** *
state 15: ** *
state 16: * ***
state 17: ** ***
state 18: *** ***
state 19: *** ***
state 20:   ** **
state 21:  * ** *
state 22:   ** **
state 23:   ** **
state 24:  * * ** *
state 25: **** **
state 26: ***** **
state 27:  * *** *
state 28:  * ***
state 29: **   ***
state 30: **   **
state 31:  *   *

```

This is 8\*8=64 bits of 8-byte input buffer:

```

state 00:  * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 01:  * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 02:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 03:  *** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 04:  **** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 05:  ***** * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 06:  ** *** * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 07:  * ** *** * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 08:  * ** *** * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 09:  *** ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 10:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 11:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 12:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 13:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 14:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 15:  ** * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 16:  * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 17:  * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *

```

```

state 18: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 19: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 20: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 21: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 22: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 23: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 24: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 25: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 26: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 27: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 28: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 29: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 30: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
state 31: * * * * * * * * * * * * * * * * * * * * * * * * * * * * *

```

### 8.2.5 Linear congruential generator

This is popular PRNG<sup>70</sup> from OpenWatcom CRT<sup>71</sup> library: <https://github.com/open-watcom/open-watcom-v2/blob/d468b609ba6ca61eeddad80dd2485e3256fc5261/bld/clib/math/c/rand.c>.

What expression it generates on each step?

```

#!/usr/bin/env python
class Expr:
    def __init__(self,s):
        self.s=s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __str__(self):
        return self.s

    def __xor__(self, other):
        return Expr("(" + self.s + "^" + self.convert_to_Expr_if_int(other).s + ")")

    def __mul__(self, other):
        return Expr("(" + self.s + "*" + self.convert_to_Expr_if_int(other).s + ")")

    def __add__(self, other):
        return Expr("(" + self.s + "+" + self.convert_to_Expr_if_int(other).s + ")")

    def __and__(self, other):
        return Expr("(" + self.s + "&" + self.convert_to_Expr_if_int(other).s + ")")

    def __rshift__(self, other):
        return Expr("(" + self.s + ">>" + self.convert_to_Expr_if_int(other).s + ")")

seed=Expr("initial_seed")

def rand():
    global seed
    seed=seed*1103515245+12345
    return (seed>>16) & 0x7fff

for i in range(10):
    print i, ":", rand()

```

```

0 : (((((initial_seed*1103515245)+12345)>>16)&32767)
1 : ((((((initial_seed*1103515245)+12345)*1103515245)+12345)>>16)&32767)
2 : (((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
3 : ((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
4 : (((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
5 : ((((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)

```

<sup>70</sup>Pseudorandom number generator

<sup>71</sup>C runtime library

```

6 : (((((((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
7 : (((((((((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
8 : (((((((((((((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
9 : (((((((((((((((((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)

```

Now if we once got several values from this PRNG, like 4583, 16304, 14440, 32315, 28670, 12568..., how would we recover the initial seed? The problem in fact is solving a system of equations:

```

((((initial_seed*1103515245)+12345)>>16)&32767)==4583
((((((((initial_seed*1103515245)+12345)*1103515245)+12345)>>16)&32767)==16304
((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)==14440
((((((((((((((((initial_seed*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)*1103515245)+12345)>>16)&32767)
==32315

```

As it turns out, Z3 can solve this system correctly using only two equations:

```

#!/usr/bin/env python
from z3 import *

s=Solver()

x=BitVec("x",32)

a=1103515245
c=12345
s.add((((x*a)+c)>>16)&32767==4583)
s.add((((((((x*a)+c)*a)+c)>>16)&32767==16304)
#s.add((((((((((((x*a)+c)*a)+c)*a)+c)>>16)&32767==14440)
#s.add((((((((((((((((x*a)+c)*a)+c)*a)+c)*a)+c)>>16)&32767==32315)

s.check()
print s.model()

```

```
[x = 11223344]
```

(Though, it takes  $\approx 20$  seconds on my ancient Intel Atom netbook.)

## 8.2.6 Path constraint

How to get weekday from UNIX timestamp?

```

#!/usr/bin/env python
input=...
SECS_DAY=24*60*60
dayno = input / SECS_DAY
wday = (dayno + 4) % 7
if wday==5:
    print "Thanks God, it's Friday!"

```

Let's say, we should find a way to run the block with print() call in it. What input value should be? First, let's build expression of *wday* variable:

```

#!/usr/bin/env python
class Expr:
    def __init__(self,s):
        self.s=s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __str__(self):
        return self.s

```

```

def __div__(self, other):
    return Expr("(" + self.s + "/" + self.convert_to_Expr_if_int(other).s + ")")

def __mod__(self, other):
    return Expr("(" + self.s + "%" + self.convert_to_Expr_if_int(other).s + ")")

def __add__(self, other):
    return Expr("(" + self.s + "+" + self.convert_to_Expr_if_int(other).s + ")")

input=Expr("input")
SECS_DAY=24*60*60
dayno = input / SECS_DAY
wday = (dayno + 4) % 7
print wday
if wday==5:
    print "Thanks God, it's Friday!"

```

```
((input/86400)+4)%7)
```

In order to execute the block, we should solve this equation:  $((\frac{input}{86400} + 4) \equiv 5 \pmod{7})$ .  
So far, this is easy task for Z3:

```

#!/usr/bin/env python
from z3 import *

s=Solver()

x=Int("x")

s.add(((x/86400)+4)%7==5)

s.check()
print s.model()

```

```
[x = 86438]
```

This is indeed correct UNIX timestamp for Friday:

```
% date --date='@86438'
Fri Jan 2 03:00:38 MSK 1970
```

Though the date back in year 1970, but it's still correct!

This is also called “path constraint”, i.e., what constraint must be satisfied to execute specific block? Several tools has “path” in their names, like “pathgrind”, [Symbolic PathFinder](#), CodeSurfer Path Inspector, etc.

Like the shell game, this task is also often encounters in practice. You can see that something dangerous can be executed inside some basic block and you're trying to deduce, what input values can cause execution of it. It may be buffer overflow, etc. Such input values are sometimes also called “inputs of death”.

Many crackmes are solved in this way, all you need is find a path into block which prints “key is correct” or something like that.

We can extend this tiny example:

```

input=...
SECS_DAY=24*60*60
dayno = input / SECS_DAY
wday = (dayno + 4) % 7
print wday
if wday==5:
    print "Thanks God, it's Friday!"
else:
    print "Got to wait a little"

```

Now we have two blocks: for the first we should solve this equation:  $((\frac{input}{86400} + 4) \equiv 5 \pmod{7})$ . But for the second we should solve inverted equation:  $((\frac{input}{86400} + 4) \not\equiv 5 \pmod{7})$ . By solving these equations, we will find two paths into both blocks.

KLEE (or similar tool) tries to find path to each [basic] block and produces “ideal” unit test. Hence, KLEE can find a path into the block which crashes everything, or reporting about correctness of the input key/license, etc. Surprisingly, KLEE can find backdoors in the very same manner.



KLEE is also called “KLEE Symbolic Virtual Machine” – by that its creators mean that the KLEE is VM<sup>72</sup> which executes a code symbolically rather than numerically (like usual CPU).

### 8.2.7 Division by zero

If division by zero is unwrapped by sanitizing check, and exception isn't caught, it can crash process.

Let's calculate simple expression  $\frac{x}{2y+4z-12}$ . We can add a warning into `__div__` method:

```
#!/usr/bin/env python
class Expr:
    def __init__(self,s):
        self.s=s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __str__(self):
        return self.s

    def __mul__(self, other):
        return Expr("(" + self.s + "*" + self.convert_to_Expr_if_int(other).s + ")")

    def __div__(self, other):
        op2=self.convert_to_Expr_if_int(other).s
        print "warning: division by zero if "+op2+"==0"
        return Expr("(" + self.s + "/" + op2 + ")")

    def __add__(self, other):
        return Expr("(" + self.s + "+" + self.convert_to_Expr_if_int(other).s + ")")

    def __sub__(self, other):
        return Expr("(" + self.s + "-" + self.convert_to_Expr_if_int(other).s + ")")

"""
    x
-----
2y + 4z - 12
"""

def f(x, y, z):
    return x/(y*2 + z*4 - 12)

print f(Expr("x"), Expr("y"), Expr("z"))
```

...so it will report about dangerous states and conditions:

```
warning: division by zero if ((y*2)+(z*4))-12==0
(x/(((y*2)+(z*4))-12))
```

This equation is easy to solve, let's try Wolfram Mathematica this time:

```
In[]:= FindInstance[{(y*2 + z*4) - 12 == 0}, {y, z}, Integers]
Out[]:= {{y -> 0, z -> 3}}
```

These values for  $y$  and  $z$  can also be called “inputs of death”.

### 8.2.8 Merge sort

How merge sort works? I have cypasted Python code from rosettacode.com almost intact:

```
#!/usr/bin/env python
class Expr:
    def __init__(self,s,i):
        self.s=s
        self.i=i
```

<sup>72</sup>Virtual Machine

```

def __str__(self):
    # return both symbolic and integer:
    return self.s+" (" + str(self.i)+")"

def __le__(self, other):
    # compare only integer parts:
    return self.i <= other.i

# copypasted from http://rosettacode.org/wiki/Sorting_algorithms/Merge_sort#Python
def merge(left, right):
    result = []
    left_idx, right_idx = 0, 0
    while left_idx < len(left) and right_idx < len(right):
        # change the direction of this comparison to change the direction of the sort
        if left[left_idx] <= right[right_idx]:
            result.append(left[left_idx])
            left_idx += 1
        else:
            result.append(right[right_idx])
            right_idx += 1

    if left_idx < len(left):
        result.extend(left[left_idx:])
    if right_idx < len(right):
        result.extend(right[right_idx:])
    return result

def tabs (t):
    return "\t"*t

def merge_sort(m, indent=0):
    print tabs(indent)+"merge_sort() begin. input:"
    for i in m:
        print tabs(indent)+str(i)

    if len(m) <= 1:
        print tabs(indent)+"merge_sort() end. returning single element"
        return m

    middle = len(m) // 2
    left = m[:middle]
    right = m[middle:]

    left = merge_sort(left, indent+1)
    right = merge_sort(right, indent+1)
    rt=list(merge(left, right))
    print tabs(indent)+"merge_sort() end. returning:"
    for i in rt:
        print tabs(indent)+str(i)
    return rt

# input buffer has both symbolic and numerical values:
input=[Expr("input1",22), Expr("input2",7), Expr("input3",2), Expr("input4",1), Expr("input5",8), Expr("input6",4)]
merge_sort(input)

```

But here is a function which compares elements. Obviously, it wouldn't work correctly without it.

So we can track both expression for each element and numerical value. Both will be printed finally. But whenever values are to be compared, only numerical parts will be used.

Result:

```

merge_sort() begin. input:
input1 (22)
input2 (7)
input3 (2)
input4 (1)
input5 (8)
input6 (4)
merge_sort() begin. input:
input1 (22)
input2 (7)
input3 (2)
merge_sort() begin. input:
input1 (22)
merge_sort() end. returning single element

```

```

merge_sort() begin. input:
input2 (7)
input3 (2)
  merge_sort() begin. input:
  input2 (7)
  merge_sort() end. returning single element
  merge_sort() begin. input:
  input3 (2)
  merge_sort() end. returning single element
merge_sort() end. returning:
input3 (2)
input2 (7)
input1 (22)
merge_sort() begin. input:
input4 (1)
input5 (8)
input6 (4)
  merge_sort() begin. input:
  input4 (1)
  merge_sort() end. returning single element
  merge_sort() begin. input:
  input5 (8)
  input6 (4)
    merge_sort() begin. input:
    input5 (8)
    merge_sort() end. returning single element
    merge_sort() begin. input:
    input6 (4)
    merge_sort() end. returning single element
  merge_sort() end. returning:
  input6 (4)
  input5 (8)
merge_sort() end. returning:
input4 (1)
input6 (4)
input5 (8)
merge_sort() end. returning:
input4 (1)
input3 (2)
input6 (4)
input2 (7)
input5 (8)
input1 (22)

```

### 8.2.9 Extending Expr class

This is somewhat senseless, nevertheless, it's easy task to extend my Expr class to support [AST](#) instead of plain strings. It's also possible to add folding steps (like I demonstrated in [Toy Decompiler: 7](#)). Maybe someone will want to do this as an exercise. By the way, the toy decompiler can be used as simple symbolic engine as well, just feed all the instructions to it and it will track contents of each register.

### 8.2.10 Conclusion

For the sake of demonstration, I made things as simple as possible. But reality is always harsh and inconvenient, so all this shouldn't be taken as a silver bullet.

The files used in this part: [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/symbolic](https://github.com/dennis714/SAT_SMT_article/tree/master/symbolic).

## 8.3 Further reading

James C. King — Symbolic Execution and Program Testing <sup>73</sup>

<sup>73</sup><https://yurichev.com/mirrors/king76symbolicexecution.pdf>

## 9 KLEE

### 9.1 Installation

KLEE building from source is tricky. Easiest way to use KLEE is to install docker<sup>74</sup> and then to run KLEE docker image<sup>75</sup>.

### 9.2 School-level equation

Let's revisit school-level system of equations from (5.2).

We will force KLEE to find a path, where all the constraints are satisfied:

```
int main()
{
    int circle, square, triangle;

    klee_make_symbolic(&circle, sizeof circle, "circle");
    klee_make_symbolic(&square, sizeof square, "square");
    klee_make_symbolic(&triangle, sizeof triangle, "triangle");

    if (circle+circle!=10) return 0;
    if (circle*square+square!=12) return 0;
    if (circle*square-triangle*circle!=circle) return 0;

    // all constraints should be satisfied at this point
    // force KLEE to produce .err file:
    klee_assert(0);
};
```

```
% clang -emit-llvm -c -g klee_eq.c
...

% klee klee_eq.bc
KLEE: output directory is "/home/klee/klee-out-93"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_eq.c:18: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 32
KLEE: done: completed paths = 1
KLEE: done: generated tests = 1
```

Let's find out, where `klee_assert()` has been triggered:

```
% ls klee-last | grep err
test000001.external.err

% ktest-tool --write-ints klee-last/test000001.ktest
ktest file : 'klee-last/test000001.ktest'
args      : ['klee_eq.bc']
num objects: 3
object 0: name: b'circle'
object 0: size: 4
object 0: data: 5
object 1: name: b'square'
object 1: size: 4
object 1: data: 2
object 2: name: b'triangle'
object 2: size: 4
object 2: data: 1
```

This is indeed correct solution to the system of equations.

KLEE has *intrinsic* `klee_assume()` which tells KLEE to cut path if some constraint is not satisfied. So we can rewrite our example in such cleaner way:

```
int main()
{
```

<sup>74</sup><https://docs.docker.com/engine/installation/linux/ubuntu/linux/>

<sup>75</sup><http://klee.github.io/docker/>

```

int circle, square, triangle;

klee_make_symbolic(&circle, sizeof circle, "circle");
klee_make_symbolic(&square, sizeof square, "square");
klee_make_symbolic(&triangle, sizeof triangle, "triangle");

klee_assume (circle+circle==10);
klee_assume (circle*square+square==12);
klee_assume (circle*square-triangle*circle==circle);

// all constraints should be satisfied at this point
// force KLEE to produce .err file:
klee_assert(0);
};

```

### 9.3 Zebra puzzle

Let's revisit zebra puzzle from (5.4).

We just define all variables and add constraints:

```

int main()
{
    int Yellow, Blue, Red, Ivory, Green;
    int Norwegian, Ukrainian, Englishman, Spaniard, Japanese;
    int Water, Tea, Milk, OrangeJuice, Coffee;
    int Kools, Chesterfield, OldGold, LuckyStrike, Parliament;
    int Fox, Horse, Snails, Dog, Zebra;

    klee_make_symbolic(&Yellow, sizeof(int), "Yellow");
    klee_make_symbolic(&Blue, sizeof(int), "Blue");
    klee_make_symbolic(&Red, sizeof(int), "Red");
    klee_make_symbolic(&Ivory, sizeof(int), "Ivory");
    klee_make_symbolic(&Green, sizeof(int), "Green");

    klee_make_symbolic(&Norwegian, sizeof(int), "Norwegian");
    klee_make_symbolic(&Ukrainian, sizeof(int), "Ukrainian");
    klee_make_symbolic(&Englishman, sizeof(int), "Englishman");
    klee_make_symbolic(&Spaniard, sizeof(int), "Spaniard");
    klee_make_symbolic(&Japanese, sizeof(int), "Japanese");

    klee_make_symbolic(&Water, sizeof(int), "Water");
    klee_make_symbolic(&Tea, sizeof(int), "Tea");
    klee_make_symbolic(&Milk, sizeof(int), "Milk");
    klee_make_symbolic(&OrangeJuice, sizeof(int), "OrangeJuice");
    klee_make_symbolic(&Coffee, sizeof(int), "Coffee");

    klee_make_symbolic(&Kools, sizeof(int), "Kools");
    klee_make_symbolic(&Chesterfield, sizeof(int), "Chesterfield");
    klee_make_symbolic(&OldGold, sizeof(int), "OldGold");
    klee_make_symbolic(&LuckyStrike, sizeof(int), "LuckyStrike");
    klee_make_symbolic(&Parliament, sizeof(int), "Parliament");

    klee_make_symbolic(&Fox, sizeof(int), "Fox");
    klee_make_symbolic(&Horse, sizeof(int), "Horse");
    klee_make_symbolic(&Snails, sizeof(int), "Snails");
    klee_make_symbolic(&Dog, sizeof(int), "Dog");
    klee_make_symbolic(&Zebra, sizeof(int), "Zebra");

    // limits.
    if (Yellow<1 || Yellow>5) return 0;
    if (Blue<1 || Blue>5) return 0;
    if (Red<1 || Red>5) return 0;
    if (Ivory<1 || Ivory>5) return 0;
    if (Green<1 || Green>5) return 0;

    if (Norwegian<1 || Norwegian>5) return 0;
    if (Ukrainian<1 || Ukrainian>5) return 0;
    if (Englishman<1 || Englishman>5) return 0;
    if (Spaniard<1 || Spaniard>5) return 0;
    if (Japanese<1 || Japanese>5) return 0;

    if (Water<1 || Water>5) return 0;
    if (Tea<1 || Tea>5) return 0;

```

```

if (Milk<1 || Milk>5) return 0;
if (OrangeJuice<1 || OrangeJuice>5) return 0;
if (Coffee<1 || Coffee>5) return 0;

if (Kools<1 || Kools>5) return 0;
if (Chesterfield<1 || Chesterfield>5) return 0;
if (OldGold<1 || OldGold>5) return 0;
if (LuckyStrike<1 || LuckyStrike>5) return 0;
if (Parliament<1 || Parliament>5) return 0;

if (Fox<1 || Fox>5) return 0;
if (Horse<1 || Horse>5) return 0;
if (Snails<1 || Snails>5) return 0;
if (Dog<1 || Dog>5) return 0;
if (Zebra<1 || Zebra>5) return 0;

// colors are distinct for all 5 houses:
if (((1<<Yellow) | (1<<Blue) | (1<<Red) | (1<<Ivory) | (1<<Green))!=0x3E) return 0; // 111110

// all nationalities are living in different houses:
if (((1<<Norwegian) | (1<<Ukrainian) | (1<<Englishman) | (1<<Spaniard) | (1<<Japanese))!=0x3E) return 0;
// 111110

// so are beverages:
if (((1<<Water) | (1<<Tea) | (1<<Milk) | (1<<OrangeJuice) | (1<<Coffee))!=0x3E) return 0; // 111110

// so are cigarettes:
if (((1<<Kools) | (1<<Chesterfield) | (1<<OldGold) | (1<<LuckyStrike) | (1<<Parliament))!=0x3E) return
0; // 111110

// so are pets:
if (((1<<Fox) | (1<<Horse) | (1<<Snails) | (1<<Dog) | (1<<Zebra))!=0x3E) return 0; // 111110

// main constraints of the puzzle:

// 2.The Englishman lives in the red house.
if (Englishman!=Red) return 0;

// 3.The Spaniard owns the dog.
if (Spaniard!=Dog) return 0;

// 4.Coffee is drunk in the green house.
if (Coffee!=Green) return 0;

// 5.The Ukrainian drinks tea.
if (Ukrainian!=Tea) return 0;

// 6.The green house is immediately to the right of the ivory house.
if (Green!=Ivory+1) return 0;

// 7.The Old Gold smoker owns snails.
if (OldGold!=Snails) return 0;

// 8.Kools are smoked in the yellow house.
if (Kools!=Yellow) return 0;

// 9.Milk is drunk in the middle house.
if (Milk!=3) return 0; // i.e., 3rd house

// 10.The Norwegian lives in the first house.
if (Norwegian!=1) return 0;

// 11.The man who smokes Chesterfields lives in the house next to the man with the fox.
if (Chesterfield!=Fox+1 && Chesterfield!=Fox-1) return 0; // left or right

// 12.Kools are smoked in the house next to the house where the horse is kept.
if (Kools!=Horse+1 && Kools!=Horse-1) return 0; // left or right

// 13.The Lucky Strike smoker drinks orange juice.
if (LuckyStrike!=OrangeJuice) return 0;

// 14.The Japanese smokes Parliaments.
if (Japanese!=Parliament) return 0;

// 15.The Norwegian lives next to the blue house.
if (Norwegian!=Blue+1 && Norwegian!=Blue-1) return 0; // left or right

```

```

// all constraints are satisfied at this point
// force KLEE to produce .err file:
klee_assert(0);

return 0;
};

```

I force KLEE to find distinct values for colors, nationalities, cigarettes, etc, in the same way as I did for Sudoku earlier (5.5).

Let's run it:

```

% clang -emit-llvm -c -g klee_zebra1.c
...

% klee klee_zebra1.bc
KLEE: output directory is "/home/klee/klee-out-97"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_zebra1.c:130: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 761
KLEE: done: completed paths = 55
KLEE: done: generated tests = 55

```

It works for  $\approx 7$  seconds on my Intel Core i3-3110M 2.4GHz notebook. Let's find out path, where `klee_assert()` has been executed:

```

% ls klee-last | grep err
test000051.external.err

% ktest-tool --write-ints klee-last/test000051.ktest | less

ktest file : 'klee-last/test000051.ktest'
args       : ['klee_zebra1.bc']
num objects: 25
object  0: name: b'Yellow'
object  0: size: 4
object  0: data: 1
object  1: name: b'Blue'
object  1: size: 4
object  1: data: 2
object  2: name: b'Red'
object  2: size: 4
object  2: data: 3
object  3: name: b'Ivory'
object  3: size: 4
object  3: data: 4
...
object 21: name: b'Horse'
object 21: size: 4
object 21: data: 2
object 22: name: b'Snails'
object 22: size: 4
object 22: data: 3
object 23: name: b'Dog'
object 23: size: 4
object 23: data: 4
object 24: name: b'Zebra'
object 24: size: 4
object 24: data: 5

```

This is indeed correct solution.

`klee_assume()` also can be used this time:

```

int main()
{
    int Yellow, Blue, Red, Ivory, Green;
    int Norwegian, Ukrainian, Englishman, Spaniard, Japanese;
    int Water, Tea, Milk, OrangeJuice, Coffee;
    int Kools, Chesterfield, OldGold, LuckyStrike, Parliament;

```

```

int Fox, Horse, Snails, Dog, Zebra;

klee_make_symbolic(&Yellow, sizeof(int), "Yellow");
klee_make_symbolic(&Blue, sizeof(int), "Blue");
klee_make_symbolic(&Red, sizeof(int), "Red");
klee_make_symbolic(&Ivory, sizeof(int), "Ivory");
klee_make_symbolic(&Green, sizeof(int), "Green");

klee_make_symbolic(&Norwegian, sizeof(int), "Norwegian");
klee_make_symbolic(&Ukrainian, sizeof(int), "Ukrainian");
klee_make_symbolic(&Englishman, sizeof(int), "Englishman");
klee_make_symbolic(&Spaniard, sizeof(int), "Spaniard");
klee_make_symbolic(&Japanese, sizeof(int), "Japanese");

klee_make_symbolic(&Water, sizeof(int), "Water");
klee_make_symbolic(&Tea, sizeof(int), "Tea");
klee_make_symbolic(&Milk, sizeof(int), "Milk");
klee_make_symbolic(&OrangeJuice, sizeof(int), "OrangeJuice");
klee_make_symbolic(&Coffee, sizeof(int), "Coffee");

klee_make_symbolic(&Kools, sizeof(int), "Kools");
klee_make_symbolic(&Chesterfield, sizeof(int), "Chesterfield");
klee_make_symbolic(&OldGold, sizeof(int), "OldGold");
klee_make_symbolic(&LuckyStrike, sizeof(int), "LuckyStrike");
klee_make_symbolic(&Parliament, sizeof(int), "Parliament");

klee_make_symbolic(&Fox, sizeof(int), "Fox");
klee_make_symbolic(&Horse, sizeof(int), "Horse");
klee_make_symbolic(&Snails, sizeof(int), "Snails");
klee_make_symbolic(&Dog, sizeof(int), "Dog");
klee_make_symbolic(&Zebra, sizeof(int), "Zebra");

// limits.
klee_assume (Yellow>=1 && Yellow<=5);
klee_assume (Blue>=1 && Blue<=5);
klee_assume (Red>=1 && Red<=5);
klee_assume (Ivory>=1 && Ivory<=5);
klee_assume (Green>=1 && Green<=5);

klee_assume (Norwegian>=1 && Norwegian<=5);
klee_assume (Ukrainian>=1 && Ukrainian<=5);
klee_assume (Englishman>=1 && Englishman<=5);
klee_assume (Spaniard>=1 && Spaniard<=5);
klee_assume (Japanese>=1 && Japanese<=5);

klee_assume (Water>=1 && Water<=5);
klee_assume (Tea>=1 && Tea<=5);
klee_assume (Milk>=1 && Milk<=5);
klee_assume (OrangeJuice>=1 && OrangeJuice<=5);
klee_assume (Coffee>=1 && Coffee<=5);

klee_assume (Kools>=1 && Kools<=5);
klee_assume (Chesterfield>=1 && Chesterfield<=5);
klee_assume (OldGold>=1 && OldGold<=5);
klee_assume (LuckyStrike>=1 && LuckyStrike<=5);
klee_assume (Parliament>=1 && Parliament<=5);

klee_assume (Fox>=1 && Fox<=5);
klee_assume (Horse>=1 && Horse<=5);
klee_assume (Snails>=1 && Snails<=5);
klee_assume (Dog>=1 && Dog<=5);
klee_assume (Zebra>=1 && Zebra<=5);

// colors are distinct for all 5 houses:
klee_assume (((1<<Yellow) | (1<<Blue) | (1<<Red) | (1<<Ivory) | (1<<Green))==0x3E); // 111110

// all nationalities are living in different houses:
klee_assume (((1<<Norwegian) | (1<<Ukrainian) | (1<<Englishman) | (1<<Spaniard) | (1<<Japanese))==0x3E);
// 111110

// so are beverages:
klee_assume (((1<<Water) | (1<<Tea) | (1<<Milk) | (1<<OrangeJuice) | (1<<Coffee))==0x3E); // 111110

// so are cigarettes:
klee_assume (((1<<Kools) | (1<<Chesterfield) | (1<<OldGold) | (1<<LuckyStrike) | (1<<Parliament))==0x3E)
; // 111110

```



```

// so are pets:
klee_assume ((1<<Fox) | (1<<Horse) | (1<<Snails) | (1<<Dog) | (1<<Zebra))==0x3E); // 111110

// main constraints of the puzzle:

// 2.The Englishman lives in the red house.
klee_assume (Englishman==Red);

// 3.The Spaniard owns the dog.
klee_assume (Spaniard==Dog);

// 4.Coffee is drunk in the green house.
klee_assume (Coffee==Green);

// 5.The Ukrainian drinks tea.
klee_assume (Ukrainian==Tea);

// 6.The green house is immediately to the right of the ivory house.
klee_assume (Green==Ivory+1);

// 7.The Old Gold smoker owns snails.
klee_assume (OldGold==Snails);

// 8.Kools are smoked in the yellow house.
klee_assume (Kools==Yellow);

// 9.Milk is drunk in the middle house.
klee_assume (Milk==3); // i.e., 3rd house

// 10.The Norwegian lives in the first house.
klee_assume (Norwegian==1);

// 11.The man who smokes Chesterfields lives in the house next to the man with the fox.
klee_assume (Chesterfield==Fox+1 || Chesterfield==Fox-1); // left or right

// 12.Kools are smoked in the house next to the house where the horse is kept.
klee_assume (Kools==Horse+1 || Kools==Horse-1); // left or right

// 13.The Lucky Strike smoker drinks orange juice.
klee_assume (LuckyStrike==OrangeJuice);

// 14.The Japanese smokes Parliaments.
klee_assume (Japanese==Parliament);

// 15.The Norwegian lives next to the blue house.
klee_assume (Norwegian==Blue+1 || Norwegian==Blue-1); // left or right

// all constraints are satisfied at this point
// force KLEE to produce .err file:
klee_assert(0);
};

```

...and this version works slightly faster ( $\approx 5$  seconds), maybe because KLEE is aware of this *intrinsic* and handles it in a special way?

## 9.4 Sudoku

I've also rewritten Sudoku example (5.5) for KLEE:

```

1 #include <stdint.h>
2
3 /*
4 coordinates:
5 -----
6 00 01 02 | 03 04 05 | 06 07 08
7 10 11 12 | 13 14 15 | 16 17 18
8 20 21 22 | 23 24 25 | 26 27 28
9 -----
10 30 31 32 | 33 34 35 | 36 37 38
11 40 41 42 | 43 44 45 | 46 47 48
12 50 51 52 | 53 54 55 | 56 57 58
13 -----
14 60 61 62 | 63 64 65 | 66 67 68

```

```

15 70 71 72 | 73 74 75 | 76 77 78
16 80 81 82 | 83 84 85 | 86 87 88
17 -----
18 */
19
20 uint8_t cells[9][9];
21
22 // http://www.norvig.com/sudoku.html
23 // http://www.mirror.co.uk/news/weird-news/worlds-hardest-sudoku-can-you-242294
24 char *puzzle="..53.....8.....2..7..1.5..4....53...1..7...6..32...8..6.5....9..4....3.....97..";
25
26 int main()
27 {
28     klee_make_symbolic(cells, sizeof cells, "cells");
29
30     // process text line:
31     for (int row=0; row<9; row++)
32         for (int column=0; column<9; column++)
33             {
34                 char c=puzzle[row*9 + column];
35                 if (c!='.')
36                     {
37                         if (cells[row][column]!=c-'0') return 0;
38                     }
39                 else
40                     {
41                         // limit cells values to 1..9:
42                         if (cells[row][column]<1) return 0;
43                         if (cells[row][column]>9) return 0;
44                     }
45             };
46
47     // for all 9 rows
48     for (int row=0; row<9; row++)
49     {
50
51         if (((1<<cells[row][0]) |
52             (1<<cells[row][1]) |
53             (1<<cells[row][2]) |
54             (1<<cells[row][3]) |
55             (1<<cells[row][4]) |
56             (1<<cells[row][5]) |
57             (1<<cells[row][6]) |
58             (1<<cells[row][7]) |
59             (1<<cells[row][8]))!=0x3FE ) return 0; // 11 1111 1110
60     };
61
62     // for all 9 columns
63     for (int c=0; c<9; c++)
64     {
65         if (((1<<cells[0][c]) |
66             (1<<cells[1][c]) |
67             (1<<cells[2][c]) |
68             (1<<cells[3][c]) |
69             (1<<cells[4][c]) |
70             (1<<cells[5][c]) |
71             (1<<cells[6][c]) |
72             (1<<cells[7][c]) |
73             (1<<cells[8][c]))!=0x3FE ) return 0; // 11 1111 1110
74     };
75
76     // enumerate all 9 squares
77     for (int r=0; r<9; r+=3)
78         for (int c=0; c<9; c+=3)
79             {
80                 // add constraints for each 3*3 square:
81                 if ((1<<cells[r+0][c+0] |
82                     1<<cells[r+0][c+1] |
83                     1<<cells[r+0][c+2] |
84                     1<<cells[r+1][c+0] |
85                     1<<cells[r+1][c+1] |
86                     1<<cells[r+1][c+2] |
87                     1<<cells[r+2][c+0] |
88                     1<<cells[r+2][c+1] |
89                     1<<cells[r+2][c+2]))!=0x3FE ) return 0; // 11 1111 1110
90             };

```

```

91 // at this point, all constraints must be satisfied
92 klee_assert(0);
93 };
94

```

Let's run it:

```

% clang -emit-llvm -c -g klee_sudoku_or1.c
...

\$ time klee klee_sudoku_or1.bc
KLEE: output directory is "/home/klee/klee-out-98"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_sudoku_or1.c:93: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 7512
KLEE: done: completed paths = 161
KLEE: done: generated tests = 161

real    3m44.111s
user    3m43.319s
sys     0m0.951s

```

Now this is really slower (on my Intel Core i3-3110M 2.4GHz notebook) in comparison to Z3Py solution (5.5).

But the answer is correct:

```

% ls klee-last | grep err
test000161.external.err

% ktest-tool --write-ints klee-last/test000161.ktest
ktest file : 'klee-last/test000161.ktest'
args       : ['klee_sudoku_or1.bc']
num objects: 1
object 0: name: b'cells'
object 0: size: 81
object 0: data: b'\x01\x04\x05\x03\x02\x07\x06\t\x08\x08\x03\t\x06\x05\x04\x01\x02\x07\x06\x07\x02\t\x01\x08\x05\x04\x03\x04\t\x06\x01\x08\x05\x03\x07\x02\x02\x01\x08\x04\x07\x03\t\x05\x06\x07\x05\x03\x02\t\x06\x04\x08\x01\x03\x06\x07\x05\x04\x02\x08\x01\t\t\x08\x04\x07\x06\x01\x02\x03\x05\x05\x02\x01\x08\x03\t\x07\x06\x04'

```

Character `\t` has code of 9 in C/C++, and KLEE prints byte array as a C/C++ string, so it shows some values in such way. We can just keep in mind that there is 9 at the each place where we see `\t`. The solution, while not properly formatted, correct indeed.

By the way, at lines 42 and 43 you may see how we tell to KLEE that all array elements must be within some limits. If we comment these lines out, we've got this:

```

% time klee klee_sudoku_or1.bc
KLEE: output directory is "/home/klee/klee-out-100"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: ERROR: /home/klee/klee_sudoku_or1.c:51: overshift error
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_sudoku_or1.c:51: overshift error
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_sudoku_or1.c:51: overshift error
KLEE: NOTE: now ignoring this error at this location
...

```

KLEE warns us that shift value at line 51 is too big. Indeed, KLEE may try all byte values up to 255 (0xFF), which are pointless to use there, and may be a symptom of error or bug, so KLEE warns about it.

Now let's use `klee_assume()` again:

```

#include <stdint.h>

/*
coordinates:
-----
00 01 02 | 03 04 05 | 06 07 08
10 11 12 | 13 14 15 | 16 17 18
20 21 22 | 23 24 25 | 26 27 28
-----

```

```

30 31 32 | 33 34 35 | 36 37 38
40 41 42 | 43 44 45 | 46 47 48
50 51 52 | 53 54 55 | 56 57 58
-----
60 61 62 | 63 64 65 | 66 67 68
70 71 72 | 73 74 75 | 76 77 78
80 81 82 | 83 84 85 | 86 87 88
-----
*/

uint8_t cells[9][9];

// http://www.norvig.com/sudoku.html
// http://www.mirror.co.uk/news/weird-news/worlds-hardest-sudoku-can-you-242294
char *puzzle="..53.....8.....2..7..1.5..4....53...1..7...6..32...8..6.5....9..4....3.....97..";

int main()
{
    klee_make_symbolic(cells, sizeof cells, "cells");

    // process text line:
    for (int row=0; row<9; row++)
        for (int column=0; column<9; column++)
            {
                char c=puzzle[row*9 + column];
                if (c!='.')
                    klee_assume (cells[row][column]==c-'0');
                else
                {
                    klee_assume (cells[row][column]>=1);
                    klee_assume (cells[row][column]<=9);
                }
            };

    // for all 9 rows
    for (int row=0; row<9; row++)
    {
        klee_assume (((1<<cells[row][0]) |
                     (1<<cells[row][1]) |
                     (1<<cells[row][2]) |
                     (1<<cells[row][3]) |
                     (1<<cells[row][4]) |
                     (1<<cells[row][5]) |
                     (1<<cells[row][6]) |
                     (1<<cells[row][7]) |
                     (1<<cells[row][8]))==0x3FE ); // 11 1111 1110

    };

    // for all 9 columns
    for (int c=0; c<9; c++)
    {
        klee_assume (((1<<cells[0][c]) |
                     (1<<cells[1][c]) |
                     (1<<cells[2][c]) |
                     (1<<cells[3][c]) |
                     (1<<cells[4][c]) |
                     (1<<cells[5][c]) |
                     (1<<cells[6][c]) |
                     (1<<cells[7][c]) |
                     (1<<cells[8][c]))==0x3FE ); // 11 1111 1110

    };

    // enumerate all 9 squares
    for (int r=0; r<9; r+=3)
        for (int c=0; c<9; c+=3)
            {
                // add constraints for each 3*3 square:
                klee_assume ((1<<cells[r+0][c+0] |
                             1<<cells[r+0][c+1] |
                             1<<cells[r+0][c+2] |
                             1<<cells[r+1][c+0] |
                             1<<cells[r+1][c+1] |

```

```

        1<<cells[r+1][c+2] |
        1<<cells[r+2][c+0] |
        1<<cells[r+2][c+1] |
        1<<cells[r+2][c+2])==0x3FE ); // 11 1111 1110
    };

    // at this point, all constraints must be satisfied
    klee_assert(0);
};

```

```

% time klee klee_sudoku_or2.bc
KLEE: output directory is "/home/klee/klee-out-99"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_sudoku_or2.c:93: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 7119
KLEE: done: completed paths = 1
KLEE: done: generated tests = 1

real    0m35.312s
user    0m34.945s
sys     0m0.318s

```

That works much faster: perhaps KLEE indeed handle this *intrinsic* in a special way. And, as we see, the only one path has been found (one we actually interesting in it) instead of 161.

It's still much slower than Z3Py solution, though.

## 9.5 Unit test: HTML/CSS color

The most popular ways to represent HTML/CSS color is by English name (e.g., "red") and by 6-digit hexadecimal number (e.g., "#0077CC"). There is third, less popular way: if each byte in hexadecimal number has two doubling digits, it can be *abbreviated*, thus, "#0077CC" can be written just as "#07C".

Let's write a function to convert 3 color components into name (if possible, first priority), 3-digit hexadecimal form (if possible, second priority), or as 6-digit hexadecimal form (as a last resort).

```

#include <string.h>
#include <stdio.h>
#include <stdint.h>

void HTML_color(uint8_t R, uint8_t G, uint8_t B, char* out)
{
    if (R==0xFF && G==0 && B==0)
    {
        strcpy (out, "red");
        return;
    };

    if (R==0x0 && G==0xFF && B==0)
    {
        strcpy (out, "green");
        return;
    };

    if (R==0 && G==0 && B==0xFF)
    {
        strcpy (out, "blue");
        return;
    };

    // abbreviated hexadecimal
    if (R>>4==(R&0xF) && G>>4==(G&0xF) && B>>4==(B&0xF))
    {
        sprintf (out, "#%X%X%X", R&0xF, G&0xF, B&0xF);
        return;
    };

    // last resort
    sprintf (out, "%02X%02X%02X", R, G, B);
};

```

```

int main()
{
    uint8_t R, G, B;
    klee_make_symbolic (&R, sizeof R, "R");
    klee_make_symbolic (&G, sizeof R, "G");
    klee_make_symbolic (&B, sizeof R, "B");

    char tmp[16];

    HTML_color(R, G, B, tmp);
};

```

There are 5 possible paths in function, and let's see, if KLEE could find them all? It's indeed so:

```

% clang -emit-llvm -c -g color.c

% klee color.bc
KLEE: output directory is "/home/klee/klee-out-134"
KLEE: WARNING: undefined reference to function: sprintf
KLEE: WARNING: undefined reference to function: strcpy
KLEE: WARNING ONCE: calling external: strcpy(51867584, 51598960)
KLEE: ERROR: /home/klee/color.c:33: external call with symbolic argument: sprintf
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/color.c:28: external call with symbolic argument: sprintf
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 479
KLEE: done: completed paths = 19
KLEE: done: generated tests = 5

```

We can ignore calls to strcpy() and sprintf(), because we are not really interesting in state of `out` variable. So there are exactly 5 paths:

```

% ls klee-last
assembly.ll  run.stats      test000003.ktest  test000005.ktest
info         test000001.ktest  test000003.pc     test000005.pc
messages.txt test000002.ktest  test000004.ktest  warnings.txt
run.istats   test000003.exec.err  test000005.exec.err

```

1st set of input variables will result in "red" string:

```

% ktest-tool --write-ints klee-last/test000001.ktest
ktest file : 'klee-last/test000001.ktest'
args       : ['color.bc']
num objects: 3
object 0: name: b'R'
object 0: size: 1
object 0: data: b'\xff'
object 1: name: b'G'
object 1: size: 1
object 1: data: b'\x00'
object 2: name: b'B'
object 2: size: 1
object 2: data: b'\x00'

```

2nd set of input variables will result in "green" string:

```

% ktest-tool --write-ints klee-last/test000002.ktest
ktest file : 'klee-last/test000002.ktest'
args       : ['color.bc']
num objects: 3
object 0: name: b'R'
object 0: size: 1
object 0: data: b'\x00'
object 1: name: b'G'
object 1: size: 1
object 1: data: b'\xff'
object 2: name: b'B'
object 2: size: 1
object 2: data: b'\x00'

```

3rd set of input variables will result in "#010000" string:

```

% ktest-tool --write-ints klee-last/test000003.ktest
ktest file : 'klee-last/test000003.ktest'
args       : ['color.bc']

```

```
num objects: 3
object 0: name: b'R'
object 0: size: 1
object 0: data: b'\x01'
object 1: name: b'G'
object 1: size: 1
object 1: data: b'\x00'
object 2: name: b'B'
object 2: size: 1
object 2: data: b'\x00'
```

4th set of input variables will result in “blue” string:

```
% ktest-tool --write-ints klee-last/test000004.ktest
ktest file : 'klee-last/test000004.ktest'
args      : ['color.bc']
num objects: 3
object 0: name: b'R'
object 0: size: 1
object 0: data: b'\x00'
object 1: name: b'G'
object 1: size: 1
object 1: data: b'\x00'
object 2: name: b'B'
object 2: size: 1
object 2: data: b'\xff'
```

5th set of input variables will result in “#F01” string:

```
% ktest-tool --write-ints klee-last/test000005.ktest
ktest file : 'klee-last/test000005.ktest'
args      : ['color.bc']
num objects: 3
object 0: name: b'R'
object 0: size: 1
object 0: data: b'\xff'
object 1: name: b'G'
object 1: size: 1
object 1: data: b'\x00'
object 2: name: b'B'
object 2: size: 1
object 2: data: b'\x11'
```

These 5 sets of input variables can form a unit test for our function.

## 9.6 Unit test: strcmp() function

The standard `strcmp()` function from C library can return 0, -1 or 1, depending of comparison result.

Here is my own implementation of `strcmp()` :

```
int my_strcmp(const char *s1, const char *s2)
{
    int ret = 0;

    while (1)
    {
        ret = *(unsigned char *) s1 - *(unsigned char *) s2;
        if (ret!=0)
            break;
        if ((*s1==0) || (*s2)==0)
            break;
        s1++;
        s2++;
    };

    if (ret < 0)
    {
        return -1;
    } else if (ret > 0)
    {
        return 1;
    }
}
```

```

    return 0;
}

int main()
{
    char input1[2];
    char input2[2];

    klee_make_symbolic(input1, sizeof input1, "input1");
    klee_make_symbolic(input2, sizeof input2, "input2");

    klee_assume((input1[0]>='a') && (input1[0]<='z'));
    klee_assume((input2[0]>='a') && (input2[0]<='z'));

    klee_assume(input1[1]==0);
    klee_assume(input2[1]==0);

    my_strcmp (input1, input2);
};

```

Let's find out, if KLEE is capable of finding all three paths? I intentionally made things simpler for KLEE by limiting input arrays to two 2 bytes or to 1 character + terminal zero byte.

```

% clang -emit-llvm -c -g strcmp.c

% klee strcmp.bc
KLEE: output directory is "/home/klee/klee-out-131"
KLEE: ERROR: /home/klee/strcmp.c:35: invalid klee_assume call (provably false)
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/strcmp.c:36: invalid klee_assume call (provably false)
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 137
KLEE: done: completed paths = 5
KLEE: done: generated tests = 5

% ls klee-last
assembly.ll  run.stats          test000002.ktest   test000004.ktest
info         test000001.ktest   test000002.pc      test000005.ktest
messages.txt test000001.pc       test000002.user.err warnings.txt
run.istats   test000001.user.err test000003.ktest

```

The first two errors are about `klee_assume()`. These are input values on which `klee_assume()` calls are stuck. We can ignore them, or take a peek out of curiosity:

```

% ktest-tool --write-ints klee-last/test000001.ktest
ktest file : 'klee-last/test000001.ktest'
args       : ['strcmp.bc']
num objects: 2
object  0: name: b'input1'
object  0: size: 2
object  0: data: b'\x00\x00'
object  1: name: b'input2'
object  1: size: 2
object  1: data: b'\x00\x00'

% ktest-tool --write-ints klee-last/test000002.ktest
ktest file : 'klee-last/test000002.ktest'
args       : ['strcmp.bc']
num objects: 2
object  0: name: b'input1'
object  0: size: 2
object  0: data: b'a\xff'
object  1: name: b'input2'
object  1: size: 2
object  1: data: b'\x00\x00'

```

Three rest files are the input values for each path inside of my implementation of `strcmp()`:

```

% ktest-tool --write-ints klee-last/test000003.ktest
ktest file : 'klee-last/test000003.ktest'
args       : ['strcmp.bc']
num objects: 2
object  0: name: b'input1'
object  0: size: 2

```



```

object 0: data: b'b\x00'
object 1: name: b'input2'
object 1: size: 2
object 1: data: b'c\x00'

% ktest-tool --write-ints klee-last/test000004.ktest
ktest file : 'klee-last/test000004.ktest'
args      : ['strcmp.bc']
num objects: 2
object 0: name: b'input1'
object 0: size: 2
object 0: data: b'c\x00'
object 1: name: b'input2'
object 1: size: 2
object 1: data: b'a\x00'

% ktest-tool --write-ints klee-last/test000005.ktest
ktest file : 'klee-last/test000005.ktest'
args      : ['strcmp.bc']
num objects: 2
object 0: name: b'input1'
object 0: size: 2
object 0: data: b'a\x00'
object 1: name: b'input2'
object 1: size: 2
object 1: data: b'a\x00'

```

3rd is about first argument (“b”) is lesser than the second (“c”). 4th is opposite (“c” and “a”). 5th is when they are equal (“a” and “a”).

Using these 3 test cases, we’ve got full coverage of our implementation of `strcmp()`.

## 9.7 UNIX date/time

UNIX date/time<sup>76</sup> is a number of seconds that have elapsed since 1-Jan-1970 00:00 UTC. C/C++ `gmtime()` function is used to decode this value into human-readable date/time.

Here is a piece of code I’ve copped from some ancient version of Minix OS (<http://www.cise.ufl.edu/~cop4600/cgi-bin/lxr/http/source.cgi/lib/ansi/gmtime.c>) and reworked slightly:

```

1 #include <stdint.h>
2 #include <time.h>
3 #include <assert.h>
4
5 /*
6  * Copypasted and reworked from
7  * http://www.cise.ufl.edu/~cop4600/cgi-bin/lxr/http/source.cgi/lib/ansi/loc_time.h
8  * http://www.cise.ufl.edu/~cop4600/cgi-bin/lxr/http/source.cgi/lib/ansi/misc.c
9  * http://www.cise.ufl.edu/~cop4600/cgi-bin/lxr/http/source.cgi/lib/ansi/gmtime.c
10 */
11
12 #define YEAR0      1900
13 #define EPOCH_YR  1970
14 #define SECS_DAY  (24L * 60L * 60L)
15 #define YEARSIZE(year) (LEAPYEAR(year) ? 366 : 365)
16
17 const int _ytab[2][12] =
18 {
19     { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 },
20     { 31, 29, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 }
21 };
22
23 const char *_days[] =
24 {
25     "Sunday", "Monday", "Tuesday", "Wednesday",
26     "Thursday", "Friday", "Saturday"
27 };
28
29 const char *_months[] =
30 {
31     "January", "February", "March",
32     "April", "May", "June",
33     "July", "August", "September",

```

<sup>76</sup>[https://en.wikipedia.org/wiki/Unix\\_time](https://en.wikipedia.org/wiki/Unix_time)

```

34     "October", "November", "December"
35 };
36
37 #define LEAPYEAR(year) (!((year) % 4) && (((year) % 100) || !((year) % 400)))
38
39 void decode_UNIX_time(const time_t time)
40 {
41     unsigned int dayclock, dayno;
42     int year = EPOCH_YR;
43
44     dayclock = (unsigned long)time % SECS_DAY;
45     dayno = (unsigned long)time / SECS_DAY;
46
47     int seconds = dayclock % 60;
48     int minutes = (dayclock % 3600) / 60;
49     int hour = dayclock / 3600;
50     int wday = (dayno + 4) % 7;
51     while (dayno >= YEARSIZE(year))
52     {
53         dayno -= YEARSIZE(year);
54         year++;
55     }
56
57     year = year - YEAR0;
58
59     int month = 0;
60
61     while (dayno >= _ytab[LEAPYEAR(year)][month])
62     {
63         dayno -= _ytab[LEAPYEAR(year)][month];
64         month++;
65     }
66
67     char *s;
68     switch (month)
69     {
70         case 0: s="January"; break;
71         case 1: s="February"; break;
72         case 2: s="March"; break;
73         case 3: s="April"; break;
74         case 4: s="May"; break;
75         case 5: s="June"; break;
76         case 6: s="July"; break;
77         case 7: s="August"; break;
78         case 8: s="September"; break;
79         case 9: s="October"; break;
80         case 10: s="November"; break;
81         case 11: s="December"; break;
82         default:
83             assert(0);
84     };
85
86     printf ("%04d-%s-%02d %02d:%02d:%02d\n", YEAR0+year, s, dayno+1, hour, minutes, seconds);
87     printf ("week day: %s\n", _days[wday]);
88 }
89
90 int main()
91 {
92     uint32_t time;
93
94     klee_make_symbolic(&time, sizeof time, "time");
95
96     decode_UNIX_time(time);
97
98     return 0;
99 }

```

Let's try it:

```

% clang -emit-llvm -c -g klee_time1.c
...
% klee klee_time1.bc
KLEE: output directory is "/home/klee/klee-out-107"
KLEE: WARNING: undefined reference to function: printf
KLEE: ERROR: /home/klee/klee_time1.c:86: external call with symbolic argument: printf

```

```
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_time1.c:83: ASSERTION FAIL: 0
KLEE: NOTE: now ignoring this error at this location
```

```
KLEE: done: total instructions = 101579
KLEE: done: completed paths = 1635
KLEE: done: generated tests = 2
```

Wow, assert() at line 83 has been triggered, why? Let's see a value of UNIX time which triggers it:

```
% ls klee-last | grep err
test000001.exec.err
test000002.assert.err

% ktest-tool --write-ints klee-last/test000002.ktest
ktest file : 'klee-last/test000002.ktest'
args      : ['klee_time1.bc']
num objects: 1
object 0: name: b'time'
object 0: size: 4
object 0: data: 978278400
```

Let's decode this value using UNIX date utility:

```
% date -u --date=@978278400
Sun Dec 31 16:00:00 UTC 2000
```

After my investigation, I've found that `month` variable can hold incorrect value of 12 (while 11 is maximal, for December), because `LEAPYEAR()` macro should receive year number as 2000, not as 100. So I've introduced a bug during rewriting this function, and KLEE found it!

Just interesting, what would be if I'll replace `switch()` to array of strings, like it usually happens in concise C/C++ code?

```
...
const char *_months[] =
{
    "January", "February", "March",
    "April", "May", "June",
    "July", "August", "September",
    "October", "November", "December"
};

...

while (dayno >= _ytab[LEAPYEAR(year)][month])
{
    dayno -= _ytab[LEAPYEAR(year)][month];
    month++;
}

char *s=_months[month];

printf ("%04d-%s-%02d %02d:%02d:%02d\n", YEAR0+year, s, dayno+1, hour, minutes, seconds);
printf ("week day: %s\n", _days[wday]);

...
```

KLEE detects attempt to read beyond array boundaries:

```
% klee klee_time2.bc
KLEE: output directory is "/home/klee/klee-out-108"
KLEE: WARNING: undefined reference to function: printf
KLEE: ERROR: /home/klee/klee_time2.c:69: external call with symbolic argument: printf
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_time2.c:67: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 101716
KLEE: done: completed paths = 1635
KLEE: done: generated tests = 2
```

This is the same UNIX time value we've already seen:

```

% ls klee-last | grep err
test000001.exec.err
test000002.ptr.err

% ktest-tool --write-ints klee-last/test000002.ktest
ktest file : 'klee-last/test000002.ktest'
args       : ['klee_time2.bc']
num objects: 1
object 0: name: b'time'
object 0: size: 4
object 0: data: 978278400

```

So, if this piece of code can be triggered on remote computer, with this input value (*input of death*), it's possible to crash the process (with some luck, though).

OK, now I'm fixing a bug by moving year subtracting expression to line 43, and let's find, what UNIX time value corresponds to some fancy date like 2022-February-2?

```

1 #include <stdint.h>
2 #include <time.h>
3 #include <assert.h>
4
5 #define YEAR0          1900
6 #define EPOCH_YR      1970
7 #define SECS_DAY      (24L * 60L * 60L)
8 #define YEARSIZE(year) (LEAPYEAR(year) ? 366 : 365)
9
10 const int _ytab[2][12] =
11 {
12     { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 },
13     { 31, 29, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 }
14 };
15
16 #define LEAPYEAR(year) (!((year) % 4) && (((year) % 100) || !((year) % 400)))
17
18 void decode_UNIX_time(const time_t time)
19 {
20     unsigned int dayclock, dayno;
21     int year = EPOCH_YR;
22
23     dayclock = (unsigned long)time % SECS_DAY;
24     dayno = (unsigned long)time / SECS_DAY;
25
26     int seconds = dayclock % 60;
27     int minutes = (dayclock % 3600) / 60;
28     int hour = dayclock / 3600;
29     int wday = (dayno + 4) % 7;
30     while (dayno >= YEARSIZE(year))
31     {
32         dayno -= YEARSIZE(year);
33         year++;
34     }
35
36     int month = 0;
37
38     while (dayno >= _ytab[LEAPYEAR(year)][month])
39     {
40         dayno -= _ytab[LEAPYEAR(year)][month];
41         month++;
42     }
43     year = year - YEAR0;
44
45     if (YEAR0+year==2022 && month==1 && dayno+1==22)
46         klee_assert(0);
47 }
48 int main()
49 {
50     uint32_t time;
51
52     klee_make_symbolic(&time, sizeof time, "time");
53
54     decode_UNIX_time(time);
55
56     return 0;

```

```
% clang -emit-llvm -c -g klee_time3.c
...

% klee klee_time3.bc
KLEE: output directory is "/home/klee/klee-out-109"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_time3.c:47: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 101087
KLEE: done: completed paths = 1635
KLEE: done: generated tests = 1635

% ls klee-last | grep err
test000587.external.err

% ktest-tool --write-ints klee-last/test000587.ktest
ktest file : 'klee-last/test000587.ktest'
args       : ['klee_time3.bc']
num objects: 1
object    0: name: b'time'
object    0: size: 4
object    0: data: 1645488640

% date -u --date='@1645488640'
Tue Feb 22 00:10:40 UTC 2022
```

Success, but hours/minutes/seconds are seems random—they are random indeed, because, KLEE satisfied all constraints we've put, nothing else. We didn't ask it to set hours/minutes/seconds to zeroes.

Let's add constraints to hours/minutes/seconds as well:

```
...

    if (YEAR0+year==2022 && month==1 && dayno+1==22 && hour==22 && minutes==22 && seconds==22)
        klee_assert(0);

...
```

Let's run it and check ...

```
% ktest-tool --write-ints klee-last/test000597.ktest
ktest file : 'klee-last/test000597.ktest'
args       : ['klee_time3.bc']
num objects: 1
object    0: name: b'time'
object    0: size: 4
object    0: data: 1645568542

% date -u --date='@1645568542'
Tue Feb 22 22:22:22 UTC 2022
```

Now that is precise.

Yes, of course, C/C++ libraries has function(s) to encode human-readable date into UNIX time value, but what we've got here is KLEE working *antipode* of decoding function, *inverse function* in a way.

## 9.8 Inverse function for base64 decoder

It's piece of cake for KLEE to reconstruct input base64 string given just base64 decoder code without corresponding encoder code. I've copy-pasted this piece of code from <http://www.opensource.apple.com/source/QuickTimeStreamingServer/QuickTimeStreamingServer-452/CommonUtilitiesLib/base64.c>.

We add constraints (lines 84, 85) so that output buffer must have byte values from 0 to 15. We also tell to KLEE that the Base64decode() function must return 16 (i.e., size of output buffer in bytes, line 82).

```
1 #include <string.h>
2 #include <stdint.h>
3 #include <stdbool.h>
4
5 // copy-pasted from http://www.opensource.apple.com/source/QuickTimeStreamingServer/QuickTimeStreamingServer-452/
   CommonUtilitiesLib/base64.c
```

```

6
7 static const unsigned char pr2six[256] =
8 {
9     /* ASCII table */
10    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
11    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
12    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 62, 64, 64, 64, 63,
13    52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 64, 64, 64, 64, 64, 64,
14    64, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14,
15    15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 64, 64, 64, 64, 64,
16    64, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40,
17    41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 64, 64, 64, 64, 64,
18    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
19    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
20    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
21    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
22    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
23    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
24    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64,
25    64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64, 64
26 };
27
28 int Base64decode(char *bufplain, const char *bufcoded)
29 {
30     int nbytesdecoded;
31     register const unsigned char *bufin;
32     register unsigned char *bufout;
33     register int nprbytes;
34
35     bufin = (const unsigned char *) bufcoded;
36     while (pr2six[*bufin] <= 63);
37     nprbytes = (bufin - (const unsigned char *) bufcoded) - 1;
38     nbytesdecoded = ((nprbytes + 3) / 4) * 3;
39
40     bufout = (unsigned char *) bufplain;
41     bufin = (const unsigned char *) bufcoded;
42
43     while (nprbytes > 4) {
44         *(bufout++) =
45             (unsigned char) (pr2six[*bufin] << 2 | pr2six[bufin[1]] >> 4);
46         *(bufout++) =
47             (unsigned char) (pr2six[bufin[1]] << 4 | pr2six[bufin[2]] >> 2);
48         *(bufout++) =
49             (unsigned char) (pr2six[bufin[2]] << 6 | pr2six[bufin[3]]);
50         bufin += 4;
51         nprbytes -= 4;
52     }
53
54     /* Note: (nprbytes == 1) would be an error, so just ignore that case */
55     if (nprbytes > 1) {
56         *(bufout++) =
57             (unsigned char) (pr2six[*bufin] << 2 | pr2six[bufin[1]] >> 4);
58     }
59     if (nprbytes > 2) {
60         *(bufout++) =
61             (unsigned char) (pr2six[bufin[1]] << 4 | pr2six[bufin[2]] >> 2);
62     }
63     if (nprbytes > 3) {
64         *(bufout++) =
65             (unsigned char) (pr2six[bufin[2]] << 6 | pr2six[bufin[3]]);
66     }
67
68     *(bufout++) = '\0';
69     nbytesdecoded -= (4 - nprbytes) & 3;
70     return nbytesdecoded;
71 }
72
73 int main()
74 {
75     char input[32];
76     uint8_t output[16+1];
77
78     klee_make_symbolic(input, sizeof input, "input");
79
80     klee_assume(input[31]==0);
81

```

```

82     klee_assume (Base64decode(output, input)==16);
83
84     for (int i=0; i<16; i++)
85         klee_assume (output[i]==i);
86
87     klee_assert(0);
88
89     return 0;
90 }

```

```

% clang -emit-llvm -c -g klee_base64.c
...
% klee klee_base64.bc
KLEE: output directory is "/home/klee/klee-out-99"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: ERROR: /home/klee/klee_base64.c:99: invalid klee_assume call (provably false)
KLEE: NOTE: now ignoring this error at this location
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_base64.c:104: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_base64.c:85: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_base64.c:81: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_base64.c:65: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
...

```

We're interested in the second error, where `klee_assert()` has been triggered:

```

% ls klee-last | grep err
test000001.user.err
test000002.external.err
test000003.ptr.err
test000004.ptr.err
test000005.ptr.err

% ktest-tool --write-ints klee-last/test000002.ktest
ktest file : 'klee-last/test000002.ktest'
args      : ['klee_base64.bc']
num objects: 1
object 0: name: b'input'
object 0: size: 32
object 0: data: b'AAECAwQFBgcICQoLDA00D4\x00\xff\xff\xff\xff\xff\xff\xff\x00'

```

This is indeed a real base64 string, terminated with the zero byte, just as it's requested by C/C++ standards. The final zero byte at 31th byte (starting at zeroth byte) is our deed: so that KLEE would report lesser number of errors.

The base64 string is indeed correct:

```

% echo AAECAwQFBgcICQoLDA00D4 | base64 -d | hexdump -C
base64: invalid input
00000000 00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f |.....|
00000010

```

base64 decoder Linux utility I've just run blaming for "invalid input"—it means the input string is not properly padded. Now let's pad it manually, and decoder utility will no complain anymore:

```

% echo AAECAwQFBgcICQoLDA00D4== | base64 -d | hexdump -C
00000000 00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f |.....|
00000010

```

The reason our generated base64 string is not padded is because base64 decoders are usually discards padding symbols ("=") at the end. In other words, they are not require them, so is the case of our decoder. Hence, padding symbols are left unnoticed to KLEE.

So we again made *antipode* or *inverse function* of base64 decoder.

## 9.9 CRC (Cyclic redundancy check)

### 9.9.1 Buffer alteration case #1

Sometimes, you need to alter a piece of data which is *protected* by some kind of checksum or **CRC**, and you can't change checksum or CRC value, but can alter piece of data so that checksum will remain the same.

Let's pretend, we've got a piece of data with "Hello, world!" string at the beginning and "and goodbye" string at the end. We can alter 14 characters at the middle, but for some reason, they must be in *a..z* limits, but we can put any characters there. CRC64 of the whole block must be `0x12345678abcdef12`.

Let's see<sup>77</sup>:

```
#include <string.h>
#include <stdint.h>

uint64_t crc64(uint64_t crc, unsigned char *buf, int len)
{
    int k;

    crc = ~crc;
    while (len--)
    {
        crc ^= *buf++;
        for (k = 0; k < 8; k++)
            crc = crc & 1 ? (crc >> 1) ^ 0x42f0e1eba9ea3693 : crc >> 1;
    }
    return crc;
}

int main()
{
#define HEAD_STR "Hello, world!.. "
#define HEAD_SIZE strlen(HEAD_STR)
#define TAIL_STR " ... and goodbye"
#define TAIL_SIZE strlen(TAIL_STR)
#define MID_SIZE 14 // work
#define BUF_SIZE HEAD_SIZE+TAIL_SIZE+MID_SIZE

    char buf[BUF_SIZE];

    klee_make_symbolic(buf, sizeof buf, "buf");

    klee_assume (memcmp (buf, HEAD_STR, HEAD_SIZE)==0);

    for (int i=0; i<MID_SIZE; i++)
        klee_assume (buf[HEAD_SIZE+i]>='a' && buf[HEAD_SIZE+i]<='z');

    klee_assume (memcmp (buf+HEAD_SIZE+MID_SIZE, TAIL_STR, TAIL_SIZE)==0);

    klee_assume (crc64 (0, buf, BUF_SIZE)==0x12345678abcdef12);

    klee_assert(0);

    return 0;
}
```

Since our code uses `memcmp()` standard C/C++ function, we need to add `--libc=uclibc` switch, so KLEE will use its own uClibc implementation.

```
% clang -emit-llvm -c -g klee_CRC64.c
% time klee --libc=uclibc klee_CRC64.bc
```

It takes about 1 minute (on my Intel Core i3-3110M 2.4GHz notebook) and we getting this:

```
...
real    0m52.643s
user    0m51.232s
sys     0m0.239s
...
% ls klee-last | grep err
test000001.user.err
test000002.user.err
```

<sup>77</sup>There are several slightly different CRC64 implementations, the one I use here can also be different from popular ones.



```

test000003.user.err
test000004.external.err

% ktest-tool --write-ints klee-last/test000004.ktest
ktest file : 'klee-last/test000004.ktest'
args      : ['klee_CRC64.bc']
num objects: 1
object   0: name: b'buf'
object   0: size: 46
object   0: data: b'Hello, world!.. qqlicayzceamyw ... and goodbye'

```

Maybe it's slow, but definitely faster than bruteforce. Indeed,  $\log_2 26^{14} \approx 65.8$  which is close to 64 bits. In other words, one need  $\approx 14$  latin characters to encode 64 bits. And KLEE + [SMT](#) solver needs 64 bits at some place it can alter to make final CRC64 value equal to what we defined.

I tried to reduce length of the *middle block* to 13 characters: no luck for KLEE then, it has no space enough.

## 9.9.2 Buffer alteration case #2

I went sadistic: what if the buffer must contain the CRC64 value which, after calculation of CRC64, will result in the same value? Fascinatedly, KLEE can solve this. The buffer will have the following format:

```
Hello, world! <8 bytes (64-bit value)> and goodbye <6 more bytes>
```

```

int main()
{
#define HEAD_STR "Hello, world!.. "
#define HEAD_SIZE strlen(HEAD_STR)
#define TAIL_STR " ... and goodbye"
#define TAIL_SIZE strlen(TAIL_STR)
// 8 bytes for 64-bit value:
#define MID_SIZE 8
#define BUF_SIZE HEAD_SIZE+TAIL_SIZE+MID_SIZE+6

    char buf[BUF_SIZE];

    klee_make_symbolic(buf, sizeof buf, "buf");

    klee_assume (memcmp (buf, HEAD_STR, HEAD_SIZE)==0);

    klee_assume (memcmp (buf+HEAD_SIZE+MID_SIZE, TAIL_STR, TAIL_SIZE)==0);

    uint64_t mid_value=*(uint64_t*)(buf+HEAD_SIZE);
    klee_assume (crc64 (0, buf, BUF_SIZE)==mid_value);

    klee_assert(0);

    return 0;
}

```

It works:

```

% time klee --libc=uclibc klee_CRC64.bc
...
real    5m17.081s
user    5m17.014s
sys     0m0.319s

% ls klee-last | grep err
test000001.user.err
test000002.user.err
test000003.external.err

% ktest-tool --write-ints klee-last/test000003.ktest
ktest file : 'klee-last/test000003.ktest'
args      : ['klee_CRC64.bc']
num objects: 1
object   0: name: b'buf'
object   0: size: 46
object   0: data: b'Hello, world!.. T+]\xb9A\x08\x0fq ... and goodbye\xb6\x8f\x9c\xd8\xc5\x00'

```

8 bytes between two strings is 64-bit value which equals to CRC64 of this whole block. Again, it's faster than brute-force way to find it. If to decrease last spare 6-byte buffer to 4 bytes or less, KLEE works so long so I've stopped it.

### 9.9.3 Recovering input data for given CRC32 value of it

I've always wanted to do so, but everyone knows this is impossible for input buffers larger than 4 bytes. As my experiments show, it's still possible for tiny input buffers of data, which is constrained in some way.

The CRC32 value of 6-byte "SILVER" string is known: `0xDFA3DFDD`. KLEE can find this 6-byte string, if it knows that each byte of input buffer is in A..Z limits:

```
1 #include <stdint.h>
2 #include <stdbool.h>
3
4 uint32_t crc32(uint32_t crc, unsigned char *buf, int len)
5 {
6     int k;
7
8     crc = ~crc;
9     while (len--)
10    {
11        crc ^= *buf++;
12        for (k = 0; k < 8; k++)
13            crc = crc & 1 ? (crc >> 1) ^ 0xedb88320 : crc >> 1;
14    }
15    return ~crc;
16 }
17
18 #define SIZE 6
19
20 bool find_string(char str[SIZE])
21 {
22     int i=0;
23     for (i=0; i<SIZE; i++)
24         if (str[i]<'A' || str[i]>'Z')
25             return false;
26
27     if (crc32(0, &str[0], SIZE)!=0xDFA3DFDD)
28         return false;
29
30     // OK, input str is valid
31     klee_assert(0); // force KLEE to produce .err file
32     return true;
33 };
34
35 int main()
36 {
37     uint8_t str[SIZE];
38
39     klee_make_symbolic(str, sizeof str, "str");
40
41     find_string(str);
42
43     return 0;
44 }
```

```
% clang -emit-llvm -c -g klee_SILVER.c
...
% klee klee_SILVER.bc
...
% ls klee-last | grep err
test000013.external.err

% ktest-tool --write-ints klee-last/test000013.ktest
ktest file : 'klee-last/test000013.ktest'
args      : ['klee_SILVER.bc']
num objects: 1
object   0: name: b'str'
object   0: size: 6
object   0: data: b'SILVER'
```

Still, it's no magic: if to remove condition at lines 23..25 (i.e., if to relax constraints), KLEE will produce some other string, which will be still correct for the CRC32 value given.

It works, because 6 Latin characters in A..Z limits contain  $\approx 28.2$  bits:  $\log_2 26^6 \approx 28.2$ , which is even smaller value than 32. In other words, the final CRC32 value holds enough bits to recover  $\approx 28.2$  bits of input.

The input buffer can be even bigger, if each byte of it will be in even tighter constraints (decimal digits, binary digits, etc).

#### 9.9.4 In comparison with other hashing algorithms

Things are that easy for some other hashing algorithms like *Fletcher checksum*, but not for cryptographically secure ones (like MD5, SHA1, etc), they are protected from such simple cryptoanalysis. See also: 10.

### 9.10 LZSS decompressor

I've googled for a very simple LZSS<sup>78</sup> decompressor and landed at this page: <http://www.opensource.apple.com/source/boot/boot-132/i386/boot2/lzss.c>.

Let's pretend, we're looking at unknown compressing algorithm with no compressor available. Will it be possible to reconstruct a compressed piece of data so that decompressor would generate data we need?

Here is my first experiment:

```
// copy-pasted from http://www.opensource.apple.com/source/boot/boot-132/i386/boot2/lzss.c
//
#include <string.h>
#include <stdint.h>
#include <stdbool.h>

#define N 4096 /* size of ring buffer - must be power of 2 */
#define N 32 /* size of ring buffer - must be power of 2 */
#define F 18 /* upper limit for match_length */
#define THRESHOLD 2 /* encode string into position and length
                    if match_length is greater than this */
#define NIL N /* index for root of binary search trees */

int
decompress_lzss(uint8_t *dst, uint8_t *src, uint32_t srclen)
{
    /* ring buffer of size N, with extra F-1 bytes to aid string comparison */
    uint8_t *dststart = dst;
    uint8_t *srcend = src + srclen;
    int i, j, k, r, c;
    unsigned int flags;
    uint8_t text_buf[N + F - 1];

    dst = dststart;
    srcend = src + srclen;
    for (i = 0; i < N - F; i++)
        text_buf[i] = ' ';
    r = N - F;
    flags = 0;
    for ( ; ; ) {
        if (((flags >>= 1) & 0x100) == 0) {
            if (src < srcend) c = *src++; else break;
            flags = c | 0xFF00; /* uses higher byte cleverly */
        } /* to count eight */
        if (flags & 1) {
            if (src < srcend) c = *src++; else break;
            *dst++ = c;
            text_buf[r++] = c;
            r &= (N - 1);
        } else {
            if (src < srcend) i = *src++; else break;
            if (src < srcend) j = *src++; else break;
            i |= ((j & 0xF0) << 4);
            j = (j & 0x0F) + THRESHOLD;
            for (k = 0; k <= j; k++) {
                c = text_buf[(i + k) & (N - 1)];
                *dst++ = c;
                text_buf[r++] = c;
                r &= (N - 1);
            }
        }
    }
    return dst - dststart;
}
```

<sup>78</sup>Lempel-Ziv-Storer-Szymanski

```

}

int main()
{
#define COMPRESSED_LEN 15
    uint8_t input[COMPRESSED_LEN];
    uint8_t plain[24];
    uint32_t size=COMPRESSED_LEN;

    klee_make_symbolic(input, sizeof input, "input");

    decompress_lzss(plain, input, size);

    // https://en.wikipedia.org/wiki/Buffalo_buffalo_Buffalo_buffalo_buffalo_buffalo_Buffalo_buffalo
    for (int i=0; i<23; i++)
        klee_assume (plain[i]=="Buffalo buffalo Buffalo"[i]);

    klee_assert(0);

    return 0;
}

```

What I did is changing size of ring buffer from 4096 to 32, because if bigger, KLEE consumes all [RAM<sup>79</sup>](#) it can. But I've found that KLEE can live with that small buffer. I've also decreased `COMPRESSED_LEN` gradually to check, whether KLEE would find compressed piece of data, and it did:

```

% clang -emit-llvm -c -g klee_lzss.c
...

% time klee klee_lzss.bc
KLEE: output directory is "/home/klee/klee-out-7"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: ERROR: /home/klee/klee_lzss.c:122: invalid klee_assume call (provably false)
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_lzss.c:47: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_lzss.c:37: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_lzss.c:124: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location

KLEE: done: total instructions = 41417919
KLEE: done: completed paths = 437820
KLEE: done: generated tests = 4

real    13m0.215s
user    11m57.517s
sys     1m2.187s

% ls klee-last | grep err
test000001.user.err
test000002.ptr.err
test000003.ptr.err
test000004.external.err

% ktest-tool --write-ints klee-last/test000004.ktest
ktest file : 'klee-last/test000004.ktest'
args      : ['klee_lzss.bc']
num objects: 1
object   0: name: b'input'
object   0: size: 15
object   0: data: b'\xffBuffalo \x01b\x0f\x03\r\x05'

```

KLEE consumed  $\approx 1GB$  of RAM and worked for  $\approx 15$  minutes (on my Intel Core i3-3110M 2.4GHz notebook), but here it is, a 15 bytes which, if decompressed by our cypasted algorithm, will result in desired text!

During my experimentation, I've found that KLEE can do even more cooler thing, to find out size of compressed piece of data:

```

int main()
{
    uint8_t input[24];

```

<sup>79</sup>Random-access memory

```

uint8_t plain[24];
uint32_t size;

klee_make_symbolic(input, sizeof input, "input");
klee_make_symbolic(&size, sizeof size, "size");

decompress_lzss(plain, input, size);

for (int i=0; i<23; i++)
    klee_assume (plain[i]=="Buffalo buffalo Buffalo"[i]);

klee_assert(0);

return 0;
}

```

...but then KLEE works much slower, consumes much more RAM and I had success only with even smaller pieces of desired text.

So how **LZSS** works? Without peeking into Wikipedia, we can say that: if **LZSS** compressor observes some data it already had, it replaces the data with a link to some place in past with size. If it observes something yet unseen, it puts data as is. This is theory. This is indeed what we've got. Desired text is three "Buffalo" words, the first and the last are equivalent, but the second is *almost* equivalent, differing with first by one character.

That's what we see:

```
'\xffBuffalo \x01b\x0f\x03\r\x05'
```

Here is some control byte (0xff), "Buffalo" word is placed *as is*, then another control byte (0x01), then we see beginning of the second word ("b") and more control bytes, perhaps, links to the beginning of the buffer. These are command to decompressor, like, in plain English, "copy data from the buffer we've already done, from that place to that place", etc.

Interesting, is it possible to meddle into this piece of compressed data? Out of whim, can we force KLEE to find a compressed data, where not just "b" character has been placed *as is*, but also the second character of the word, i.e., "bu"?

I've modified main() function by adding `klee_assume()`: now the 11th byte of input (compressed) data (right after "b" byte) must have "u". I has no luck with 15 byte of compressed data, so I increased it to 16 bytes:

```

int main()
{
#define COMPRESSED_LEN 16
uint8_t input[COMPRESSED_LEN];
uint8_t plain[24];
uint32_t size=COMPRESSED_LEN;

klee_make_symbolic(input, sizeof input, "input");

klee_assume(input[11]=='u');

decompress_lzss(plain, input, size);

for (int i=0; i<23; i++)
    klee_assume (plain[i]=="Buffalo buffalo Buffalo"[i]);

klee_assert(0);

return 0;
}

```

...and voilà: KLEE found a compressed piece of data which satisfied our whimsical constraint:

```

% time klee klee_lzss.bc
KLEE: output directory is "/home/klee/klee-out-9"
KLEE: WARNING: undefined reference to function: klee_assert
KLEE: ERROR: /home/klee/klee_lzss.c:97: invalid klee_assume call (provably false)
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_lzss.c:47: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
KLEE: ERROR: /home/klee/klee_lzss.c:37: memory error: out of bound pointer
KLEE: NOTE: now ignoring this error at this location
KLEE: WARNING ONCE: calling external: klee_assert(0)
KLEE: ERROR: /home/klee/klee_lzss.c:99: failed external call: klee_assert

```

```
KLEE: NOTE: now ignoring this error at this location
```

```
KLEE: done: total instructions = 36700587
```

```
KLEE: done: completed paths = 369756
```

```
KLEE: done: generated tests = 4
```

```
real    12m16.983s
```

```
user    11m17.492s
```

```
sys     0m58.358s
```

```
% ktest-tool --write-ints klee-last/test000004.ktest
```

```
ktest file : 'klee-last/test000004.ktest'
```

```
args      : ['klee_lzss.bc']
```

```
num objects: 1
```

```
object   0: name: b'input'
```

```
object   0: size: 16
```

```
object   0: data: b'\xffBuffalo \x13bu\x10\x02\r\x05'
```

So now we find a piece of compressed data where two strings are placed as *is*: “Buffalo” and “bu”.

```
'\xffBuffalo \x13bu\x10\x02\r\x05'
```

Both pieces of compressed data, if feeded into our cypasted function, produce “Buffalo buffalo Buffalo” text string.

Please note, I still have no access to LZSS compressor code, and I didn’t get into LZSS decompressor details yet.

Unfortunately, things are not that cool: KLEE is very slow and I had success only with small pieces of text, and also ring buffer size had to be decreased significantly (original LZSS decompressor with ring buffer of 4096 bytes cannot decompress correctly what we found).

Nevertheless, it’s very impressive, taking into account the fact that we’re not getting into internals of this specific LZSS decompressor. Once more time, we’ve created *antipode* of decompressor, or *inverse function*.

Also, as it seems, KLEE isn’t very good so far with decompression algorithms (but who’s good then?). I’ve also tried various JPEG/PNG/GIF decoders (which, of course, has decompressors), starting with simplest possible, and KLEE had stuck.

## 9.11 strtodx() from RetroBSD

Just found this function in RetroBSD: <https://github.com/RetroBSD/retrobsd/blob/master/src/libc/stdlib/strtod.c>. It converts a string into floating point number for given radix.

```
1 #include <stdio.h>
2
3 // my own version, only for radix 10:
4 int isdigitx (char c, int radix)
5 {
6     if (c>='0' && c<='9')
7         return 1;
8     return 0;
9 };
10
11 /*
12 * double strtodx (char *string, char **endPtr, int radix)
13 *     This procedure converts a floating-point number from an ASCII
14 *     decimal representation to internal double-precision format.
15 *
16 * Original sources taken from 386bsd and modified for variable radix
17 * by Serge Vakulenko, <vak@kiae.su>.
18 *
19 * Arguments:
20 * string
21 *     A decimal ASCII floating-point number, optionally preceded
22 *     by white space. Must have form "-I.FE-X", where I is the integer
23 *     part of the mantissa, F is the fractional part of the mantissa,
24 *     and X is the exponent. Either of the signs may be "+", "-", or
25 *     omitted. Either I or F may be omitted, or both. The decimal point
26 *     isn't necessary unless F is present. The "E" may actually be an "e",
27 *     or "E", "S", "s", "F", "f", "D", "d", "L", "l".
28 *     E and X may both be omitted (but not just one).
29 *
30 * endPtr
31 *     If non-NULL, store terminating character's address here.
```

```

32 *
33 * radix
34 * Radix of floating point, one of 2, 8, 10, 16.
35 *
36 * The return value is the double-precision floating-point
37 * representation of the characters in string. If endPtr isn't
38 * NULL, then *endPtr is filled in with the address of the
39 * next character after the last one that was part of the
40 * floating-point number.
41 */
42 double strtodx (char *string, char **endPtr, int radix)
43 {
44     int sign = 0, expSign = 0, fracSz, fracOff, i;
45     double fraction, dblExp, *powTab;
46     register char *p;
47     register char c;
48
49     /* Exponent read from "EX" field. */
50     int exp = 0;
51
52     /* Exponent that derives from the fractional part. Under normal
53     * circumstances, it is the negative of the number of digits in F.
54     * However, if I is very long, the last digits of I get dropped
55     * (otherwise a long I with a large negative exponent could cause an
56     * unnecessary overflow on I alone). In this case, fracExp is
57     * incremented one for each dropped digit. */
58     int fracExp = 0;
59
60     /* Number of digits in mantissa. */
61     int mantSize;
62
63     /* Number of mantissa digits BEFORE decimal point. */
64     int decPt;
65
66     /* Temporarily holds location of exponent in string. */
67     char *pExp;
68
69     /* Largest possible base 10 exponent.
70     * Any exponent larger than this will already
71     * produce underflow or overflow, so there's
72     * no need to worry about additional digits. */
73     static int maxExponent = 307;
74
75     /* Table giving binary powers of 10.
76     * Entry is 102i. Used to convert decimal
77     * exponents into floating-point numbers. */
78     static double powersOf10[] = {
79         1e1, 1e2, 1e4, 1e8, 1e16, 1e32, //1e64, 1e128, 1e256,
80     };
81     static double powersOf2[] = {
82         2, 4, 16, 256, 65536, 4.294967296e9, 1.8446744073709551616e19,
83         //3.4028236692093846346e38, 1.1579208923731619542e77, 1.3407807929942597099e154,
84     };
85     static double powersOf8[] = {
86         8, 64, 4096, 2.81474976710656e14, 7.9228162514264337593e28,
87         //6.2771017353866807638e57, 3.9402006196394479212e115, 1.5525180923007089351e231,
88     };
89     static double powersOf16[] = {
90         16, 256, 65536, 1.8446744073709551616e19,
91         //3.4028236692093846346e38, 1.1579208923731619542e77, 1.3407807929942597099e154,
92     };
93
94     /*
95     * Strip off leading blanks and check for a sign.
96     */
97     p = string;
98     while (*p == ' ' || *p == '\t')
99         ++p;
100     if (*p == '-') {
101         sign = 1;
102         ++p;
103     } else if (*p == '+')
104         ++p;
105
106     /*
107     * Count the number of digits in the mantissa (including the decimal

```

```

108     * point), and also locate the decimal point.
109     */
110     decPt = -1;
111     for (mantSize=0; ; ++mantSize) {
112         c = *p;
113         if (!isdigitx (c, radix)) {
114             if (c != '.' || decPt >= 0)
115                 break;
116             decPt = mantSize;
117         }
118         ++p;
119     }
120
121     /*
122     * Now suck up the digits in the mantissa. Use two integers to
123     * collect 9 digits each (this is faster than using floating-point).
124     * If the mantissa has more than 18 digits, ignore the extras, since
125     * they can't affect the value anyway.
126     */
127     pExp = p;
128     p -= mantSize;
129     if (decPt < 0)
130         decPt = mantSize;
131     else
132         --mantSize;          /* One of the digits was the point. */
133
134     switch (radix) {
135     default:
136     case 10: fracSz = 9;  fracOff = 1000000000; powTab = powersOf10; break;
137     case 2:  fracSz = 30; fracOff = 1073741824; powTab = powersOf2;  break;
138     case 8:  fracSz = 10; fracOff = 1073741824; powTab = powersOf8;  break;
139     case 16: fracSz = 7;  fracOff = 268435456;  powTab = powersOf16; break;
140     }
141     if (mantSize > 2 * fracSz)
142         mantSize = 2 * fracSz;
143     fracExp = decPt - mantSize;
144     if (mantSize == 0) {
145         fraction = 0.0;
146         p = string;
147         goto done;
148     } else {
149         int frac1, frac2;
150
151         for (frac1=0; mantSize>fracSz; --mantSize) {
152             c = *p++;
153             if (c == '.')
154                 c = *p++;
155             frac1 = frac1 * radix + (c - '0');
156         }
157         for (frac2=0; mantSize>0; --mantSize) {
158             c = *p++;
159             if (c == '.')
160                 c = *p++;
161             frac2 = frac2 * radix + (c - '0');
162         }
163         fraction = (double) fracOff * frac1 + frac2;
164     }
165
166     /*
167     * Skim off the exponent.
168     */
169     p = pExp;
170     if (*p=='E' || *p=='e' || *p=='S' || *p=='s' || *p=='F' || *p=='f' ||
171         *p=='D' || *p=='d' || *p=='L' || *p=='l') {
172         ++p;
173         if (*p == '-') {
174             expSign = 1;
175             ++p;
176         } else if (*p == '+')
177             ++p;
178         while (isdigitx (*p, radix))
179             exp = exp * radix + (*p++ - '0');
180     }
181     if (expSign)
182         exp = fracExp - exp;
183     else

```



```

184         exp = fracExp + exp;
185
186     /*
187     * Generate a floating-point number that represents the exponent.
188     * Do this by processing the exponent one bit at a time to combine
189     * many powers of 2 of 10. Then combine the exponent with the
190     * fraction.
191     */
192     if (exp < 0) {
193         expSign = 1;
194         exp = -exp;
195     } else
196         expSign = 0;
197     if (exp > maxExponent)
198         exp = maxExponent;
199     dblExp = 1.0;
200     for (i=0; exp; exp>>=1, ++i)
201         if (exp & 01)
202             dblExp *= powTab[i];
203     if (expSign)
204         fraction /= dblExp;
205     else
206         fraction *= dblExp;
207
208 done:
209     if (endPtr)
210         *endPtr = p;
211
212     return sign ? -fraction : fraction;
213 }
214
215 #define BUFSIZE 10
216 int main()
217 {
218     char buf[BUFSIZE];
219     klee_make_symbolic (buf, sizeof buf, "buf");
220     klee_assume(buf[9]==0);
221
222     strtodx (buf, NULL, 10);
223 };

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/KLEE/strtodx.c](https://github.com/dennis714/SAT_SMT_article/blob/master/KLEE/strtodx.c) )

Interestingly, KLEE cannot handle floating-point arithmetic, but nevertheless, found something:

```

...
KLEE: ERROR: /home/klee/klee_test.c:202: memory error: out of bound pointer
...
% ktest-tool klee-last/test003483.ktest
ktest file : 'klee-last/test003483.ktest'
args       : ['klee_test.bc']
num objects: 1
object  0: name: b'buf'
object  0: size: 10
object  0: data: b'-.0E-66\x00\x00\x00'

```

As it seems, string “-.0E-66” makes out of bounds array access (read) at line 202. While further investigation, I’ve found that `powersOf10[]` array is too short: 6th element (started at 0th) has been accessed. And here we see part of array commented (line 79)! Probably someone’s mistake?

## 9.12 Unit testing: simple expression evaluator (calculator)

I has been looking for simple expression evaluator (calculator in other words) which takes expression like “2+2” on input and gives answer. I’ve found one at <http://stackoverflow.com/a/13895198>. Unfortunately, it has no bugs, so I’ve introduced one: a token buffer ( `buf[]` at line 31) is smaller than input buffer ( `input[]` at line 19).

```

1 // copied from http://stackoverflow.com/a/13895198 and reworked
2
3 // Bare bones scanner and parser for the following LL(1) grammar:

```

```

4 // expr -> term { [+ -] term } ; An expression is terms separated by add ops.
5 // term -> factor { [* /] factor } ; A term is factors separated by mul ops.
6 // factor -> unsigned_factor ; A signed factor is a factor,
7 // | - unsigned_factor ; possibly with leading minus sign
8 // unsigned_factor -> ( expr ) ; An unsigned factor is a parenthesized expression
9 // | NUMBER ; or a number
10 //
11 // The parser returns the floating point value of the expression.
12
13 #include <string.h>
14 #include <stdio.h>
15 #include <stdlib.h>
16 #include <stdint.h>
17 #include <stdbool.h>
18
19 char input[128];
20 int input_idx=0;
21
22 char my_getchar()
23 {
24     char rt=input[input_idx];
25     input_idx++;
26     return rt;
27 };
28
29 // The token buffer. We never check for overflow! Do so in production code.
30 // it's deliberately smaller than input[] so KLEE could find buffer overflow
31 char buf[64];
32 int n = 0;
33
34 // The current character.
35 int ch;
36
37 // The look-ahead token. This is the 1 in LL(1).
38 enum { ADD_OP, MUL_OP, LEFT_PAREN, RIGHT_PAREN, NOT_OP, NUMBER, END_INPUT } look_ahead;
39
40 // Forward declarations.
41 void init(void);
42 void advance(void);
43 int expr(void);
44 void error(char *msg);
45
46 // Parse expressions, one per line.
47 int main(void)
48 {
49     // take input expression from input[]
50     //input[0]=0;
51     //strcpy (input, "2+2");
52     klee_make_symbolic(input, sizeof input, "input");
53     input[127]=0;
54
55     init();
56     while (1)
57     {
58         int val = expr();
59         printf("%d\n", val);
60
61         if (look_ahead != END_INPUT)
62             error("junk after expression");
63         advance(); // past end of input mark
64     }
65     return 0;
66 }
67
68 // Just die on any error.
69 void error(char *msg)
70 {
71     fprintf(stderr, "Error: %s. Exiting.\n", msg);
72     exit(1);
73 }
74
75 // Buffer the current character and read a new one.
76 void read()
77 {
78     buf[n++] = ch;
79     buf[n] = '\0'; // Terminate the string.

```

```

80     ch = my_getchar();
81 }
82
83 // Ignore the current character.
84 void ignore()
85 {
86     ch = my_getchar();
87 }
88
89 // Reset the token buffer.
90 void reset()
91 {
92     n = 0;
93     buf[0] = '\0';
94 }
95
96 // The scanner. A tiny deterministic finite automaton.
97 int scan()
98 {
99     reset();
100 START:
101     // first character is digit?
102     if (isdigit (ch))
103         goto DIGITS;
104
105     switch (ch)
106     {
107         case ' ': case '\t': case '\r':
108             ignore();
109             goto START;
110
111         case '-': case '+': case '^':
112             read();
113             return ADD_OP;
114
115         case '~':
116             read();
117             return NOT_OP;
118
119         case '*': case '/': case '%':
120             read();
121             return MUL_OP;
122
123         case '(':
124             read();
125             return LEFT_PAREN;
126
127         case ')':
128             read();
129             return RIGHT_PAREN;
130
131         case 0:
132         case '\n':
133             ch = ' '; // delayed ignore()
134             return END_INPUT;
135
136         default:
137             printf ("bad character: 0x%x\n", ch);
138             exit(0);
139     }
140
141 DIGITS:
142     if (isdigit (ch))
143     {
144         read();
145         goto DIGITS;
146     }
147     else
148         return NUMBER;
149 }
150
151 // To advance is just to replace the look-ahead.
152 void advance()
153 {
154     look_ahead = scan();
155 }

```

```

156
157 // Clear the token buffer and read the first look-ahead.
158 void init()
159 {
160     reset();
161     ignore(); // junk current character
162     advance();
163 }
164
165 int get_number(char *buf)
166 {
167     char *endptr;
168
169     int rt=strtoul (buf, &endptr, 10);
170
171     // is the whole buffer has been processed?
172     if (strlen(buf)!=endptr-buf)
173     {
174         fprintf (stderr, "invalid number: %s\n", buf);
175         exit(0);
176     };
177     return rt;
178 };
179
180 int unsigned_factor()
181 {
182     int rtn = 0;
183     switch (look_ahead)
184     {
185         case NUMBER:
186             rtn=get_number(buf);
187             advance();
188             break;
189
190         case LEFT_PAREN:
191             advance();
192             rtn = expr();
193             if (look_ahead != RIGHT_PAREN) error("missing ')");
194             advance();
195             break;
196
197         default:
198             printf("unexpected token: %d\n", look_ahead);
199             exit(0);
200     }
201     return rtn;
202 }
203
204 int factor()
205 {
206     int rtn = 0;
207     // If there is a leading minus...
208     if (look_ahead == ADD_OP && buf[0] == '-')
209     {
210         advance();
211         rtn = -unsigned_factor();
212     }
213     else
214         rtn = unsigned_factor();
215
216     return rtn;
217 }
218
219 int term()
220 {
221     int rtn = factor();
222     while (look_ahead == MUL_OP)
223     {
224         switch(buf[0])
225         {
226             case '*':
227                 advance();
228                 rtn *= factor();
229                 break;
230
231             case '/':

```



Maybe this is not impressive result, nevertheless, it's yet another reminder that division and remainder operations must be wrapped somehow in production code to avoid possible crash.

## 9.13 Regular expressions

I've always wanted to generate possible strings for given regular expression. This is not so hard if to dive into regular expression matcher theory and details, but can we force RE matcher to do this?

I took lightest RE engine I've found: <https://github.com/cesanta/slre>, and wrote this:

```
int main(void)
{
    char s[6];
    klee_make_symbolic(s, sizeof s, "s");
    s[5]=0;
    if (slre_match("^\\d[a-c]+(x|y|z)", s, 5, NULL, 0, 0)==5)
        klee_assert(0);
}
```

So I wanted a string consisting of digit, "a" or "b" or "c" (at least one character) and "x" or "y" or "z" (one character). The whole string must have size of 5 characters.

```
% klee --libc=uclibc slre.bc
...
KLEE: ERROR: /home/klee/slre.c:445: failed external call: klee_assert
KLEE: NOTE: now ignoring this error at this location
...

% ls klee-last | grep err
test000014.external.err

% ktest-tool --write-ints klee-last/test000014.ktest
ktest file : 'klee-last/test000014.ktest'
args      : ['slre.bc']
num objects: 1
object   0: name: b's'
object   0: size: 6
object   0: data: b'5aaax\xff'
```

This is indeed correct string. "\xff" is at the place where terminal zero byte should be, but RE engine we use ignores the last zero byte, because it has buffer length as a passed parameter. Hence, KLEE doesn't *reconstruct* final byte.

Can we get more? Now we add additional constraint:

```
int main(void)
{
    char s[6];
    klee_make_symbolic(s, sizeof s, "s");
    s[5]=0;
    if (slre_match("^\\d[a-c]+(x|y|z)", s, 5, NULL, 0, 0)==5 &&
        strcmp(s, "5aaax")!=0)
        klee_assert(0);
}
```

```
% ktest-tool --write-ints klee-last/test000014.ktest
ktest file : 'klee-last/test000014.ktest'
args      : ['slre.bc']
num objects: 1
object   0: name: b's'
object   0: size: 6
object   0: data: b'7aaax\xff'
```

Let's say, out of whim, we don't like "a" at the 2nd position (starting at 0th):

```
int main(void)
{
    char s[6];
    klee_make_symbolic(s, sizeof s, "s");
    s[5]=0;
    if (slre_match("^\\d[a-c]+(x|y|z)", s, 5, NULL, 0, 0)==5 &&
        strcmp(s, "5aaax")!=0 &&
        s[2]!='a')
        klee_assert(0);
}
```

KLEE found a way to satisfy our new constraint:

```
% ktest-tool --write-ints klee-last/test000014.ktest
ktest file : 'klee-last/test000014.ktest'
args      : ['slre.bc']
num objects: 1
object   0: name: b's'
object   0: size: 6
object   0: data: b'7abax\xff'
```

Let's also define constraint KLEE cannot satisfy:

```
int main(void)
{
    char s[6];
    klee_make_symbolic(s, sizeof s, "s");
    s[5]=0;
    if (slre_match("^\\d[a-c]+(x|y|z)", s, 5, NULL, 0, 0)==5 &&
        strcmp(s, "5aaax")!=0 &&
        s[2]!='a' &&
        s[2]!='b' &&
        s[2]!='c')
        klee_assert(0);
}
```

It cannot indeed, and KLEE finished without reporting about `klee_assert()` triggering.

## 9.14 Exercise

Here is my crackme/keygenme, which may be tricky, but easy to solve using KLEE: <http://challenges.re/74/>.

# 10 (Amateur) cryptography

## 10.1 Serious cryptography

Let's back to the method we previously used (8.2) to construct expressions using running Python function. We can try to build expression for the output of XXTEA encryption algorithm:

```
#!/usr/bin/env python

class Expr:
    def __init__(self,s):
        self.s=s

    def __str__(self):
        return self.s

    def convert_to_Expr_if_int(self, n):
        if isinstance(n, int):
            return Expr(str(n))
        if isinstance(n, Expr):
            return n
        raise AssertionError # unsupported type

    def __xor__(self, other):
        return Expr("(" + self.s + "^" + self.convert_to_Expr_if_int(other).s + ")")

    def __mul__(self, other):
        return Expr("(" + self.s + "*" + self.convert_to_Expr_if_int(other).s + ")")

    def __add__(self, other):
        return Expr("(" + self.s + "+" + self.convert_to_Expr_if_int(other).s + ")")

    def __and__(self, other):
        return Expr("(" + self.s + "&" + self.convert_to_Expr_if_int(other).s + ")")

    def __lshift__(self, other):
        return Expr("(" + self.s + "<<" + self.convert_to_Expr_if_int(other).s + ")")

    def __rshift__(self, other):
```

```

    return Expr("(" + self.s + ">>" + self.convert_to_Expr_if_int(other).s + ")")

def __getitem__(self, d):
    return Expr("(" + self.s + "[" + d.s + "]")")

# reworked from:

# Pure Python (2.x) implementation of the XXTEA cipher
# (c) 2009. Ivan Voras <ivoras@gmail.com>
# Released under the BSD License.

def raw_xxtea(v, n, k):

    def MX():
        return ((z>>5)^(y<<2)) + ((y>>3)^(z<<4))^(sum^y) + (k[(Expr(str(p)) & 3)^e]^z)

    y = v[0]
    sum = Expr("0")
    DELTA = 0x9e3779b9
    # Encoding only
    z = v[n-1]

    # number of rounds:
    #q = 6 + 52 / n
    q=1

    while q > 0:
        q -= 1
        sum = sum + DELTA
        e = (sum >> 2) & 3
        p = 0
        while p < n - 1:
            y = v[p+1]
            z = v[p] = v[p] + MX()
            p += 1
        y = v[0]
        z = v[n-1] = v[n-1] + MX()
    return 0

v=[Expr("input1"), Expr("input2"), Expr("input3"), Expr("input4")]
k=Expr("key")

raw_xxtea(v, 4, k)

for i in range(4):
    print i, ":", v[i]
#print len(str(v[0]))+len(str(v[1]))+len(str(v[2]))+len(str(v[3]))

```

A key is chosen according to input data, and, obviously, we can't know it during symbolic execution, so we leave expression like `k[...]`.

Now results for just one round, for each of 4 outputs:

```

0 : (input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^input2)+
((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))

1 : (input2+((((input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^
input2)+((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))>>5)^(input3<<2))+((input3>>3)^(input1+
((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^input2)+((key[((0&3)^
((0+2654435769)>>2)&3]))^input4))))<<4)))^(((0+2654435769)^input3)+((key[((1&3)^((0+2654435769)>>2)&
3]))^input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^input2)+
((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))))

2 : (input3+((((input2+((((input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^
(((0+2654435769)^input2)+((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))>>5)^(input3<<2))+
((input3>>3)^(input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^
input2)+((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))<<4)))^(((0+2654435769)^input3)+
((key[((1&3)^((0+2654435769)>>2)&3]))^input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))
^(((0+2654435769)^input2)+((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))>>5)^(input4<<2))+
((input4>>3)^(input2+((((input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^
(((0+2654435769)^input2)+((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))>>3)^(input4<<4)))^
(((0+2654435769)^input2)+((key[((0&3)^((0+2654435769)>>2)&3]))^input4))))>>5)^(input3<<2))+
((input3>>3)^(input1+
((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^input2)+((key[((0&3)^
((0+2654435769)>>2)&3]))^input4))))<<4)))^(((0+2654435769)^input3)+((key[((1&3)^((0+2654435769)>>2)&3]))^
input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^input2)+((key[((0&3)^
((0+2654435769)>>2)&3]))^input4))))))<<4)))^(((0+2654435769)^input4)+((key[((2&3)^((0+2654435769)>>2)&
3]))^input2+((((input1+((((input4>>5)^(input2<<2))+((input2>>3)^(input4<<4)))^(((0+2654435769)^

```





In other words, theoretically, you can build system of equation like this:  $MD5(x) = 12341234\dots$ , but expressions are so huge so it's impossible to solve them. Yes, cryptographers are fully aware of this and one of the goals of the successful cipher is to make expressions as big as possible, using reasonable time and size of algorithm.

Nevertheless, you can find numerous papers about breaking these cryptosystems with reduced number of rounds: when expression isn't *exploded* yet, sometimes it's possible. This cannot be applied in practice, but such experience has some interesting theoretical results.

### 10.1.1 Attempts to break “serious” crypto

CryptoMiniSat itself exist to support XOR operation, which is ubiquitous in cryptography.

- Bitcoin mining with SAT solver: <http://jheusser.github.io/2013/02/03/satcoin.html>, <https://github.com/msoos/sha256-sat-bitcoin>.
- [Alexander Semenov, attempts to break A5/1, etc. \(Russian presentation\)](#)
- [Vegard Nossum - SAT-based preimage attacks on SHA-1](#)
- [Algebraic Attacks on the Crypto-1 Stream Cipher in MiFare Classic and Oyster Cards](#)
- [Attacking Bivium Using SAT Solvers](#)
- [Extending SAT Solvers to Cryptographic Problems](#)
- [Applications of SAT Solvers to Cryptanalysis of Hash Functions](#)
- [Algebraic-Differential Cryptanalysis of DES](#)

## 10.2 Amateur cryptography

This is what you can find in serial numbers, license keys, executable file packers, [CTF<sup>80</sup>](#), malware, etc. Sometimes even ransomware (but rarely nowadays, in 2017).

Amateur cryptography is often can be broken using SMT solver, or even KLEE.

Amateur cryptography is usually based not on theory, but on visual complexity: if its creator getting results which are seems chaotic enough, often, one stops to improve it further. This is security not even on obscurity, but on chaotic mess. This is also sometimes called “The Fallacy of Complex Manipulation” (see [RFC4086](#)).

Devising your own cryptoalgorithm is a very tricky thing to do. This can be compared to devising your own [PRNG](#). Even famous Donald Knuth in 1959 constructed one, and it was visually very complex, but, as it turns out in practice, it has very short cycle of length 3178. [See also: The Art of Computer Programming vol.II page 4, (1997).]

The very first problem is that making an algorithm which can generate very long expressions is tricky thing itself. Common error is to use operations like XOR and rotations/permutations, which can't help much. Even worse: some people think that XORing a value several times can be better, like:  $(x \oplus 1234) \oplus 5678$ . Obviously, these two XOR operations (or more precisely, any number of it) can be reduced to a single one. Same story about applied operations like addition and subtraction—they all also can be reduced to single one.

Real cryptoalgorithms, like IDEA, can use several operations from different groups, like XOR, addition and multiplication. Applying them all in specific order will make resulting expression irreducible.

When I prepared this part, I tried to make an example of such amateur hash function:

```
// copied from http://blog.regehr.org/archives/1063
uint32_t rotl32b (uint32_t x, uint32_t n)
{
    assert (n<32);
    if (!n) return x;
    return (x<<n) | (x>>(32-n));
}

uint32_t rotr32b (uint32_t x, uint32_t n)
{
    assert (n<32);
    if (!n) return x;
    return (x>>n) | (x<<(32-n));
}
```

<sup>80</sup>Capture the Flag

```

void megahash (uint32_t buf[4])
{
    for (int i=0; i<4; i++)
    {
        uint32_t t0=buf[0]^0x12345678^buf[1];
        uint32_t t1=buf[1]^0xabcdef01^buf[2];
        uint32_t t2=buf[2]^0x23456789^buf[3];
        uint32_t t3=buf[3]^0x0abcdef0^buf[0];

        buf[0]=rotl32b(t0, 1);
        buf[1]=rotr32b(t1, 2);
        buf[2]=rotl32b(t2, 3);
        buf[3]=rotr32b(t3, 4);
    };
};

int main()
{
    uint32_t buf[4];
    klee_make_symbolic(buf, sizeof buf);
    megahash (buf);
    if (buf[0]==0x18f71ce6 // or whatever
        && buf[1]==0xf37c2fc9
        && buf[2]==0x1cfe96fe
        && buf[3]==0x8c02c75e)
        klee_assert(0);
};

```

KLEE can break it with little effort. Functions of such complexity is common in shareware, which checks license keys, etc.

Here is how we can make its work harder by making rotations dependent of inputs, and this makes number of possible inputs much, much bigger:

```

void megahash (uint32_t buf[4])
{
    for (int i=0; i<16; i++)
    {
        uint32_t t0=buf[0]^0x12345678^buf[1];
        uint32_t t1=buf[1]^0xabcdef01^buf[2];
        uint32_t t2=buf[2]^0x23456789^buf[3];
        uint32_t t3=buf[3]^0x0abcdef0^buf[0];

        buf[0]=rotl32b(t0, t1&0x1F);
        buf[1]=rotr32b(t1, t2&0x1F);
        buf[2]=rotl32b(t2, t3&0x1F);
        buf[3]=rotr32b(t3, t0&0x1F);
    };
};

```

Addition (or [modular addition](#), as cryptographers say) can make thing even harder:

```

void megahash (uint32_t buf[4])
{
    for (int i=0; i<4; i++)
    {
        uint32_t t0=buf[0]^0x12345678^buf[1];
        uint32_t t1=buf[1]^0xabcdef01^buf[2];
        uint32_t t2=buf[2]^0x23456789^buf[3];
        uint32_t t3=buf[3]^0x0abcdef0^buf[0];

        buf[0]=rotl32b(t0, t2&0x1F)+t1;
        buf[1]=rotr32b(t1, t3&0x1F)+t2;
        buf[2]=rotl32b(t2, t1&0x1F)+t3;
        buf[3]=rotr32b(t3, t2&0x1F)+t4;
    };
};

```

As an exercise, you can try to make a block cipher which KLEE wouldn't break. This is quite sobering experience. But even if you can, this is not a panacea, there is an array of other cryptoanalytical methods to break it.

Summary: if you deal with amateur cryptography, you can always give KLEE and SMT solver a try. Even more: sometimes you have only decryption function, and if algorithm is simple enough, KLEE or SMT solver can reverse things back.

One fun thing to mention: if you try to implement amateur cryptoalgorithm in Verilog/VHDL language to run it on [FPGA<sup>81</sup>](#), maybe in brute-force way, you can find that [EDA<sup>82</sup>](#) tools can optimize many things during synthesis (this is the word they use for “compilation”) and can leave cryptoalgorithm much smaller/simpler than it was. Even if you try to implement DES algorithm *in bare metal* with a fixed key, Altera Quartus can optimize first round of it and make it smaller than others.

### 10.2.1 Bugs

Another prominent feature of amateur cryptography is bugs. Bugs here often left uncaught because output of encrypting function visually looked “good enough” or “obfuscated enough”, so a developer stopped to work on it.

This is especially feature of hash functions, because when you work on block cipher, you have to do two functions (encryption/decryption), while hashing function is single.

Weirdest ever amateur encryption algorithm I once saw, encrypted only odd bytes of input block, while even bytes left untouched, so the input plain text has been partially seen in the resulting encrypted block. It was encryption routine used in license key validation. Hard to believe someone did this on purpose. Most likely, it was just an unnoticed bug.

### 10.2.2 XOR ciphers

Simplest possible amateur cryptography is just application of XOR operation using some kind of table. Maybe even simpler. This is a real algorithm I once saw:

```
for (i=0; i<size; i++)
    buf[i]=buf[i]^(31*(i+1));
```

This is not even encryption, rather concealing or hiding.

### 10.2.3 Other features

**Tables** There are often table(s) with pseudorandom data, which is/are used chaotically.

**Checksumming** End-users can have proclivity to changing license codes, serial numbers, etc., with a hope this could affect behaviour of software. So there is often some kind of checksum: starting at simple summing and [CRC](#). This is close to [MAC<sup>83</sup>](#) in real cryptography.

### 10.2.4 Examples

- A popular FLEXIm license manager was based on a simple amateur cryptoalgorithm (before they switched to [ECC<sup>84</sup>](#)), which can be cracked easily.
- Pegasus Mail Password Decoder: <http://phrack.org/issues/52/3.html> - a very typical example.
- You can find a lot of blog posts about breaking CTF-level crypto using Z3, etc. Here is one of them: <http://doar-e.github.io/blog/2015/08/18/keygenning-with-klée/>.
- Another: [Automated algebraic cryptanalysis with OpenREIL and Z3](#). By the way, this solution tracks state of each register at each EIP/RIP, this is almost the same as [SSA](#), which is heavily used in compilers and worth learning.
- Many examples of amateur cryptography I've taken from an old Fravia website: [https://yurichev.com/mirrors/amateur\\_crypto\\_examples\\_from\\_Fravia/](https://yurichev.com/mirrors/amateur_crypto_examples_from_Fravia/).

<sup>81</sup>Field-programmable gate array

<sup>82</sup>Electronic design automation

<sup>83</sup>Message authentication code

<sup>84</sup>Elliptic curve cryptography

## 10.3 Case study: simple hash function

(This piece of text was initially added to my “Reverse Engineering for Beginners” book ([beginners.re](http://beginners.re)) at March 2014) <sup>85</sup>.

Here is one-way hash function, that converted a 64-bit value to another and we need to try to reverse its flow back.

### 10.3.1 Manual decompiling

Here its assembly language listing in IDA:

```
sub_401510    proc near
              ; ECX = input
              mov     rdx, 5D7E0D1F2E0F1F84h
              mov     rax, rcx          ; input
              imul   rax, rdx
              mov     rdx, 388D76AEE8CB1500h
              mov     ecx, eax
              and     ecx, 0Fh
              ror     rax, cl
              xor     rax, rdx
              mov     rdx, 0D2E9EE7E83C4285Bh
              mov     ecx, eax
              and     ecx, 0Fh
              rol     rax, cl
              lea    r8, [rax+rdx]
              mov     rdx, 8888888888888889h
              mov     rax, r8
              mul     rdx
              shr     rdx, 5
              mov     rax, rdx
              lea    rcx, [r8+rdx*4]
              shl     rax, 6
              sub     rcx, rax
              mov     rax, r8
              rol     rax, cl
              ; EAX = output
              retn
sub_401510    endp
```

The example was compiled by GCC, so the first argument is passed in ECX.

If you don't have Hex-Rays, or if you distrust to it, you can try to reverse this code manually. One method is to represent the CPU registers as local C variables and replace each instruction by a one-line equivalent expression, like:

```
uint64_t f(uint64_t input)
{
    uint64_t rax, rbx, rcx, rdx, r8;

    ecx=input;

    rdx=0x5D7E0D1F2E0F1F84;
    rax=rcx;
    rax*=rdx;
    rdx=0x388D76AEE8CB1500;
    rax=_rotr(rax, rax&0xF); // rotate right
    rax^=rdx;
    rdx=0xD2E9EE7E83C4285B;
    rax=_rotl(rax, rax&0xF); // rotate left
    r8=rax+rdx;
    rdx=0x8888888888888889;
    rax=r8;
    rax*=rdx;
    rdx=rdx>>5;
    rax=rdx;
    rcx=r8+rdx*4;
    rax=rax<<6;
    rcx=rcx-rax;
    rax=r8
    rax=_rotl (rax, rcx&0xFF); // rotate left
```

<sup>85</sup>This example was also used by Murphy Berzish in his lecture about SAT and SMT: <http://mirror.csclub.uwaterloo.ca/csclub/mtrberzi-sat-smt-slides.pdf>, <http://mirror.csclub.uwaterloo.ca/csclub/mtrberzi-sat-smt.mp4>

```
};
    return rax;
};
```

If you are careful enough, this code can be compiled and will even work in the same way as the original. Then, we are going to rewrite it gradually, keeping in mind all registers usage. Attention and focus is very important here—any tiny typo may ruin all your work! Here is the first step:

```
uint64_t f(uint64_t input)
{
    uint64_t rax, rbx, rcx, rdx, r8;

    ecx=input;

    rdx=0x5D7E0D1F2E0F1F84;
    rax=rcx;
    rax*=rdx;
    rdx=0x388D76AEE8CB1500;
    rax=_rotr(rax, rax&0xF); // rotate right
    rax^=rdx;
    rdx=0xD2E9EE7E83C4285B;
    rax=_lrotl(rax, rax&0xF); // rotate left
    r8=rax+rdx;

    rdx=0x8888888888888889;
    rax=r8;
    rax*=rdx;
    // RDX here is a high part of multiplication result
    rdx=rdx>>5;
    // RDX here is division result!
    rax=rdx;

    rcx=r8+rdx*4;
    rax=rax<<6;
    rcx=rcx-rax;
    rax=r8
    rax=_lrotl (rax, rcx&0xFF); // rotate left
    return rax;
};
```

Next step:

```
uint64_t f(uint64_t input)
{
    uint64_t rax, rbx, rcx, rdx, r8;

    ecx=input;

    rdx=0x5D7E0D1F2E0F1F84;
    rax=rcx;
    rax*=rdx;
    rdx=0x388D76AEE8CB1500;
    rax=_rotr(rax, rax&0xF); // rotate right
    rax^=rdx;
    rdx=0xD2E9EE7E83C4285B;
    rax=_lrotl(rax, rax&0xF); // rotate left
    r8=rax+rdx;

    rdx=0x8888888888888889;
    rax=r8;
    rax*=rdx;
    // RDX here is a high part of multiplication result
    rdx=rdx>>5;
    // RDX here is division result!
    rax=rdx;

    rcx=(r8+rdx*4)-(rax<<6);
    rax=r8
    rax=_lrotl (rax, rcx&0xFF); // rotate left
    return rax;
};
```

We can spot the division using multiplication. Indeed, let's calculate the divider in Wolfram Mathematica:

Listing 1: Wolfram Mathematica

```
In[1]:=N[2^(64 + 5)/16^8888888888888889]
Out[1]:=60.
```

We get this:

```
uint64_t f(uint64_t input)
{
    uint64_t rax, rbx, rcx, rdx, r8;

    ecx=input;

    rdx=0x5D7E0D1F2E0F1F84;
    rax=rcx;
    rax*=rdx;
    rdx=0x388D76AEE8CB1500;
    rax=_rotr(rax, rax&0xF); // rotate right
    rax^=rdx;
    rdx=0xD2E9EE7E83C4285B;
    rax=_lrotl(rax, rax&0xF); // rotate left
    r8=rax+rdx;

    rax=rdx=r8/60;

    rcx=(r8+rax*4)-(rax*64);
    rax=r8
    rax=_lrotl (rax, rcx&0xFF); // rotate left
    return rax;
};
```

One more step:

```
uint64_t f(uint64_t input)
{
    uint64_t rax, rbx, rcx, rdx, r8;

    rax=input;
    rax*=0x5D7E0D1F2E0F1F84;
    rax=_rotr(rax, rax&0xF); // rotate right
    rax^=0x388D76AEE8CB1500;
    rax=_lrotl(rax, rax&0xF); // rotate left
    r8=rax+0xD2E9EE7E83C4285B;

    rcx=r8-(r8/60)*60;
    rax=r8
    rax=_lrotl (rax, rcx&0xFF); // rotate left
    return rax;
};
```

By simple reducing, we finally see that it's calculating the remainder, not the quotient:

```
uint64_t f(uint64_t input)
{
    uint64_t rax, rbx, rcx, rdx, r8;

    rax=input;
    rax*=0x5D7E0D1F2E0F1F84;
    rax=_rotr(rax, rax&0xF); // rotate right
    rax^=0x388D76AEE8CB1500;
    rax=_lrotl(rax, rax&0xF); // rotate left
    r8=rax+0xD2E9EE7E83C4285B;

    return _lrotl (r8, r8 % 60); // rotate left
};
```

We end up with this fancy formatted source-code:

```
#include <stdio.h>
#include <stdint.h>
#include <stdbool.h>
#include <string.h>
#include <intrin.h>

#define C1 0x5D7E0D1F2E0F1F84
#define C2 0x388D76AEE8CB1500
#define C3 0xD2E9EE7E83C4285B
```

```

uint64_t hash(uint64_t v)
{
    v*=C1;
    v=_rotr(v, v&0xF); // rotate right
    v^=C2;
    v=_lrotl(v, v&0xF); // rotate left
    v+=C3;
    v=_lrotl(v, v % 60); // rotate left
    return v;
};

int main()
{
    printf ("%llu\n", hash(...));
};

```

Since we are not cryptanalysts we can't find an easy way to generate the input value for some specific output value. The rotate instruction's coefficients look frightening—it's a warranty that the function is not bijective, it has collisions, or, speaking more simply, many inputs may be possible for one output.

Brute-force is not solution because values are 64-bit ones, that's beyond reality.

### 10.3.2 Now let's use the Z3

Still, without any special cryptographic knowledge, we may try to break this algorithm using Z3.

Here is the Python source code:

```

1 #!/usr/bin/env python
2
3 from z3 import *
4
5 C1=0x5D7E0D1F2E0F1F84
6 C2=0x388D76AEE8CB1500
7 C3=0xD2E9EE7E83C4285B
8
9 inp, i1, i2, i3, i4, i5, i6, outp = BitVecs('inp i1 i2 i3 i4 i5 i6 outp', 64)
10
11 s = Solver()
12 s.add(i1==inp*C1)
13 s.add(i2==RotateRight (i1, i1 & 0xF))
14 s.add(i3==i2 ^ C2)
15 s.add(i4==RotateLeft(i3, i3 & 0xF))
16 s.add(i5==i4 + C3)
17 s.add(outp==RotateLeft (i5, URem(i5, 60)))
18
19 s.add(outp==10816636949158156260)
20
21 print s.check()
22 m=s.model()
23 print m
24 print (" inp=0x%X" % m[inp].as_long())
25 print (" outp=0x%X" % m[outp].as_long())

```

This is going to be our first solver.

We see the variable definitions on line 7. These are just 64-bit variables. `i1..i6` are intermediate variables, representing the values in the registers between instruction executions.

Then we add the so-called constraints on lines 10..15. The last constraint at 17 is the most important one: we are going to try to find an input value for which our algorithm will produce 10816636949158156260.

*RotateRight*, *RotateLeft*, *URem*—are functions from the Z3 Python API, not related to Python language.

Then we run it:

```

...>python.exe 1.py
sat
[i1 = 3959740824832824396,
 i3 = 8957124831728646493,
 i5 = 10816636949158156260,
 inp = 1364123924608584563,
 outp = 10816636949158156260,
 i4 = 14065440378185297801,
 i2 = 4954926323707358301]
inp=0x12EE577B63E80B73
outp=0x961C69FF0AEFD7E4

```



“sat” mean “satisfiable”, i.e., the solver was able to find at least one solution. The solution is printed in the square brackets. The last two lines are the input/output pair in hexadecimal form. Yes, indeed, if we run our function with `0x12EE577B63E80B73` as input, the algorithm will produce the value we were looking for.

But, as we noticed before, the function we work with is not bijective, so there may be other correct input values. The Z3 is not capable of producing more than one result, but let’s hack our example slightly, by adding line 19, which implies “look for any other results than this”:

```

1  #!/usr/bin/env python
2
3  from z3 import *
4
5  C1=0x5D7E0D1F2E0F1F84
6  C2=0x388D76AEE8CB1500
7  C3=0xD2E9EE7E83C4285B
8
9  inp, i1, i2, i3, i4, i5, i6, outp = BitVecs('inp i1 i2 i3 i4 i5 i6 outp', 64)
10
11 s = Solver()
12 s.add(i1==inp*C1)
13 s.add(i2==RotateRight (i1, i1 & 0xF))
14 s.add(i3==i2 ^ C2)
15 s.add(i4==RotateLeft(i3, i3 & 0xF))
16 s.add(i5==i4 + C3)
17 s.add(outp==RotateLeft (i5, URem(i5, 60)))
18
19 s.add(outp==10816636949158156260)
20
21 s.add(inp!=0x12EE577B63E80B73)
22
23 print s.check()
24 m=s.model()
25 print m
26 print (" inp=0x%X" % m[inp].as_long())
27 print ("outp=0x%X" % m[outp].as_long())

```

Indeed, it finds another correct result:

```

...>python.exe 2.py
sat
[i1 = 3959740824832824396,
 i3 = 8957124831728646493,
 i5 = 10816636949158156260,
 inp = 10587495961463360371,
 outp = 10816636949158156260,
 i4 = 14065440378185297801,
 i2 = 4954926323707358301]
inp=0x92EE577B63E80B73
outp=0x961C69FF0AEFD7E4

```

This can be automated. Each found result can be added as a constraint and then the next result will be searched for. Here is a slightly more sophisticated example:

```

1  #!/usr/bin/env python
2
3  from z3 import *
4
5  C1=0x5D7E0D1F2E0F1F84
6  C2=0x388D76AEE8CB1500
7  C3=0xD2E9EE7E83C4285B
8
9  inp, i1, i2, i3, i4, i5, i6, outp = BitVecs('inp i1 i2 i3 i4 i5 i6 outp', 64)
10
11 s = Solver()
12 s.add(i1==inp*C1)
13 s.add(i2==RotateRight (i1, i1 & 0xF))
14 s.add(i3==i2 ^ C2)
15 s.add(i4==RotateLeft(i3, i3 & 0xF))
16 s.add(i5==i4 + C3)
17 s.add(outp==RotateLeft (i5, URem(i5, 60)))
18
19 s.add(outp==10816636949158156260)
20
21 # cypasted from http://stackoverflow.com/questions/11867611/z3py-checking-all-solutions-for-equation
22 result=[]

```

```

23 while True:
24     if s.check() == sat:
25         m = s.model()
26         print m[inp]
27         result.append(m)
28         # Create a new constraint the blocks the current model
29         block = []
30         for d in m:
31             # d is a declaration
32             if d.arity() > 0:
33                 raise Z3Exception("uninterpreted functions are not supported")
34             # create a constant from declaration
35             c=d()
36             if is_array(c) or c.sort().kind() == Z3_UNINTERPRETED_SORT:
37                 raise Z3Exception("arrays and uninterpreted sorts are not supported")
38             block.append(c != m[d])
39         s.add(Or(block))
40     else:
41         print "results total=",len(result)
42         break

```

We got:

```

1364123924608584563
1234567890
9223372038089343698
4611686019661955794
13835058056516731602
3096040143925676201
12319412180780452009
7707726162353064105
16931098199207839913
1906652839273745429
11130024876128521237
15741710894555909141
6518338857701133333
5975809943035972467
15199181979890748275
10587495961463360371
results total= 16

```

So there are 16 correct input values for `0x92EE577B63E80B73` as a result.

The second is 1234567890—it is indeed the value which was used by me originally while preparing this example.

Let's also try to research our algorithm a bit more. Acting on a sadistic whim, let's find if there are any possible input/output pairs in which the lower 32-bit parts are equal to each other?

Let's remove the *outp* constraint and add another, at line 17:

```

1  #!/usr/bin/env python
2
3  from z3 import *
4
5  C1=0x5D7E0D1F2E0F1F84
6  C2=0x388D76AEE8CB1500
7  C3=0xD2E9EE7E83C4285B
8
9  inp, i1, i2, i3, i4, i5, i6, outp = BitVecs('inp i1 i2 i3 i4 i5 i6 outp', 64)
10
11 s = Solver()
12 s.add(i1==inp*C1)
13 s.add(i2==RotateRight (i1, i1 & 0xF))
14 s.add(i3==i2 ^ C2)
15 s.add(i4==RotateLeft(i3, i3 & 0xF))
16 s.add(i5==i4 + C3)
17 s.add(outp==RotateLeft (i5, URem(i5, 60)))
18
19 s.add(outp & 0xFFFFFFFF == inp & 0xFFFFFFFF)
20
21 print s.check()
22 m=s.model()
23 print m
24 print (" inp=0x%X" % m[inp].as_long())
25 print ("outp=0x%X" % m[outp].as_long())

```

It is indeed so:

```

sat
[i1 = 14869545517796235860,
 i3 = 8388171335828825253,
 i5 = 6918262285561543945,
 inp = 1370377541658871093,
 outp = 14543180351754208565,
 i4 = 10167065714588685486,
 i2 = 5541032613289652645]
inp=0x13048F1D12C00535
outp=0xC9D3C17A12C00535

```

Let's be more sadistic and add another constraint: last 16 bits must be `0x1234` :

```

1  #!/usr/bin/env python
2
3  from z3 import *
4
5  C1=0x5D7E0D1F2E0F1F84
6  C2=0x388D76AEE8CB1500
7  C3=0xD2E9EE7E83C4285B
8
9  inp, i1, i2, i3, i4, i5, i6, outp = BitVecs('inp i1 i2 i3 i4 i5 i6 outp', 64)
10
11  s = Solver()
12  s.add(i1==inp*C1)
13  s.add(i2==RotateRight (i1, i1 & 0xF))
14  s.add(i3==i2 ^ C2)
15  s.add(i4==RotateLeft(i3, i3 & 0xF))
16  s.add(i5==i4 + C3)
17  s.add(outp==RotateLeft (i5, URem(i5, 60)))
18
19  s.add(outp & 0xFFFFFFFF == inp & 0xFFFFFFFF)
20  s.add(outp & 0xFFFF == 0x1234)
21
22  print s.check()
23  m=s.model()
24  print m
25  print (" inp=0x%X" % m[inp].as_long())
26  print ("outp=0x%X" % m[outp].as_long())

```

Oh yes, this possible as well:

```

sat
[i1 = 2834222860503985872,
 i3 = 2294680776671411152,
 i5 = 17492621421353821227,
 inp = 461881484695179828,
 outp = 419247225543463476,
 i4 = 2294680776671411152,
 i2 = 2834222860503985872]
inp=0x668EEC35F961234
outp=0x5D177215F961234

```

Z3 works very fast and it implies that the algorithm is weak, it is not cryptographic at all (like the most of the amateur cryptography).

## 11 SAT-solvers

SMT vs. SAT is like high level PL vs. assembly language. The latter can be much more efficient, but it's hard to program in it.

### 11.1 CNF form

CNF<sup>86</sup> is a *normal form*.

Any boolean expression can be converted to *normal form* and CNF is one of them. The CNF expression is a bunch of clauses (sub-expressions) consisting of terms (variables), ORs and NOTs, all of which are then

<sup>86</sup>[https://en.wikipedia.org/wiki/Conjunctive\\_normal\\_form](https://en.wikipedia.org/wiki/Conjunctive_normal_form)

glued together with AND into a full expression. There is a way to memorize it: CNF is “AND of ORs” (or “product of sums”) and DNF<sup>87</sup> is “OR of ANDs” (or “sum of products”).

Example is:  $(\neg A \vee B) \wedge (C \vee \neg D)$ .

$\vee$  stands for OR (logical disjunction<sup>88</sup>), “+” sign is also sometimes used for OR.

$\wedge$  stands for AND (logical conjunction<sup>89</sup>). It is easy to memorize:  $\wedge$  looks like “A” letter. “.” is also sometimes used for AND.

$\neg$  is negation (NOT).

## 11.2 Example: 2-bit adder

SAT-solver is merely a solver of huge boolean equations in CNF form. It just gives the answer, if there is a set of input values which can satisfy CNF expression, and what input values must be.

Here is a 2-bit adder for example:




Figure 9: 2-bit adder circuit

The adder in its simplest form: it has no carry-in and carry-out, and it has 3 XOR gates and one AND gate. Let’s try to figure out, which sets of input values will force adder to set both two output bits? By doing quick memory calculation, we can see that there are 4 ways to do so:  $0 + 3 = 3$ ,  $1 + 2 = 3$ ,  $2 + 1 = 3$ ,  $3 + 0 = 3$ . Here is also truth table, with these rows highlighted:

<sup>87</sup>Disjunctive normal form

<sup>88</sup>[https://en.wikipedia.org/wiki/Logical\\_disjunction](https://en.wikipedia.org/wiki/Logical_disjunction)

<sup>89</sup>[https://en.wikipedia.org/wiki/Logical\\_conjunction](https://en.wikipedia.org/wiki/Logical_conjunction)

	aH	aL	bH	bL	qH	qL
$3+3 = 6 \equiv 2 \pmod{4}$	1	1	1	1	1	0
$3+2 = 5 \equiv 1 \pmod{4}$	1	1	1	0	0	1
$3+1 = 4 \equiv 0 \pmod{4}$	1	1	0	1	0	0
$3+0 = 3 \equiv 3 \pmod{4}$	1	1	0	0	1	1
$2+3 = 5 \equiv 1 \pmod{4}$	1	0	1	1	0	1
$2+2 = 4 \equiv 0 \pmod{4}$	1	0	1	0	0	0
$2+1 = 3 \equiv 3 \pmod{4}$	1	0	0	1	1	1
$2+0 = 2 \equiv 2 \pmod{4}$	1	0	0	0	1	0
$1+3 = 4 \equiv 0 \pmod{4}$	0	1	1	1	0	0
$1+2 = 3 \equiv 3 \pmod{4}$	0	1	1	0	1	1
$1+1 = 2 \equiv 2 \pmod{4}$	0	1	0	1	1	0
$1+0 = 1 \equiv 1 \pmod{4}$	0	1	0	0	0	1
$0+3 = 3 \equiv 3 \pmod{4}$	0	0	1	1	1	1
$0+2 = 2 \equiv 2 \pmod{4}$	0	0	1	0	1	0
$0+1 = 1 \equiv 1 \pmod{4}$	0	0	0	1	0	1
$0+0 = 0 \equiv 0 \pmod{4}$	0	0	0	0	0	0

Let's find, what SAT-solver can say about it?

First, we should represent our 2-bit adder as CNF expression.

Using Wolfram Mathematica, we can express 1-bit expression for both adder outputs:

```
In[ ]:=AdderQ0[aL_,bL_]=Xor[aL,bL]
```

```
Out[ ]:=aL ∨ bL
```

```
In[ ]:=AdderQ1[aL_,aH_,bL_,bH_]=Xor[And[aL,bL],Xor[aH,bH]]
```

```
Out[ ]:=aH ∨ bH ∨ (aL && bL)
```

We need such expression, where both parts will generate 1's. Let's use Wolfram Mathematica find all instances of such expression (I glued both parts with And):

```
In[ ]:=Boole[SatisfiabilityInstances[And[AdderQ0[aL,bL],AdderQ1[aL,aH,bL,bH]],{aL,aH,bL,bH},4]]
```

```
Out[ ]:={1,1,0,0},{1,0,0,1},{0,1,1,0},{0,0,1,1}
```

Yes, indeed, Mathematica says, there are 4 inputs which will lead to the result we need. So, Mathematica can also be used as SAT solver.

Nevertheless, let's proceed to CNF form. Using Mathematica again, let's convert our expression to CNF form:

```
In[ ]:=cnf=BooleanConvert[And[AdderQ0[aL,bL],AdderQ1[aL,aH,bL,bH]],`CNF']
```

```
Out[ ]:=(!aH ∥ !bH) && (aH ∥ bH) && (!aL ∥ !bL) && (aL ∥ bL)
```

Looks more complex. The reason of such verbosity is that CNF form doesn't allow XOR operations.

### 11.2.1 MiniSat

For the starters, we can try MiniSat<sup>90</sup>. The standard way to encode CNF expression for MiniSat is to enumerate all OR parts at each line. Also, MiniSat doesn't support variable names, just numbers. Let's enumerate our variables: 1 will be aH, 2 - aL, 3 - bH, 4 - bL.

Here is what I've got when I converted Mathematica expression to the MiniSat input file:

```
p cnf 4 4
-1 -3 0
1 3 0
-2 -4 0
2 4 0
```

Two 4's at the first lines are number of variables and number of clauses respectively. There are 4 lines then, each for each OR clause. Minus before variable number meaning that the variable is negated. Absence of minus - not negated. Zero at the end is just terminating zero, meaning end of the clause.

In other words, each line is OR-clause with optional negations, and the task of MiniSat is to find such set of input, which can satisfy all lines in the input file.

That file I named as *adder.cnf* and now let's try MiniSat:

```
% minisat -verb=0 adder.cnf results.txt
SATISFIABLE
```

The results are in *results.txt* file:

```
SAT
-1 -2 3 4 0
```

This means, if the first two variables (aH and aL) will be *false*, and the last two variables (bH and bL) will be set to *true*, the whole CNF expression is satisfiable. Seems to be true: if bH and bL are the only inputs set to *true*, both resulting bits are also has *true* states.

Now how to get other instances? SAT-solvers, like SMT solvers, produce only one solution (or *instance*).

MiniSat uses PRNG and its initial seed can be set explicitly. I tried different values, but result is still the same. Nevertheless, CryptoMiniSat in this case was able to show all possible 4 instances, in chaotic order, though. So this is not very robust way.

Perhaps, the only known way is to negate solution clause and add it to the input expression. We've got `-1 -2 3 4`, now we can negate all values in it (just toggle minuses: `1 2 -3 -4`) and add it to the end of the input file:

```
p cnf 4 5
-1 -3 0
1 3 0
-2 -4 0
2 4 0
1 2 -3 -4
```

Now we've got another result:

```
SAT
1 2 -3 -4 0
```

This means, aH and aL must be both *true* and bH and bL must be *false*, to satisfy the input expression. Let's negate this clause and add it again:

```
p cnf 4 6
-1 -3 0
1 3 0
-2 -4 0
2 4 0
1 2 -3 -4
-1 -2 3 4 0
```

The result is:

```
SAT
-1 2 3 -4 0
```

aH=false, aL=true, bH=true, bL=false. This is also correct, according to our truth table.

Let's add it again:

<sup>90</sup><http://minisat.se/MiniSat.html>

```
p cnf 4 7
-1 -3 0
1 3 0
-2 -4 0
2 4 0
1 2 -3 -4
-1 -2 3 4 0
1 -2 -3 4 0
```

```
SAT
1 -2 -3 4 0
```

$aH=true$ ,  $aL=false$ ,  $bH=false$ ,  $bL=true$ . This is also correct.  
This is fourth result. There are shouldn't be more. What if to add it?

```
p cnf 4 8
-1 -3 0
1 3 0
-2 -4 0
2 4 0
1 2 -3 -4
-1 -2 3 4 0
1 -2 -3 4 0
-1 2 3 -4 0
```

Now MiniSat just says "UNSATISFIABLE" without any additional information in the resulting file.  
Our example is tiny, but MiniSat can work with huge CNF expressions.

### 11.2.2 CryptoMiniSat

XOR operation is absent in CNF form, but crucial in cryptographical algorithms. Simplest possible way to represent single XOR operation in CNF form is:  $(\neg x \vee \neg y) \wedge (x \vee y)$  - not that small expression, though, many XOR operations in single expression can be optimized better.

One significant difference between MiniSat and CryptoMiniSat is that the latter supports clauses with XOR operations instead of ORs, because CryptoMiniSat has aim to analyze crypto algorithms<sup>91</sup>. XOR clauses are handled by CryptoMiniSat in a special way without translating to OR clauses.

You need just to prepend a clause with "x" in CNF file and OR clause is then treated as XOR clause by CryptoMiniSat. As of 2-bit adder, this smallest possible XOR-CNF expression can be used to find all inputs where both output adder bits are set:

$$(aH \oplus bH) \wedge (aL \oplus bL)$$

This is .cnf file for CryptoMiniSat:

```
p cnf 4 2
x1 3 0
x2 4 0
```

Now I run CryptoMiniSat with various random values to initialize its PRNG ...

```
% cryptominisat4 --verb 0 --random 0 XOR_adder.cnf
s SATISFIABLE
v 1 2 -3 -4 0
% cryptominisat4 --verb 0 --random 1 XOR_adder.cnf
s SATISFIABLE
v -1 -2 3 4 0
% cryptominisat4 --verb 0 --random 2 XOR_adder.cnf
s SATISFIABLE
v 1 -2 -3 4 0
% cryptominisat4 --verb 0 --random 3 XOR_adder.cnf
s SATISFIABLE
v 1 2 -3 -4 0
% cryptominisat4 --verb 0 --random 4 XOR_adder.cnf
s SATISFIABLE
v -1 2 3 -4 0
% cryptominisat4 --verb 0 --random 5 XOR_adder.cnf
s SATISFIABLE
v -1 2 3 -4 0
% cryptominisat4 --verb 0 --random 6 XOR_adder.cnf
s SATISFIABLE
```

<sup>91</sup><http://www.msoos.org/xor-clauses/>

```
v -1 -2 3 4 0
% cryptominisat4 --verb 0 --random 7 XOR_adder.cnf
s SATISFIABLE
v 1 -2 -3 4 0
% cryptominisat4 --verb 0 --random 8 XOR_adder.cnf
s SATISFIABLE
v 1 2 -3 -4 0
% cryptominisat4 --verb 0 --random 9 XOR_adder.cnf
s SATISFIABLE
v 1 2 -3 -4 0
```

Nevertheless, all 4 possible solutions are:

```
v -1 -2 3 4 0
v -1 2 3 -4 0
v 1 -2 -3 4 0
v 1 2 -3 -4 0
```

...the same as reported by MiniSat.

## 11.3 Cracking Minesweeper with SAT solver

See also about cracking it using Z3: [5.11](#).

### 11.3.1 Simple *population count* function

First of all, somehow we need to count neighbour bombs. The counting function is similar to *population count* function.

We can try to make **CNF** expression using Wolfram Mathematica. This will be a function, returning *True* if any of 2 bits of 8 inputs bits are *True* and others are *False*. First, we make truth table of such function:

```
In[]:= tbl2 =
Table[PadLeft[IntegerDigits[i, 2], 8] ->
If[Equal[DigitCount[i, 2][[1]], 2], 1, 0], {i, 0, 255}]

Out[]= {{0, 0, 0, 0, 0, 0, 0, 0} -> 0, {0, 0, 0, 0, 0, 0, 0, 1} -> 0,
{0, 0, 0, 0, 0, 0, 1, 0} -> 0, {0, 0, 0, 0, 0, 0, 1, 1} -> 1,
{0, 0, 0, 0, 0, 1, 0, 0} -> 0, {0, 0, 0, 0, 0, 1, 0, 1} -> 1,
{0, 0, 0, 0, 0, 1, 1, 0} -> 1, {0, 0, 0, 0, 0, 1, 1, 1} -> 0,
{0, 0, 0, 0, 1, 0, 0, 0} -> 0, {0, 0, 0, 0, 1, 0, 0, 1} -> 1,
{0, 0, 0, 0, 1, 0, 1, 0} -> 1, {0, 0, 0, 0, 1, 0, 1, 1} -> 0,
...
{1, 1, 1, 1, 1, 0, 1, 0} -> 0, {1, 1, 1, 1, 1, 0, 1, 1} -> 0,
{1, 1, 1, 1, 1, 1, 0, 0} -> 0, {1, 1, 1, 1, 1, 1, 0, 1} -> 0,
{1, 1, 1, 1, 1, 1, 1, 0} -> 0, {1, 1, 1, 1, 1, 1, 1, 1} -> 0}
```

Now we can make **CNF** expression using this truth table:

```
In[]:= BooleanConvert[
BooleanFunction[tbl2, {a, b, c, d, e, f, g, h}], "CNF"]

Out[]= (! a || ! b || ! c) && (! a || ! b || ! d) && (! a || !
b || ! e) && (! a || ! b || ! f) && (! a || ! b || ! g) && (!
a || ! b || ! h) && (! a || ! c || ! d) && (! a || ! c || !
e) && (! a || ! c || ! f) && (! a || ! c || ! g) && (! a || !
c || ! h) && (! a || ! d || ! e) && (! a || ! d || ! f) && (!
a || ! d || ! g) && (! a || ! d || ! h) && (! a || ! e || !
f) && (! a || ! e || ! g) && (! a || ! e || ! h) && (! a || !
f || ! g) && (! a || ! f || ! h) && (! a || ! g || ! h) && (a ||
b || c || d || e || f || g) && (a || b || c || d || e || f ||
h) && (a || b || c || d || e || g || h) && (a || b || c || d || f ||
g || h) && (a || b || c || e || f || g || h) && (a || b || d ||
e || f || g || h) && (a || c || d || e || f || g ||
h) && (! b || ! c || ! d) && (! b || ! c || ! e) && (! b || !
c || ! f) && (! b || ! c || ! g) && (! b || ! c || ! h) && (!
b || ! d || ! e) && (! b || ! d || ! f) && (! b || ! d || !
g) && (! b || ! d || ! h) && (! b || ! e || ! f) && (! b || !
e || ! g) && (! b || ! e || ! h) && (! b || ! f || ! g) && (!
b || ! f || ! h) && (! b || ! g || ! h) && (b || c || d || e ||
f || g ||
h) && (! c || ! d || ! e) && (! c || ! d || ! f) && (! c || !
d || ! g) && (! c || ! d || ! h) && (! c || ! e || ! f) && (!
```



```

c || ! e || ! g) && (! c || ! e || ! h) && (! c || ! f || !
g) && (! c || ! f || ! h) && (! c || ! g || ! h) && (! d || !
e || ! f) && (! d || ! e || ! g) && (! d || ! e || ! h) && (!
d || ! f || ! g) && (! d || ! f || ! h) && (! d || ! g || !
h) && (! e || ! f || ! g) && (! e || ! f || ! h) && (! e || !
g || ! h) && (! f || ! g || ! h)

```

The syntax is similar to C/C++. Let's check it.

I wrote a Python function to convert Mathematica's output into CNF file which can be feeded to SAT solver:

```

#!/usr/bin/python

import subprocess

def mathematica_to_CNF (s, a):
    s=s.replace("a", a[0]).replace("b", a[1]).replace("c", a[2]).replace("d", a[3])
    s=s.replace("e", a[4]).replace("f", a[5]).replace("g", a[6]).replace("h", a[7])
    s=s.replace("!", "-").replace("||", " ").replace("(", "").replace(")", "")
    s=s.split("&&")
    return s

def POPCNT2 (a):
    s="(!a||!b||!c)&&(!a||!b||!d)&&(!a||!b||!e)&&(!a||!b||!f)&&(!a||!b||!g)&&(!a||!b||!h)&&(!a||!c||!d)&& \
    "(!a||!c||!e)&&(!a||!c||!f)&&(!a||!c||!g)&&(!a||!c||!h)&&(!a||!d||!e)&&(!a||!d||!f)&&(!a||!d||!g)&& \
    "(!a||!d||!h)&&(!a||!e||!f)&&(!a||!e||!g)&&(!a||!e||!h)&&(!a||!f||!g)&&(!a||!f||!h)&&(!a||!g||!h)&& \
    "(a||b||c||d||e||f||g)&&(a||b||c||d||e||f||h)&&(a||b||c||d||e||g||h)&&(a||b||c||d||f||g||h)&& \
    "(a||b||c||e||f||g||h)&&(a||b||d||e||f||g||h)&&(a||c||d||e||f||g||h)&&(!b||!c||!d)&&(!b||!c||!e)&& \
    "(!b||!c||!f)&&(!b||!c||!g)&&(!b||!c||!h)&&(!b||!d||!e)&&(!b||!d||!f)&&(!b||!d||!g)&&(!b||!d||!h)&& \
    "(!b||!e||!f)&&(!b||!e||!g)&&(!b||!e||!h)&&(!b||!f||!g)&&(!b||!f||!h)&&(!b||!g||!h)&&(b||c||d||e||f||g||h) \
    "&& \
    "(!c||!d||!e)&&(!c||!d||!f)&&(!c||!d||!g)&&(!c||!d||!h)&&(!c||!e||!f)&&(!c||!e||!g)&&(!c||!e||!h)&& \
    "(!c||!f||!g)&&(!c||!f||!h)&&(!c||!g||!h)&&(!d||!e||!f)&&(!d||!e||!g)&&(!d||!e||!h)&&(!d||!f||!g)&& \
    "(!d||!f||!h)&&(!d||!g||!h)&&(!e||!f||!g)&&(!e||!f||!h)&&(!e||!g||!h)&&(!f||!g||!h)"
    return mathematica_to_CNF(s, a)

clauses=POPCNT2(["1","2","3","4","5","6","7","8"])

f=open("tmp.cnf", "w")
f.write ("p cnf 8 "+str(len(clauses))+"\n")
for c in clauses:
    f.write(c+" 0\n")
f.close()

```

It replaces a/b/c/... variables to the variable names passed (1/2/3...), reworks syntax, etc. Here is a result:

```

p cnf 8 64
-1 -2 -3 0
-1 -2 -4 0
-1 -2 -5 0
-1 -2 -6 0
-1 -2 -7 0
-1 -2 -8 0
-1 -3 -4 0
-1 -3 -5 0
-1 -3 -6 0
-1 -3 -7 0
-1 -3 -8 0
-1 -4 -5 0
-1 -4 -6 0
-1 -4 -7 0
-1 -4 -8 0
-1 -5 -6 0
-1 -5 -7 0
-1 -5 -8 0
-1 -6 -7 0
-1 -6 -8 0
-1 -7 -8 0
1 2 3 4 5 6 7 0
1 2 3 4 5 6 8 0
1 2 3 4 5 7 8 0
1 2 3 4 6 7 8 0
1 2 3 5 6 7 8 0
1 2 4 5 6 7 8 0
1 3 4 5 6 7 8 0
-2 -3 -4 0
-2 -3 -5 0

```

```

-2 -3 -6 0
-2 -3 -7 0
-2 -3 -8 0
-2 -4 -5 0
-2 -4 -6 0
-2 -4 -7 0
-2 -4 -8 0
-2 -5 -6 0
-2 -5 -7 0
-2 -5 -8 0
-2 -6 -7 0
-2 -6 -8 0
-2 -7 -8 0
2 3 4 5 6 7 8 0
-3 -4 -5 0
-3 -4 -6 0
-3 -4 -7 0
-3 -4 -8 0
-3 -5 -6 0
-3 -5 -7 0
-3 -5 -8 0
-3 -6 -7 0
-3 -6 -8 0
-3 -7 -8 0
-4 -5 -6 0
-4 -5 -7 0
-4 -5 -8 0
-4 -6 -7 0
-4 -6 -8 0
-4 -7 -8 0
-5 -6 -7 0
-5 -6 -8 0
-5 -7 -8 0
-6 -7 -8 0

```

I can run it:

```

% minisat -verb=0 tst1.cnf results.txt
SATISFIABLE

% cat results.txt
SAT
1 -2 -3 -4 -5 -6 -7 8 0

```

The variable name in results lacking minus sign is *True*. Variable name with minus sign is *False*. We see there are just two variables are *True*: 1 and 8. This is indeed correct: MiniSat solver found a condition, for which our function returns *True*. Zero at the end is just a terminal symbol which means nothing.

We can ask MiniSat for another solution, by adding current solution to the input CNF file, but with all variables negated:

```

...
-5 -6 -8 0
-5 -7 -8 0
-6 -7 -8 0
-1 2 3 4 5 6 7 -8 0

```

In plain English language, this means “give me ANY solution which can satisfy all clauses, but also not equal to the last clause we’ve just added”.

MiniSat, indeed, found another solution, again, with only 2 variables equal to *True*:

```

% minisat -verb=0 tst2.cnf results.txt
SATISFIABLE

% cat results.txt
SAT
1 2 -3 -4 -5 -6 -7 -8 0

```

By the way, *population count* function for 8 neighbours (POPCNT8) in CNF form is simplest:

```
a&&b&&c&&d&&e&&f&&g&&h
```

Indeed: it’s true if all 8 input bits are *True*.

The function for 0 neighbours (POPCNT0) is also simple:

```
!a&&!b&&!c&&!d&&!e&&!f&&!g&&!h
```

It means, it will return *True*, if all input variables are *False*.  
By the way, POPCNT1 function is also simple:

```
(!a||!b)&&(!a||!c)&&(!a||!d)&&(!a||!e)&&(!a||!f)&&(!a||!g)&&(!a||!h)&&(a||b||c||d||e||f||g||h)&&  
(!b||!c)&&(!b||!d)&&(!b||!e)&&(!b||!f)&&(!b||!g)&&(!b||!h)&&(!c||!d)&&(!c||!e)&&(!c||!f)&&(!c||!g)&&  
(!c||!h)&&(!d||!e)&&(!d||!f)&&(!d||!g)&&(!d||!h)&&(!e||!f)&&(!e||!g)&&(!e||!h)&&(!f||!g)&&(!f||!h)&&(!g||!h)
```

There is just enumeration of all possible pairs of 8 variables (a/b, a/c, a/d, etc), which implies: no two bits must be present simultaneously in each possible pair. And there is another clause: "(a||b||c||d||e||f||g||h)", which implies: at least one bit must be present among 8 variables.

And yes, you can ask Mathematica for finding [CNF](#) expressions for any other truth table.

### 11.3.2 Minesweeper

Now we can use Mathematica to generate all *population count* functions for 0..8 neighbours.

For 9 · 9 Minesweeper matrix including invisible border, there will be 11 · 11 = 121 variables, mapped to Minesweeper matrix like this:

```
 1  2  3  4  5  6  7  8  9 10 11  
12 13 14 15 16 17 18 19 20 21 22  
23 24 25 26 27 28 29 30 31 32 33  
34 35 36 37 38 39 40 41 42 43 44  
  
...  
  
100 101 102 103 104 105 106 107 108 109 110  
111 112 113 114 115 116 117 118 119 120 121
```

Then we write a Python script which stacks all *population count* functions: each function for each known number of neighbours (digit on Minesweeper field). Each POPCNTx() function takes list of variable numbers and outputs list of clauses to be added to the final [CNF](#) file.

As of empty cells, we also add them as clauses, but with minus sign, which means, the variable must be *False*. Whenever we try to place bomb, we add its variable as clause without minus sign, this means the variable must be *True*.

Then we execute external minisat process. The only thing we need from it is exit code. If an input [CNF](#) is [UNSAT](#), it returns 20:

We use here the information from the previous solving of Minesweeper: [5.11](#).

```
#!/usr/bin/python  
  
import subprocess  
  
WIDTH=9  
HEIGHT=9  
VARS_TOTAL=(WIDTH+2)*(HEIGHT+2)  
  
known=[  
"01?10001?",  
"01?100011",  
"011100000",  
"000000000",  
"111110011",  
"?????1001?",  
"?????3101?",  
"?????211?",  
"?????????"]  
  
def mathematica_to_CNF (s, a):  
    s=s.replace("a", a[0]).replace("b", a[1]).replace("c", a[2]).replace("d", a[3])  
    s=s.replace("e", a[4]).replace("f", a[5]).replace("g", a[6]).replace("h", a[7])  
    s=s.replace("!", "-").replace("||", " ").replace("(", "").replace(")", "").replace(")", "")  
    s=s.split("&&")  
    return s  
  
def POPCNT0 (a):  
    s="!a&&!b&&!c&&!d&&!e&&!f&&!g&&!h"  
    return mathematica_to_CNF(s, a)
```

```

def POPCNT1 (a):
s="(!a|!b)&&(!a|!c)&&(!a|!d)&&(!a|!e)&&(!a|!f)&&(!a|!g)&&(!a|!h)&&(a|b|c|d|e|f|g|h)&&" \
"!b|!c)&&(!b|!d)&&(!b|!e)&&(!b|!f)&&(!b|!g)&&(!b|!h)&&(!c|!d)&&(!c|!e)&&(!c|!f)&&(!c|!g)&&" \
"!c|!h)&&(!d|!e)&&(!d|!f)&&(!d|!g)&&(!d|!h)&&(!e|!f)&&(!e|!g)&&(!e|!h)&&(!f|!g)&&(!f|!h)&&(!g|!h)"
return mathematica_to_CNF(s, a)

def POPCNT2 (a):
s="(!a|!b|!c)&&(!a|!b|!d)&&(!a|!b|!e)&&(!a|!b|!f)&&(!a|!b|!g)&&(!a|!b|!h)&&(!a|!c|!d)&&" \
"!a|!c|!e)&&(!a|!c|!f)&&(!a|!c|!g)&&(!a|!c|!h)&&(!a|!d|!e)&&(!a|!d|!f)&&(!a|!d|!g)&&" \
"!a|!d|!h)&&(!a|!e|!f)&&(!a|!e|!g)&&(!a|!e|!h)&&(!a|!f|!g)&&(!a|!f|!h)&&(!a|!g|!h)&&" \
"(a|b|c|d|e|f|g)&&(a|b|c|d|e|f|h)&&(a|b|c|d|e|g|h)&&(a|b|c|d|f|g|h)&&" \
"(a|b|c|e|f|g|h)&&(a|b|d|e|f|h)&&(a|c|d|e|f|g|h)&&(!b|!c|!d)&&(!b|!c|!e)&&" \
"!b|!c|!f)&&(!b|!c|!g)&&(!b|!c|!h)&&(!b|!d|!e)&&(!b|!d|!f)&&(!b|!d|!g)&&(!b|!d|!h)&&" \
"!b|!e|!f)&&(!b|!e|!g)&&(!b|!e|!h)&&(!b|!f|!g)&&(!b|!f|!h)&&(!b|!g|!h)&&(!b|!c|!d|!e|f|g|h)"
&&" \
"!c|!d|!e)&&(!c|!d|!f)&&(!c|!d|!g)&&(!c|!d|!h)&&(!c|!e|!f)&&(!c|!e|!g)&&(!c|!e|!h)&&" \
"!c|!f|!g)&&(!c|!f|!h)&&(!c|!g|!h)&&(!d|!e|!f)&&(!d|!e|!g)&&(!d|!e|!h)&&(!d|!f|!g)&&" \
"!d|!f|!h)&&(!d|!g|!h)&&(!e|!f|!g)&&(!e|!f|!h)&&(!e|!g|!h)&&(!f|!g|!h)"
return mathematica_to_CNF(s, a)

def POPCNT3 (a):
s="(!a|!b|!c|!d)&&(!a|!b|!c|!e)&&(!a|!b|!c|!f)&&(!a|!b|!c|!g)&&(!a|!b|!c|!h)&&" \
"!a|!b|!d|!e)&&(!a|!b|!d|!f)&&(!a|!b|!d|!g)&&(!a|!b|!d|!h)&&(!a|!b|!e|!f)&&" \
"!a|!b|!e|!g)&&(!a|!b|!e|!h)&&(!a|!b|!f|!g)&&(!a|!b|!f|!h)&&(!a|!b|!g|!h)&&" \
"!a|!c|!d|!e)&&(!a|!c|!d|!f)&&(!a|!c|!d|!g)&&(!a|!c|!d|!h)&&(!a|!c|!e|!f)&&" \
"!a|!c|!e|!g)&&(!a|!c|!e|!h)&&(!a|!c|!f|!g)&&(!a|!c|!f|!h)&&(!a|!c|!g|!h)&&" \
"!a|!d|!e|!f)&&(!a|!d|!e|!g)&&(!a|!d|!e|!h)&&(!a|!d|!f|!g)&&(!a|!d|!f|!h)&&" \
"!a|!d|!g|!h)&&(!a|!e|!f|!g)&&(!a|!e|!f|!h)&&(!a|!e|!g|!h)&&(!a|!f|!g|!h)&&" \
"(a|b|c|d|e|f)&&(a|b|c|d|e|g)&&(a|b|c|d|e|h)&&(a|b|c|d|f|g)&&(a|b|c|d|f|h)&&" \
"(a|b|c|d|g|h)&&(a|b|c|e|f|g)&&(a|b|c|e|h)&&(a|b|c|e|g|h)&&(a|b|c|f|g|h)&&" \
"(a|b|d|e|f|g)&&(a|b|d|e|h)&&(a|b|d|e|g|h)&&(a|b|d|f|g|h)&&(a|b|e|f|g|h)&&" \
"(a|c|d|e|f|g)&&(a|c|d|h)&&(a|c|d|e|f|g)&&(a|c|d|e|h)&&(a|c|d|f|g|h)&&(a|c|e|f|g|h)&&" \
"(a|d|e|f|g|h)&&(!b|!c|!d|!e)&&(!b|!c|!d|!f)&&(!b|!c|!d|!g)&&(!b|!c|!d|!h)&&" \
"!b|!c|!e|!f)&&(!b|!c|!e|!g)&&(!b|!c|!e|!h)&&(!b|!c|!f|!g)&&(!b|!c|!f|!h)&&" \
"!b|!c|!g|!h)&&(!b|!d|!e|!f)&&(!b|!d|!e|!g)&&(!b|!d|!e|!h)&&(!b|!d|!f|!g)&&" \
"!b|!d|!f|!h)&&(!b|!d|!g|!h)&&(!b|!e|!f|!g)&&(!b|!e|!f|!h)&&(!b|!e|!g|!h)&&" \
"!b|!f|!g|!h)&&(!b|!c|!d|!e|f|g)&&(!b|!c|!d|!e|f|h)&&(!b|!c|!d|!e|g|h)&&(!b|!c|!d|!e|h)&&" \
"(b|c|e|f|g|h)&&(b|d|e|f|g|h)&&(!c|!d|!e|!f)&&(!c|!d|!e|!g)&&(!c|!d|!e|!h)&&" \
"!c|!d|!f|!g)&&(!c|!d|!f|!h)&&(!c|!d|!g|!h)&&(!c|!e|!f|!g)&&(!c|!e|!f|!h)&&" \
"!c|!e|!g|!h)&&(!c|!f|!g|!h)&&(c|d|e|f|g|h)&&(!d|!e|!f|!g)&&(!d|!e|!f|!h)&&" \
"!d|!e|!g|!h)&&(!d|!f|!g|!h)&&(!e|!f|!g|!h)"
return mathematica_to_CNF(s, a)

def POPCNT4 (a):
s="(!a|!b|!c|!d|!e)&&(!a|!b|!c|!d|!f)&&(!a|!b|!c|!d|!g)&&(!a|!b|!c|!d|!h)&&" \
"!a|!b|!c|!e|!f)&&(!a|!b|!c|!e|!g)&&(!a|!b|!c|!e|!h)&&(!a|!b|!c|!f|!g)&&" \
"!a|!b|!c|!f|!h)&&(!a|!b|!c|!g|!h)&&(!a|!b|!d|!e|!f)&&(!a|!b|!d|!e|!g)&&" \
"!a|!b|!d|!e|!h)&&(!a|!b|!d|!f|!g)&&(!a|!b|!d|!f|!h)&&(!a|!b|!d|!g|!h)&&" \
"!a|!b|!e|!f|!g)&&(!a|!b|!e|!f|!h)&&(!a|!b|!e|!g|!h)&&(!a|!b|!f|!g|!h)&&" \
"!a|!c|!d|!e|!f)&&(!a|!c|!d|!e|!g)&&(!a|!c|!d|!e|!h)&&(!a|!c|!d|!f|!g)&&" \
"!a|!c|!d|!f|!h)&&(!a|!c|!d|!g|!h)&&(!a|!c|!e|!f|!g)&&(!a|!c|!e|!f|!h)&&" \
"!a|!c|!e|!g|!h)&&(!a|!c|!f|!g|!h)&&(!a|!d|!e|!f|!g)&&(!a|!d|!e|!f|!h)&&" \
"!a|!d|!e|!g|!h)&&(!a|!d|!f|!g|!h)&&(!a|!e|!f|!g|!h)&&(a|b|c|d|e)&&(a|b|c|d|f)&&" \
"(a|b|c|d|g)&&(a|b|c|d|h)&&(a|b|c|e|f)&&(a|b|c|e|g)&&(a|b|c|e|h)&&(a|b|c|f|g)&&" \
"(a|b|c|f|h)&&(a|b|c|g|h)&&(a|b|d|e|f)&&(a|b|d|e|g)&&(a|b|d|e|h)&&(a|b|d|f|g)&&" \
"(a|b|d|f|h)&&(a|b|d|g|h)&&(a|b|e|f|g)&&(a|b|e|f|h)&&(a|b|e|g|h)&&(a|b|f|g|h)&&" \
"(a|c|d|e|f)&&(a|c|d|e|g)&&(a|c|d|e|h)&&(a|c|d|f|g)&&(a|c|d|f|h)&&(a|c|d|g|h)&&" \
"(a|c|e|f|g)&&(a|c|e|h)&&(a|c|e|g|h)&&(a|c|f|g|h)&&(a|d|e|f|g)&&(a|d|e|f|h)&&" \
"(a|d|e|g|h)&&(a|d|f|g|h)&&(a|e|f|g|h)&&(!b|!c|!d|!e|!f)&&(!b|!c|!d|!e|!g)&&" \
"!b|!c|!d|!e|!h)&&(!b|!c|!d|!f|!g)&&(!b|!c|!d|!f|!h)&&(!b|!c|!d|!g|!h)&&" \
"!b|!c|!e|!f|!g)&&(!b|!c|!e|!f|!h)&&(!b|!c|!e|!g|!h)&&(!b|!c|!f|!g|!h)&&" \
"!b|!d|!e|!f|!g)&&(!b|!d|!e|!f|!h)&&(!b|!d|!e|!g|!h)&&(!b|!d|!f|!g|!h)&&" \
"!b|!e|!f|!g|!h)&&(b|c|d|e|f)&&(b|c|d|e|g)&&(b|c|d|e|h)&&(b|c|d|f|g)&&" \
"(b|c|d|f|h)&&(b|c|d|g|h)&&(b|c|e|f|g)&&(b|c|e|f|h)&&(b|c|e|g|h)&&" \
"(b|c|f|g|h)&&(b|d|e|f|g)&&(b|d|e|f|h)&&(b|d|e|g|h)&&(b|d|f|g|h)&&" \
"(b|e|f|g|h)&&(!c|!d|!e|!f|!g)&&(!c|!d|!e|!f|!h)&&(!c|!d|!e|!g|!h)&&" \
"!c|!d|!f|!g|!h)&&(!c|!e|!f|!g|!h)&&(c|d|e|f|g)&&(c|d|e|f|h)&&(c|d|e|g|h)&&" \
"(c|d|f|g|h)&&(c|e|f|g|h)&&(!d|!e|!f|!g|!h)&&(d|e|f|g|h)"
return mathematica_to_CNF(s, a)

def POPCNT5 (a):
s="(!a|!b|!c|!d|!e|!f)&&(!a|!b|!c|!d|!e|!g)&&(!a|!b|!c|!d|!e|!h)&&" \
"!a|!b|!c|!d|!f|!g)&&(!a|!b|!c|!d|!f|!h)&&(!a|!b|!c|!d|!g|!h)&&" \
"!a|!b|!c|!e|!f|!g)&&(!a|!b|!c|!e|!f|!h)&&(!a|!b|!c|!e|!g|!h)&&" \
"!a|!b|!c|!f|!g|!h)&&(!a|!b|!d|!e|!f|!g)&&(!a|!b|!d|!e|!f|!h)&&"

```

```

"!a|!b|!d|!e|!g|!h)&&!a|!b|!d|!f|!g|!h)&&!a|!b|!e|!f|!g|!h)&&" \
"!a|!c|!d|!e|!f|!g)&&!a|!c|!d|!e|!f|!h)&&!a|!c|!d|!e|!g|!h)&&" \
"!a|!c|!d|!f|!g|!h)&&!a|!c|!e|!f|!g|!h)&&!a|!d|!e|!f|!g|!h)&&" \
"(a|b|c|d)&&(a|b|c|e)&&(a|b|c|f)&&(a|b|c|g)&&(a|b|c|h)&&(a|b|d|e)&&" \
"(a|b|d|f)&&(a|b|d|g)&&(a|b|d|h)&&(a|b|e|f)&&(a|b|e|g)&&(a|b|e|h)&&" \
"(a|b|f|g)&&(a|b|f|h)&&(a|b|g|h)&&(a|c|d|e)&&(a|c|d|f)&&(a|c|d|g)&&" \
"(a|c|d|h)&&(a|c|e|f)&&(a|c|e|g)&&(a|c|e|h)&&(a|c|f|g)&&(a|c|f|h)&&" \
"(a|c|g|h)&&(a|d|e|f)&&(a|d|e|g)&&(a|d|e|h)&&(a|d|f|g)&&(a|d|f|h)&&" \
"(a|d|g|h)&&(a|e|f|g)&&(a|e|f|h)&&(a|e|g|h)&&(a|f|g|h)&&!b|!c|!d|!e|!f|!g)&&" \
"!b|!c|!d|!e|!f|!g|!h)&&!b|!c|!d|!e|!f|!g|!h)&&!b|!c|!d|!e|!f|!g|!h)&&" \
"!b|!c|!e|!f|!g|!h)&&!b|!d|!e|!f|!g|!h)&&(b|c|d|e)&&(b|c|d|f)&&" \
"(b|c|d|g)&&(b|c|d|h)&&(b|c|e|f)&&(b|c|e|g)&&(b|c|e|h)&&(b|c|f|g)&&" \
"(b|c|f|h)&&(b|c|g|h)&&(b|d|e|f)&&(b|d|e|g)&&(b|d|e|h)&&(b|d|f|g)&&" \
"(b|d|f|h)&&(b|d|g|h)&&(b|e|f|g)&&(b|e|f|h)&&(b|e|g|h)&&(b|f|g|h)&&" \
"!c|!d|!e|!f|!g|!h)&&(c|d|e|f)&&(c|d|e|g)&&(c|d|e|h)&&(c|d|f|g)&&" \
"(c|d|f|h)&&(c|d|g|h)&&(c|e|f|g)&&(c|e|f|h)&&(c|e|g|h)&&(c|f|g|h)&&" \
"(d|e|f|g)&&(d|e|f|h)&&(d|e|g|h)&&(d|f|g|h)&&(e|f|g|h)"
return mathematica_to_CNF(s, a)

```

```

def POPCNT6 (a):
s="(!a|!b|!c|!d|!e|!f|!g)&&!a|!b|!c|!d|!e|!f|!h)&&!a|!b|!c|!d|!e|!g|!h)&&" \
"!a|!b|!c|!d|!f|!g|!h)&&!a|!b|!c|!e|!f|!g|!h)&&!a|!b|!d|!e|!f|!g|!h)&&" \
"!a|!c|!d|!e|!f|!g|!h)&&(a|b|c)&&(a|b|d)&&(a|b|e)&&(a|b|f)&&(a|b|g)&&(a|b|h)&&" \
"(a|c|d)&&(a|c|e)&&(a|c|f)&&(a|c|g)&&(a|c|h)&&(a|d|e)&&(a|d|f)&&(a|d|g)&&" \
"(a|d|h)&&(a|e|f)&&(a|e|g)&&(a|e|h)&&(a|f|g)&&(a|f|h)&&(a|g|h)&&" \
"!b|!c|!d|!e|!f|!g|!h)&&(b|c|d)&&(b|c|e)&&(b|c|f)&&(b|c|g)&&(b|c|h)&&(b|d|e)&&" \
"(b|d|f)&&(b|d|g)&&(b|d|h)&&(b|e|f)&&(b|e|g)&&(b|e|h)&&(b|f|g)&&(b|f|h)&&(b|g|h)&&" \
"(c|d|e)&&(c|d|f)&&(c|d|g)&&(c|d|h)&&(c|e|f)&&(c|e|g)&&(c|e|h)&&(c|f|g)&&(c|f|h)&&" \
"(c|g|h)&&(d|e|f)&&(d|e|g)&&(d|e|h)&&(d|f|g)&&(d|f|h)&&(d|g|h)&&" \
"(e|f|g)&&(e|f|h)&&(e|g|h)&&(f|g|h)"
return mathematica_to_CNF(s, a)

```

```

def POPCNT7 (a):
s="(!a|!b|!c|!d|!e|!f|!g|!h)&&(a|b)&&(a|c)&&(a|d)&&(a|e)&&(a|f)&&(a|g)&&(a|h)&&(b|c)&&" \
"(b|d)&&(b|e)&&(b|f)&&(b|g)&&(b|h)&&(c|d)&&(c|e)&&(c|f)&&(c|g)&&(c|h)&&(d|e)&&(d|f)&&(d|g)&&" \
"(d|h)&&(e|f)&&(e|g)&&(e|h)&&(f|g)&&(f|h)&&(g|h)"
return mathematica_to_CNF(s, a)

```

```

def POPCNT8 (a):
s="a&&b&&c&&d&&e&&f&&g&&h"
return mathematica_to_CNF(s, a)

```

```

POPCNT_functions=[POPCNT0, POPCNT1, POPCNT2, POPCNT3, POPCNT4, POPCNT5, POPCNT6, POPCNT7, POPCNT8]

```

```

def coords_to_var (row, col):
# we always use SAT variables as strings, anyway.
# the 1st variables is 1, not 0
return str(row*(WIDTH+2)+col+1)

```

```

def chk_bomb(row, col):
clauses=[]

# make empty border
# all variables are negated (because they must be False)
for c in range(WIDTH+2):
    clauses.append ("-"+coords_to_var(0,c))
    clauses.append ("-"+coords_to_var(HEIGHT+1,c))
for r in range(HEIGHT+2):
    clauses.append ("-"+coords_to_var(r,0))
    clauses.append ("-"+coords_to_var(r,WIDTH+1))

for r in range(1,HEIGHT+1):
    for c in range(1,WIDTH+1):
        t=known[r-1][c-1]
        if t in "012345678":
            # cell at r, c is empty (False):
            clauses.append ("-"+coords_to_var(r,c))
            # we need an empty border so the following expression would work for all possible cells:
            neighbours=[coords_to_var(r-1, c-1), coords_to_var(r-1, c), coords_to_var(r-1, c+1),
            coords_to_var(r, c-1),
            coords_to_var(r, c+1), coords_to_var(r+1, c-1), coords_to_var(r+1, c), coords_to_var(r
            +1, c+1)]
            clauses=clauses+POPCNT_functions[int(t)](neighbours)

# place a bomb

```

```

clauses.append (coords_to_var(row,col))

f=open("tmp.cnf", "w")
f.write ("p cnf "+str(VARS_TOTAL)+" "+str(len(clauses))+"\n")
for c in clauses:
    f.write(c+" 0\n")
f.close()

child = subprocess.Popen(["minisat", "tmp.cnf"], stdout=subprocess.PIPE)
child.wait()
# 10 is SAT, 20 is UNSAT
if child.returncode==20:
    print "row=%d, col=%d, unsat!" % (row, col)

for r in range(1,HEIGHT+1):
    for c in range(1,WIDTH+1):
        if known[r-1][c-1]=="?":
            chk_bomb(r, c)

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SAT/minesweeper/minesweeper\\_SAT.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SAT/minesweeper/minesweeper_SAT.py) )

The output CNF file can be large, up to  $\approx 2000$  clauses, or more, here is an example: [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SAT/minesweeper/sample.cnf](https://github.com/dennis714/SAT_SMT_article/blob/master/SAT/minesweeper/sample.cnf).

Anyway, it works just like my previous Z3Py script:

```

row=1, col=3, unsat!
row=6, col=2, unsat!
row=6, col=3, unsat!
row=7, col=4, unsat!
row=7, col=9, unsat!
row=8, col=9, unsat!

```

...but it runs way faster, even considering overhead of executing external program. Perhaps, Z3Py version could be optimized much better?

The files, including Wolfram Mathematica notebook: [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/SAT/minesweeper](https://github.com/dennis714/SAT_SMT_article/tree/master/SAT/minesweeper).

## 11.4 Conway's "Game of Life"

### 11.4.1 Reversing back state of "Game of Life"

How could we reverse back a known state of GoL? This can be solved by brute-force, but this is extremely slow and inefficient.

Let's try to use SAT solver.

First, we need to define a function which will tell, if the new cell will be created/born, preserved/stay or died. Quick refresher: cell is born if it has 3 neighbours, it stays alive if it has 2 or 3 neighbours, it dies in any other case.

This is how I can define a function reflecting state of a new cell in the next state:

```

if center==true:
    return popcnt2(neighbours) || popcnt3(neighbours)
if center==false
    return popcnt3(neighbours)

```

We can get rid of "if" construction:

```

result=(center=true && (popcnt2(neighbours) || popcnt3(neighbours))) || (center=false && popcnt3(neighbours))

```

...where "center" is state of central cell, "neighbours" are 8 neighbouring cells, popcnt2 is a function which returns True if it has exactly 2 bits on input, popcnt3 is the same, but for 3 bits (just like these were used in my "Minesweeper" example (11.3)).

Using Wolfram Mathematica, I first create all helper functions and truth table for the function, which returns *true*, if a cell must be present in the next state, or *false* if not:

```

In[1]:= popcount[n_Integer]:=IntegerDigits[n,2] // Total
In[2]:= popcount2[n_Integer]:=Equal[popcount[n],2]
In[3]:= popcount3[n_Integer]:=Equal[popcount[n],3]

```

```
In[4]:= newcell[center_Integer,neighbours_Integer]:=(center==1 && (popcount2[neighbours]|| popcount3[neighbours
]))||
(center==0 && popcount3[neighbours])

In[13]:= NewCellIsTrue=Flatten[Table[Join[{center},PadLeft[IntegerDigits[neighbours,2],8]] ->
Boole[newcell[center, neighbours]],{neighbours,0,255},{center,0,1}]]

Out[13]= {{0,0,0,0,0,0,0,0,0,0}->0,
{1,0,0,0,0,0,0,0,0,0}->0,
{0,0,0,0,0,0,0,0,0,1}->0,
{1,0,0,0,0,0,0,0,0,1}->0,
{0,0,0,0,0,0,0,0,1,0}->0,
{1,0,0,0,0,0,0,0,1,0}->0,
{0,0,0,0,0,0,0,0,1,1}->0,
{1,0,0,0,0,0,0,0,1,1}->1,
...

```

Now we can create a **CNF**-expression out of truth table:

```
In[14]:= BooleanConvert[BooleanFunction[NewCellIsTrue,{center,a,b,c,d,e,f,g,h}], "CNF"]
Out[14]= (!a||!b||!c||!d)&&(!a||!b||!c||!e)&&(!a||!b||!c||!f)&&(!a||!b||!c||!g)&&(!a||!b||!c||!h)&&
(!a||!b||!d||!e)&&(!a||!b||!d||!f)&&(!a||!b||!d||!g)&&(!a||!b||!d||!h)&&(!a||!b||!e||!f)&&
(!a||!b||!e||!g)&&(!a||!b||!e||!h)&&(!a||!b||!f||!g)&&(!a||!b||!f||!h)&&(!a||!b||!g||!h)&&
(!a||!c||!d||!e)&&(!a||!c||!d||!f)&&(!a||!c||!d||!g)&&(!a||!c||!d||!h)&&(!a||!c||!e||!f)&&
(!a||!c||!e||!g)&&(!a||!c||!e||!h)&&(!a||!c||!f||!g)&&(!a||!c||!f||!h)&&
...

```

Also, we need a second function, *inverted one*, which will return *true* if the cell must be absent in the next state, or *false* otherwise:

```
In[15]:= NewCellIsFalse=Flatten[Table[Join[{center},PadLeft[IntegerDigits[neighbours,2],8]] ->
Boole[Not[newcell[center, neighbours]],{neighbours,0,255},{center,0,1}]]
Out[15]= {{0,0,0,0,0,0,0,0,0,0}->1,
{1,0,0,0,0,0,0,0,0,0}->1,
{0,0,0,0,0,0,0,0,0,1}->1,
{1,0,0,0,0,0,0,0,0,1}->1,
{0,0,0,0,0,0,0,0,1,0}->1,
...

In[16]:= BooleanConvert[BooleanFunction[NewCellIsFalse,{center,a,b,c,d,e,f,g,h}], "CNF"]
Out[16]= (!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&
(!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&
(!a||!b||!center||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&
(!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&
(!a||!b||!c||!d||!e||!f||!g||!h)&&(!a||!b||!c||!d||!e||!f||!g||!h)&&
...

```

Using the very same way as in my “Minesweeper” example, I can convert **CNF** expression to list of clauses:

```
def mathematica_to_CNF (s, center, a):
    s=s.replace("center", center)
    s=s.replace("a", a[0]).replace("b", a[1]).replace("c", a[2]).replace("d", a[3])
    s=s.replace("e", a[4]).replace("f", a[5]).replace("g", a[6]).replace("h", a[7])
    s=s.replace("!", "-").replace("||", " ").replace("(", "").replace(")", "")
    s=s.split("&&")
    return s

```

And again, as in “Minesweeper”, there is an invisible border, to make processing simpler. **SAT** variables are also numbered as in previous example:

```
1 2 3 4 5 6 7 8 9 10 11
12 13 14 15 16 17 18 19 20 21 22
23 24 25 26 27 28 29 30 31 32 33
34 35 36 37 38 39 40 41 42 43 44
...
100 101 102 103 104 105 106 107 108 109 110
111 112 113 114 115 116 117 118 119 120 121

```





```

#!/usr/bin/python

import os
from GoL_SAT_utils import *

final_state=[
" * ",
"* * ",
" * "]

H=len(final_state) # HEIGHT
W=len(final_state[0]) # WIDTH

print "HEIGHT=", H, "WIDTH=", W

VARS_TOTAL=W*H+1
VAR_FALSE=str(VARS_TOTAL)

def try_again (clauses):
# rules for the main part of grid
for r in range(H):
for c in range(W):
if final_state[r][c]=="*":
clauses=clauses+cell_is_true(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))
else:
clauses=clauses+cell_is_false(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))

# cells behind visible grid must always be false:
for c in range(-1, W+1):
for r in [-1,H]:
clauses=clauses+cell_is_false(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))
for c in [-1,W]:
for r in range(-1, H+1):
clauses=clauses+cell_is_false(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))

write_CNF("tmp.cnf", clauses, VARS_TOTAL)

print "%d clauses" % len(clauses)

solution=run_minisat ("tmp.cnf")
os.remove("tmp.cnf")
if solution==None:
print "unsat!"
exit(0)

grid=SAT_solution_to_grid(solution, H, W)

print_grid(grid)
write_RLE(grid)

return grid

clauses=[]
# always false:
clauses.append ("-"+VAR_FALSE)

while True:
solution=try_again(clauses)
clauses.append(negate_clause(grid_to_clause(solution, H, W)))
print ""

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SAT/GoL/reverse1.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SAT/GoL/reverse1.py) )

Here is the result:

```

HEIGHT= 3 WIDTH= 3
2525 clauses
.*
*.*
.*
1.rle written

2526 clauses
.**
*..
*.*
2.rle written

```

```

2527 clauses
**
.*
**
3.rle written

2528 clauses
**
.*
**
4.rle written

2529 clauses
**
.*
**
5.rle written

2530 clauses
**
.*
**
6.rle written

2531 clauses
unsat!

```

The first result is the same as initial state. Indeed: this is “still life”, i.e., state which will never change, and it is correct solution. The last solution is also valid.

Now the problem: 2nd, 3rd, 4th and 5th solutions are equivalent to each other, they just mirrored or rotated. In fact, this is reflectional<sup>92</sup> (like in mirror) and rotational<sup>93</sup> symmetries. We can solve this easily: we will take each solution, reflect and rotate it and add them negated to the list of clauses, so minisat will skip them during its work:

```

...
while True:
    solution=try_again(clauses)
    clauses.append(negate_clause(grid_to_clause(solution, H, W)))
    clauses.append(negate_clause(grid_to_clause(reflect_vertically(solution), H, W)))
    clauses.append(negate_clause(grid_to_clause(reflect_horizontally(solution), H, W)))
    # is this square?
    if W==H:
        clauses.append(negate_clause(grid_to_clause(rotate_square_array(solution,1), H, W)))
        clauses.append(negate_clause(grid_to_clause(rotate_square_array(solution,2), H, W)))
        clauses.append(negate_clause(grid_to_clause(rotate_square_array(solution,3), H, W)))
    print ""
...

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SAT/GoL/reverse2.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SAT/GoL/reverse2.py) )

Functions `reflect_vertically()`, `reflect_horizontally` and `rotate_squarearray()` are simple array manipulation routines.

Now we get just 3 solutions:

```

HEIGHT= 3 WIDTH= 3
2525 clauses
.*
**
.*
1.rle written

2531 clauses
**
.*
**
2.rle written

```

<sup>92</sup>[https://en.wikipedia.org/wiki/Reflection\\_symmetry](https://en.wikipedia.org/wiki/Reflection_symmetry)

<sup>93</sup>[https://en.wikipedia.org/wiki/Rotational\\_symmetry](https://en.wikipedia.org/wiki/Rotational_symmetry)

```
2537 clauses
*. *
.*
.*
*. *
3.rle written

2543 clauses
unsat!
```

This one has only one single ancestor:

```
final_state=[
" * ",
" * ",
" * "]
_PRE_END

_PRE_BEGIN
HEIGHT= 3 WIDTH= 3
2503 clauses
...
***
...
1.rle written

2509 clauses
unsat!
```

This is oscillator, of course.  
How many states can lead to such picture?

```
final_state=[
" * ",
"  ",
" ** ",
" * ",
" * ",
" *** "]
```

28, these are few:

```
HEIGHT= 6 WIDTH= 5
5217 clauses
.*.*
..*..
.**.*
..*..
..*.*
.**..
1.rle written

5220 clauses
.*.*
..*..
.**.*
..*..
*.*.*
.**..
2.rle written

5223 clauses
.*.*
..*..
.**.*
....*
*.*.*
.**..
3.rle written

5226 clauses
.*.*
..*..
.**.*
*.*.*
..*.*
.**..
```

```
4.rle written
```

```
...
```

Now the biggest, “space invader”:

```
final_state=[
"
 *      *
"
 *      *
"
*****
"
** *** **
"
*****
"
* ***** *
"
* *      * *
"
 ** **
"
"]
```

```
HEIGHT= 10 WIDTH= 13
```

```
16469 clauses
```

```
..*.***.....
.....*****
.....**.*.....
.....*.*.....
..*.*.*.*.....
.*.*.***.*
*.....*.*.*
.*.*.*.....
.*.....**.....
.....**.*.....
1.rle written
```

```
16472 clauses
```

```
*.*.***.....
.....*****
.....**.*.....
.....*.*.....
..*.*.*.*.....
.*.*.***.*
*.....*.*.*
.*.*.*.....
.*.....**.....
.....**.*.....
2.rle written
```

```
16475 clauses
```

```
..*.***.....
*.....*****
.....**.*.....
.....*.*.....
..*.*.*.*.....
.*.*.***.*
*.....*.*.*
.*.*.*.....
.*.....**.....
.....**.*.....
3.rle written
```

```
...
```

I don't know how many possible states can lead to “space invader”, perhaps, too many. Had to stop it. And it slows down during execution, because number of clauses is increasing (because of negating solutions addition).

All solutions are also exported to RLE files, which can be opened by Golly<sup>94</sup>.

### 11.4.2 Finding “still lives”

“Still life” in terms of GoL is a state which doesn't change at all.

First, using previous definitions, we will define a truth table of function, which will return *true*, if the center cell of the next state is the same as it has been in the previous state, i.e., hasn't been changed:

<sup>94</sup><http://golly.sourceforge.net/>



```

"!b||!center||!d||!e||!g)&&!b||!center||!d||!e||!h)&&!b||!center||!d||!f||!g)&&" \
"!b||!center||!d||!f||!h)&&!b||!center||!d||!g||!h)&&!b||!center||!e||!f||!g)&&" \
"!b||!center||!e||!f||!h)&&!b||!center||!e||!g||!h)&&!b||!center||!f||!g||!h)&&" \
"(b||c||!center||!d||!e||!f||!g||!h)&&!c||!center||!d||!e||!f)&&!c||!center||!d||!e||!g)&&" \
"!c||!center||!d||!e||!h)&&!c||!center||!d||!f||!g)&&!c||!center||!d||!f||!h)&&" \
"!c||!center||!d||!g||!h)&&!c||!center||!e||!f||!g)&&!c||!center||!e||!f||!h)&&" \
"!c||!center||!e||!g||!h)&&!c||!center||!f||!g||!h)&&!center||!d||!e||!f||!g)&&" \
"!center||!d||!e||!f||!h)&&!center||!d||!e||!g||!h)&&!center||!d||!f||!g||!h)&&" \
"!center||!e||!f||!g||!h)"

return mathematica_to_CNF(s, center, a)

def try_again (clauses):
# rules for the main part of grid
for r in range(H):
    for c in range(W):
        clauses=clauses+stillife(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))

# cells behind visible grid must always be false:
for c in range(-1, W+1):
    for r in [-1,H]:
        clauses=clauses+cell_is_false(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))
for c in [-1,W]:
    for r in range(-1, H+1):
        clauses=clauses+cell_is_false(coords_to_var(r, c, H, W), get_neighbours(r, c, H, W))

write_CNF("tmp.cnf", clauses, VARS_TOTAL)

print "%d clauses" % len(clauses)

solution=run_minisat ("tmp.cnf")
os.remove("tmp.cnf")
if solution==None:
    print "unsat!"
    exit(0)

grid=SAT_solution_to_grid(solution, H, W)
print_grid(grid)
write_RLE(grid)

return grid

clauses=[]
# always false:
clauses.append ("-"+VAR_FALSE)

while True:
    solution=try_again(clauses)
    clauses.append(negate_clause(grid_to_clause(solution, H, W)))
    clauses.append(negate_clause(grid_to_clause(reflect_vertically(solution), H, W)))
    clauses.append(negate_clause(grid_to_clause(reflect_horizontally(solution), H, W)))
    # is this square?
    if W==H:
        clauses.append(negate_clause(grid_to_clause(rotate_square_array(solution,1), H, W)))
        clauses.append(negate_clause(grid_to_clause(rotate_square_array(solution,2), H, W)))
        clauses.append(negate_clause(grid_to_clause(rotate_square_array(solution,3), H, W)))
    print ""

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SAT/GoL/stillife1.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SAT/GoL/stillife1.py) )

What we've got for 2 · 2?

```

1881 clauses
..
..
1.rle written

1887 clauses
**
**
2.rle written

1893 clauses
unsat!

```

Both solutions are correct: empty square will progress into empty square (no cells are born). 2 · 2 box is

also known “still life”.

What about 3 · 3 square?

```
2887 clauses
...
...
...
1.rle written

2893 clauses
.*
.*
...
2.rle written

2899 clauses
.*
.*
*.*
**
3.rle written

2905 clauses
.*
*.*
**
4.rle written

2911 clauses
.*
*.*
*.*
5.rle written

2917 clauses
unsat!
```

Here is a problem: we see familiar 2 · 2 box, but shifted. This is indeed correct solution, but we don't interested in it, because it has been already seen.

What we can do is add another condition. We can force minisat to find solutions with no empty rows and columns. This is easy. These are SAT variables for 5 · 5 square:

```
1  2  3  4  5
6  7  8  9 10
11 12 13 14 15
16 17 18 19 20
21 22 23 24 25
```

Each clause is “OR” clause, so all we have to do is to add 5 clauses:

```
1 OR 2 OR 3 OR 4 OR 5
6 OR 7 OR 8 OR 9 OR 10
...

```

That means that each row must have at least one *True* value somewhere. We can also do this for each column as well.

```
...
# each row must contain at least one cell!
for r in range(H):
    clauses.append(" ".join([coords_to_var(r, c, H, W) for c in range(W)]))

# each column must contain at least one cell!
for c in range(W):
    clauses.append(" ".join([coords_to_var(r, c, H, W) for r in range(H)]))
...

```

( [https://github.com/dennis714/SAT\\_SMT\\_article/blob/master/SAT/GoL/stillife2.py](https://github.com/dennis714/SAT_SMT_article/blob/master/SAT/GoL/stillife2.py) )

Now we can see that 3 · 3 square has 3 possible “still lives”:

```
2893 clauses
.*
.*

```

```

*.*
**
1.rle written

2899 clauses
*.*
*.*
*.*
2.rle written

2905 clauses
**
*.*
**
3.rle written

2911 clauses
unsat!

```

4.4 has 7:

```

4169 clauses
.*.*
...*
***.
*...
1.rle written

4175 clauses
.*.*
.*.*
*.*.
**..
2.rle written

4181 clauses
.*.*
*.*
*.*.
**..
3.rle written

4187 clauses
.*.*
*.*
*.*.
**..
4.rle written

4193 clauses
.*.*
*.*
*.*.
*.*.
5.rle written

4199 clauses
.*.*
*.*
*.*.
*.*.
6.rle written

4205 clauses
.*.*
*.*
*.*
*.*.
7.rle written

4211 clauses
unsat!

```

When I try large squares, like  $20 \cdot 20$ , funny things happen. First of all, minisat finds solutions not very pleasing aesthetically, but still correct, like:









### 11.4.3 The source code

Source code and Wolfram Mathematica notebook: [https://github.com/dennis714/SAT\\_SMT\\_article/tree/master/SAT/GoL](https://github.com/dennis714/SAT_SMT_article/tree/master/SAT/GoL).

## 12 Acronyms used

<b>CNF</b> Conjunctive normal form.....	3
<b>DNF</b> Disjunctive normal form.....	124
<b>DSL</b> Domain-specific language.....	4
<b>CPRNG</b> Cryptographically Secure Pseudorandom Number Generator .....	20
<b>SMT</b> Satisfiability modulo theories .....	1
<b>SAT</b> Boolean satisfiability problem.....	1
<b>LCG</b> Linear congruential generator .....	1
<b>PL</b> Programming Language.....	4
<b>OOP</b> Object-oriented programming .....	40
<b>SSA</b> Static single assignment form.....	32
<b>CPU</b> Central processing unit.....	35
<b>FPU</b> Floating-point unit.....	60
<b>PRNG</b> Pseudorandom number generator.....	70
<b>CRT</b> C runtime library .....	70
<b>CRC</b> Cyclic redundancy check .....	67
<b>AST</b> Abstract syntax tree.....	41
<b>AKA</b> Also Known As .....	1
<b>CTF</b> Capture the Flag .....	114
<b>ISA</b> Instruction Set Architecture.....	35
<b>CSP</b> Constraint satisfaction problem .....	6

<b>CS</b> Computer science .....	3
<b>DAG</b> Directed acyclic graph.....	30
<b>NOP</b> No Operation .....	39
<b>JVM</b> Java Virtual Machine.....	60
<b>VM</b> Virtual Machine .....	73
<b>LZSS</b> Lempel-Ziv-Storer-Szymanski .....	99
<b>RAM</b> Random-access memory .....	100
<b>FPGA</b> Field-programmable gate array .....	116
<b>EDA</b> Electronic design automation.....	116
<b>MAC</b> Message authentication code .....	116
<b>ECC</b> Elliptic curve cryptography .....	116
<b>API</b> Application programming interface .....	5
<b>NSA</b> National Security Agency .....	20