

# SATISFIABILITY

AND

## THE ART OF COMPUTER PROGRAMMING

Don Knuth ~ SAT 2012 ~ TRENTO

THE ART OF  
COMPUTER PROGRAMMING  
VOLUME 4      PRE-FASCICLE 6A

A (VERY INCOMPLETE)  
DRAFT OF SECTION 7.2.2.2:  
SATISFIABILITY

DONALD E. KNUTH *Stanford University*

ADDISON-WESLEY



[http://www-cs-faculty.stanford.edu/~knuth/  
fasc6a.ps.gz](http://www-cs-faculty.stanford.edu/~knuth/fasc6a.ps.gz)

\* \* \*

Zeroth printing (revision -93), 10 June 2012

\* \* \*

THIS BOOKLET contains draft material that I'm circulating to experts in the field, in hopes that they can help remove its most egregious errors before too many other people see it. I am also, however, posting it on the Internet for courageous and/or random readers who don't mind the risk of reading a few pages that have not yet reached a very mature state. *Beware:* This material has not yet been proofread as thoroughly as ...

At the moment I'm just beginning to write this material, so it's in an *extremely* raw state. But hey, I had to start somewhere. As I finish drafting small pieces of the final big picture, I'm trying them out by making these test pages. You will soon see, however, that I've got a *long* way to go before I'll have anything coherent. These are scraps that I hope to refine and polish (if the FORCE stays with me). You'll also notice that many of the equation numbers and exercise numbers are flaky, because I keep changing them frequently.

\* \* \*

My notes on combinatorial algorithms have been accumulating for more than fifty years, yet I fear that in many respects my knowledge is woefully behind the times. Please look, for example, at the exercises that I've classed as research problems (rated with difficulty level 46 or higher), namely exercises 90, 91, 108, . . . ; I've also implicitly mentioned or posed additional

unsolved questions in the answers to exercises 64, 80, 81, 90, 108, 185, 300, 323, . . . . Are those problems still open? Please inform me if you know of a solution to any of these intriguing questions. And of course if no solution is known today but you do make progress on any of them in the future, I hope you'll let me know.

I urgently need your help also with respect to some exercises that I made up as I was preparing this material. I certainly don't like to receive credit for things that have already been published by others, and most of these results are quite natural "fruits" that were just waiting to be "plucked." Therefore please tell me if you know who deserves to be credited, with respect to the ideas found in exercises 1, 2, 3, 4, 14, 61, 62, 68, 80(b,c,d), 82, 83, 93, 160, 168, 178, 179, 180, 235, 236, 237, 240, 251(b), 252, 280, 293, 320, 321, 324, . . . , and/or the answers to exercises . . . .

Special thanks are due to Armin Biere, Niklas Eén, Svante Janson, Oliver Kullmann, Wes Pegden, Niklas Sörensson, and ... for their detailed comments on my early attempts at exposition, as well as to numerous other correspondents who have contributed crucial corrections.

I happily offer a “finder’s fee” of \$2.56 for each error in this draft when it is first reported to me, whether that error be typographical, technical, or historical. The same reward holds for items that I forgot to put in the index. And valuable suggestions for improvements to the text are worth 32¢ each. (Furthermore, if you find a better solution to an exercise, I’ll actually do my best to give you immortal glory, by publishing your name in the eventual book:—)

<http://www-cs-faculty.stanford.edu/~knuth/fasc6a.ps.gz>

\* \* \*

So far there are 75 pages (37 for text, 13 for exercises, 25 for answers); 15 of those pages are about basic probability techniques, not about satisfiability.

I think the section on SAT will eventually amount to about 150 pages in Volume 4B. (For comparison, Volume 4A already contains about 135 pages about BDDs and ZDDs.)

bytes	date	program	
28336	2011-08-26	<b>sat0.w</b>	tabula rasa
29623	2011-09-03	<b>sat1.w</b>	truth table
22240	2011-09-13	<b>sat2.w</b>	incremental truth table
39733	2011-09-18	<b>sat3.w</b>	DPL 1962 “ready list”
39727	2011-11-24	<b>sat4.w</b>	prefer binary clauses
42912	2011-12-04	<b>sat5.w</b>	more efficient
55812	2011-12-28	<b>sat6.w</b>	Cook’s clause learning
28303	2012-02-06	<b>sat7.w</b>	Papadimitriou’s “Walk”
29380	2012-03-07	<b>sat8.w</b>	“WalkSAT”
38981	2012-05-10	<b>sat9.w</b>	Survey Propagation

Stephen A. Cook, CSC 2409S lecture notes  
(University of Toronto, 17 January 1972)

It operates on an expanding set of clauses  $c_1, c_2, \dots, c_n$   
(starting with the initial ones)

It has a push down stack of literals  $\ell_1, \dots, \ell_m$ , consist-  
ent with no repetitions.

At every time we have a partial truth assignment  
obtained by assigning true to  $\ell_1, \dots, \ell_m$  on the stack. We  
think of the objective as finding a satisfying assignment.

1)  $m \leftarrow 0$  ;  $c_1, \dots, c_n$  given (we start with initial clauses  
 $c_1, \dots, c_n$  and empty stack)

Comment:

At every time before we enter step 2) no clause is  
completely falsified by the partial truth assignment which  
verifies all  $\ell_1, \dots, \ell_m$ .

2)  $m \leftarrow m+1$

Select a literal  $\ell_m$  such that neither  $\ell_m$  nor  $\bar{\ell}_m$  appears  
on the stack.

If no such  $\ell_m$  exists then we have found a satisfying



0.0.1	0.0.2	0.0.3	0.0.4	0.0.5
0.1.1	0.1.2	0.1.3	0.1.4	0.1.5
0.2.1	0.2.2	0.2.3	0.2.4	0.2.5
.	.	.	.	.
4.3.1	4.3.2	4.3.3	4.3.4	4.3.5
4.4.1	4.4.2	4.4.3	4.4.4	4.4.5
~0.0.1	~1.1.1			
~0.0.1	~2.2.1			
~0.0.1	~3.3.1			
.	.	.	.	.
~4.2.5	~4.3.5			
~4.2.5	~4.4.5			
~4.3.5	~4.4.5			

<http://www-cs-faculty.stanford.edu/~knuth/programs/sat-to-dimacs.w>

<http://www-cs-faculty.stanford.edu/~knuth/programs/dimacs-to-sat.w>

```
c file created by SAT-TO-DIMACS Wed Jun 6 10:23:47 2012
c 4.4.5 -> 125
c 4.4.4 -> 124
.
.
.
c 0.0.2 -> 2
c 0.0.1 -> 1
p cnf 125 825
-125 -120 0
-125 -115 0
.
.
.
15 14 13 12 11 0
10 9 8 7 6 0
5 4 3 2 1 0
```

```
[1]> sat0 < queen5x5color.dat
(125 variables, 825 clauses, 1725 literals)
~0.0.1 ~0.0.2 ~0.0.3 ~0.0.4 0.0.5 ... ~4.4.5
Altogether 19|700+18|785 mems, 39|208 bytes.

[2]> sat1 < queen5x5color.dat
(125 variables, 825 clauses, 1725 literals)
~0.0.1 0.0.2 ~0.0.3 ~0.0.4 ~0.0.5 ... 4.4.5
Altogether 12|652+2|709 mems, 40|624 bytes.

[3]> sat2 < queen5x5color.dat
(125 variables, 825 clauses, 1725 literals)
~0.0.1 ~0.0.2 ~0.0.3 ~0.0.4 0.0.5 ... ~4.4.5
Altogether 4|526+3|517|938|425 mems, 15|516 bytes.
```

```
#define o  mems++      /* count one mem */
#define oo mems += 2     /* count two mems */
```

46. As an experiment, I'm swapping the first true literal into the first position of its clause, hoping that subsequent "decrease" loops will thereby be shortened.

⟨ Increase the breakcount of  $c$ 's critical variable 46 ⟩ ≡

```
{  
  for (o, ii = i = cmem[c].start; ; i++) {  
    o, q = mem[i];  
    if (o, value(q)) break;  
  }  
  o, vmem[q >> 1].breakcount++;  
  if i ≠ ii oo, mem[i] = mem[ii], mem[ii] = q;  
}
```

This code is used in section 44.

[4]> sat3 < queen5x5color.dat  
(125 variables, 825 clauses, 1725 literals)  
~0.0.1 ~0.0.2 ~0.0.3 ~0.0.4 0.0.5 ... ~1.3.4  
Altogether 19|952+11|647 mems, 42|724 bytes.

[5]> sat5 < queen5x5color.dat  
(125 variables, 825 clauses, 1725 literals)  
~0.0.1 ~0.1.1 ~0.2.1 ~0.3.1 0.4.1 ... 0.2.3  
Altogether 18|851+688|489|345 mems, 73|248 bytes.

[6]> sat6 < queen5x5color.dat  
(125 variables, 825 clauses, 1725 literals)  
~1.1.1 ~2.2.1 ~3.3.1 ~4.4.1 ~0.1.1 ... 0.3.5  
Altogether 0 new clauses, 0 recycled,  
36|176+17|047 mems, 74|748 bytes.

```
[7]> sat3 d10|000|000 < rand-3-125-600-0.dat
(125 variables, 600 clauses, 1800 literals)
after 10|002|287 mems:0043103310050244541525454
after 20|001|484 mems:0043233352305014452151514
after 30|002|010 mems:033140403243032441550535
after 40|000|099 mems:3045100324034535554155454
after 50|005|241 mems:3302043304005454315125
after 60|000|099 mems:3332353351050545440445544
~
Altogether 27|638+60|489|039 mems, 42|124 bytes.
```

```
[8]> sat5 d10|000|000 < rand-3-125-600-0.dat
(125 variables, 600 clauses, 1800 literals)
after 10|002|965 mems:0acd545d54a456daa544c47ad
after 20|003|037 mems:3bcb44b554dc4babca
after 30|000|035 mems:3c5ba445a54bd4a4dd4544d44
~
Altogether 18|551+31|987|450 mems, 69|648 bytes.
```

```
[9]> sat3 d2000 < queen5x5color.dat
(125 variables, 825 clauses, 1725 literals)
after 2|019 mems:11114555
after 4|027 mems:11114555555555555511145555
after 6|026 mems:111145555555555555111455555555
after 8|043 mems:111145555555555555111455555555
after 10|032 mems:1111455555555555111455555555
~0.0.1 ~0.0.2 ~0.0.3 ~0.0.4 0.0.5 ... ~1.3.4
Altogether 19|952+11|647 mems, 42|724 bytes.
```

```
[10]> sat5 d100|000|000 < queen5x5color.dat
(125 variables, 825 clauses, 1725 literals)
after 100|000|033 mems:bbbbbbbbbbbbb|bc55555577
after 200|005|826 mems:bbbbbbbbbbbc55555557777
after 300|000|049 mems:bbbbbbbbb|bc5555557777bc55
after 400|000|033 mems:bbbbbbb|bc555555557777bbb
after 500|000|050 mems:bbbbbb|bc5555555557777bbb
after 600|000|005 mems:bbbbbc5555555557777bc55
~0.0.1 ~0.1.1 ~0.2.1 ~0.3.1 0.4.1 ... 0.2.3
Altogether 18|851+688|489|345 mems, 73|248 bytes.
```

```
[11]> sat3-rand d100|000 < queen5x5color.dat
(125 variables, 825 clauses, 1725 literals)
after 100|023 mems:1111111111111111111111111114
after 202|114 mems:11111111111111111111111111125
after 300|010 mems:111111111111111111111111125555
after 400|042 mems:11111111111111111111111112555555
after 501|349 mems:111111111111111111111111125555555
after 603|041 mems:1111111111111111111111111255555555
after 701|449 mems:1111111111111111111111125555555555
~1.0.2 ~2.3.1 ~2.1.5 ~2.0.2 ... ~0.2.5
Altogether 21|679+795|425 mems, 42|724 bytes.
```

10000	01000	00100	00010	00001
00010	00001	10000	01000	00100
10000	01000	00100	00010	00001
00001	10000	01000	00100	00010
00100	00010	00001	10000	01000

000	001	010	011	100
011	100	000	001	010
000	001	010	011	100
100	000	001	010	011
010	011	100	000	001

~0.0.3	~0.0.1	
~0.1.3	~0.1.1	
.	.	.
~4.4.3	~4.4.1	
~0.0.2	~0.0.1	
~0.1.2	~0.1.1	
.	.	.
~4.4.2	~4.4.1	
~0.0^0.11	0.0.1	0.1.1
0.0^0.11	0.0.1	~0.1.1
0.0^0.11	~0.0.1	0.1.1
~0.0^0.11	~0.0.1	~0.1.1
.	.	.
0.0^4.41	0.0^4.42	0.0^4.43

```
[12]> sat5 d500|000 < queen5x5color-binary.dat  
(555 variables, 2130 clauses, 6340 literals)  
after 500|597 mems:ba45a45a4a55a55d554a5554a55  
after 1|000|633 mems:ba45a45a4a55a55d554a5554a55  
after 1|500|013 mems:ba45a45a4a55a55d554a5554a55  
after 2|000|006 mems:ba45a45a4a55a55d554a5554a55  
~0.0.3 0.0^4.43 4.4.3 ~4.4.1 ... ~1.3^1.41  
Altogether 67|011+2|312|306 mems, 256|608 bytes.  
  
[13]> sat3 d1|000|000|000|000 < queen5x5color-bin  
(555 variables, 2130 clauses, 6340 literals)  
after 1|000|000|000|063 mems:11151541551545424445  
after 2|000|000|000|122 mems:11151541552445155515  
~0.0.3 ~0.0.1 ~0.1.3 ~0.0^0.13 ... 0.2^0.32  
Altogether 116|372+2|189|929|944|193 mems, 156|244
```

```
[14]> sat3-rand d1|000|000|000|000 < queen5x5color.cnf  
(555 variables, 2130 clauses, 6340 literals)  
after 1|000|000|001|844 mems:00000000000000000000000000000000  
after 2|000|000|012|419 mems:00000000000000000000000000000003  
after 3|000|000|002|278 mems:00000000000000000000000000000003  
after 4|000|000|000|002 mems:000000000000000000000000000000030  
after 5|000|000|000|557 mems:000000000000000000000000000000030  
after 6|000|000|004|086 mems:000000000000000000000000000000033  
after 7|000|000|011|274 mems:000000000000000000000000000000033  
after 8|000|000|000|156 mems:0000000000000000000000000000000300  
after 9|000|000|000|495 mems:0000000000000000000000000000000303  
after 10|000|000|004|416 mems:000000000000000000000000000000030  
after 11|000|000|000|191 mems:000000000000000000000000000000033  
etc... estimated 3 petamems to finish (3e15)!
```

~0.0.3 ~0.0.1

. . . .

~4.4.2 ~4.4.1

0.0.3 0.1.3 0.0.2 0.1.2 0.0.1 0.1.1

~0.0.3 ~0.1.3 0.0.2 0.1.2 0.0.1 0.1.1

0.0.3 0.1.3 ~0.0.2 ~0.1.2 0.0.1 0.1.1

~0.0.3 ~0.1.3 ~0.0.2 ~0.1.2 0.0.1 0.1.1

0.0.3 0.1.3 0.0.2 0.1.2 ~0.0.1 ~0.1.1

. . . .

0.0.3 4.4.3 0.0.2 4.4.2 0.0.1 4.4.1

~0.0.3 ~4.4.3 0.0.2 4.4.2 0.0.1 4.4.1

0.0.3 4.4.3 ~0.0.2 ~4.4.2 0.0.1 4.4.1

~0.0.3 ~4.4.3 ~0.0.2 ~4.4.2 0.0.1 4.4.1

0.0.3 4.4.3 0.0.2 4.4.2 ~0.0.1 ~4.4.1

```
[15]> sat5 < queen5x5color-altbinary.dat
(75 variables, 850 clauses, 4900 literals)
0.0.3 ~0.0.1 0.1.3 ~0.1.1 0.0.2 ... ~2.1.1
Altogether 43|651+76|175 mems, 145|248 bytes.

[16]> sat3 < queen5x5color-altbinary.dat
(75 variables, 850 clauses, 4900 literals)
0.0.3 ~0.0.1 0.1.3 ~0.1.1 ~0.2.3 ... ~0.4.2
Altogether 77|652+675|724|229|162 mems, 90|342 byt

[17]> sat3-rand < queen5x5color-altbinary.dat
(75 variables, 850 clauses, 4900 literals)
3.4.2 ~3.4.1 0.1.2 ~0.1.1 4.2.3 ... ~0.4.2
Altogether 78|695+6426|416 mems, 90|324 bytes.
```

$$\begin{array}{cccccccccc}
 & & & & & x_5 & x_4 & x_3 & x_2 & x_1 & x_0 \\
 & & & & y_7 & y_6 & y_5 & y_4 & y_3 & y_2 & y_1 & y_0 \\
 \hline
 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1
 \end{array}$$

$$(X \times Y = 6401)$$

$\sim Z_{13}$   
 $A_2:6 \quad \sim Y_4 \quad \sim X_2$   
 $\sim P_0:6 \quad \sim A_1:6 \quad \sim A_0:6$   
 $R_0:6 \quad \sim A_2:6 \quad \sim P_0:6$   
 $Q_1:9 \quad \sim A_4:9 \quad \sim A_3:9$   
 $\sim A_9:11 \quad R_1:10 \quad Q_1:10$   
 $A_{10:5} \quad A_8:5 \quad \sim P_2:5$   
 $P_3:7 \quad A_{10:7} \quad \sim A_9:7$   
 $\sim R_3:3 \quad Q_2:2$   
 $\sim U_{11} \quad A_{13:11} \quad P_3:11$   
 $C_8:3 \quad \sim C_8:2$   
 $\sim C_{11:2} \quad B_{11:2} \quad V_{10}$   
 $\sim D_{13:2} \quad U_{12}$

. . . .

[18]> sat3 < 6401.dat

(286 variables, 911 clauses, 2309 literals)

~Z13 ... ~A8:9 A10:9 ~A6:8 ~Q2:8 ~A7:7 ~Y4

Altogether 40|217+183|094 mems, 63|704 bytes.

[19]> sat5 < 6401.dat

(286 variables, 911 clauses, 2309 literals)

~Z13 ... A0:2 A2:2 ~P2:3 A0:5 X2 Y2

Altogether 26|121+12|307|600 mems, 104|784 bytes.

[20]> sat6 < 6401.dat

(286 variables, 911 clauses, 2309 literals)

~Z13 ... A2:5 P2:8 ~A10:8 ~A1:4 ~A1:6 ~B8:5

Altogether 2|019 new clauses, 439 recycled,

40|415+3|159|955 mems, 720|648 bytes.

[21]> sat5 d5|000|000 < 6401.dat  
(286 variables, 911 clauses, 2309 literals)  
after 5|008|203 mems:544554555555444a4455d55\  
    d55ba4b44d54d54d5455b5555555d545c5  
after 10|001|285 mems:544554555555444d54a445\  
    5d555d55bd5555555d555ba4b44d54d54a445\  
    5d555555c54a44555555555455b555ba4d4\  
    55b55555555555455555d5  
~Z13 ... A0:2 A2:2 ~P2:3 A0:5 X2 Y2  
Altogether 26|121+12|307|600 mems, 104|784 bytes.



voted to making the program fast. For large problems, nearly all the time was spent doing intermediate predicate evaluations, so optimization of predicate evaluations received particular attention. Backtracking proceeds by setting or changing one variable at a time. For our purposes, a clause is true if at least one literal in it is true or unknown. We keep a list for each literal. Each clause is on the list for one of the literals that causes it to be true. Initially, when no variables are set, any literal in a clause causes it to be true. When a variable is set or changed, one literal becomes false. The clauses on the list for that literal are examined. If a clause contains other literals that cause it to be true, then the clause is moved to the list for that literal. If all the literals in a clause are false, then the predicate is false, and it is time to backtrack. After backtracking, the clause that caused the predicate to become false is again true, so it is not necessary to move it (or any other clauses remaining on its list) to a new list. Also, there is no need to move any other clauses when backtracking. Since a clause can be left on the list of *any* literal that makes it true, there is no need to restore other clauses to their former lists. This method of evaluating intermediate predicates allows us to evaluate large predicates rapidly. For  $v = 256$ ,  $t = 4096$ , and  $s = 3$ , we estimate that only  $35 \mu s$  are required for a typical evaluation.

I haven't written any of this stuff up yet.

But the current draft has lots of cool material about Monte Carlo / Las Vegas approaches to SAT solving, like WalkSAT.

Here's an excerpt that you might find interesting:

Exercise 175 proves that such an algorithm always has an optimum cutoff value  $N = N^*$ , which minimizes the expected time to success when the algorithm is restarted after each failure. Sometimes  $N^* = \infty$  is the best choice, meaning that we should always keep plowing ahead; in other cases  $N^*$  is quite small.

But  $N^*$  exists only in theory, and the theory requires perfect knowledge of the algorithm's behavior. In practice we usually have little or no information about how  $N$  should best be specified. Fortunately there's still an effective way to pro-

ceed, by using the notion of *reluctant doubling* introduced by M. Luby, A. Sinclair, and D. Zuckerman [*Information Proc. Letters* **47** (1993), 173–180], who defined the interesting sequence

$$S_1, S_2, \dots = 1, 1, 2, 1, 1, 2, 4, 1, 1, 2, 1, 1, 2, 4, 8, \\ 1, 1, 2, 1, 1, 2, 4, 1, 1, 2, 1, 1, 2, 4, 8, 16, \dots \quad (97)$$

The elements of this sequence are all powers of 2; and we have  $S_{n+1} = 2S_n$  if and only if the number  $S_n$  has already occurred an even number of times. A convenient way to generate this sequence is to work with two integers  $(u, v)$ , and to start with  $(u_1, v_1) = (1, 1)$ ; then

$$(u_{n+1}, v_{n+1}) = (u_n \& -u_n = v_n ? (u_n + 1, 1) : (u_n, 2v_n)). \quad (98)$$

The successive pairs are  $(1, 1)$ ,  $(2, 1)$ ,  $(2, 2)$ ,  $(3, 1)$ ,  $(4, 1)$ ,  $(4, 2)$ ,  $(4, 4)$ ,  $(5, 1)$ ,  $\dots$ , and we have  $S_n = v_n$  for all  $n \geq 1$ .

Exercise 178 introduces the “reluctant Fibonacci sequence”

1, 1, 2, 1, 2, 3, 1, 1, 2, 3, 5, 1, 1, 2, 1, 2, 3, 5, 8,

1, 1, 2, 1, 2, 3, 1, 1, 2, 3, 5, 8, 13, 1, . . .

and its answer explains how to generate it with an amazingly short sequence of bitwise operations.

**185.** [46] If the given clauses are satisfiable, and if  $p > 0$ , can there be an initial  $x$  for which Algorithm W always loops forever?

(answer)

**185.** *Note to the reader:* I've been tearing my hair out trying to resolve this question. If there is a counterexample, I definitely ought to mention that fact in my book. If not, the result still is essential for a basic understanding of this important algorithm. I've been alternating between thinking I had a counterexample and thinking that there is a simple way to prove that no counterexamples exist. I can't figure out the essential reason why counterexamples have escaped all of my devious constructions. Hopefully somebody will already have resolved this issue? Help!

**185.** [46] If the given clauses are satisfiable, and if  $p > 0$ , can there be an initial  $x$  for which Algorithm W always loops forever?

(revised answer)

**185.** Equivalently, consider the following digraph on  $2^n$  vertices, one vertex for each  $x = x_1 \dots x_n$ : For all clauses  $C_j$  not satisfied by  $x$ , there's an arc from  $x$  to all vertices  $y$  that are obtainable by a flip that Algorithm W might make. (In particular, if at least one literal of  $C_j$  has cost zero, the literals of nonzero cost are *not* flipped.) Does every strong component of this digraph have an exit, unless it represents a solution?

**185.** [36] (H. H. Hoos, 1998.) If the given clauses are satisfiable, and if  $p > 0$ , can there be an initial  $x$  for which Algorithm W always loops forever?

(current answer)

**185.** (Solution by Bram Cohen, 2012.) Consider the 10 clauses  $\bar{1}234\bar{5}67$ ,  $\bar{1}2\bar{3}4\bar{5}67$ ,  $123\bar{4}5$ ,  $123\bar{4}6$ ,  $123\bar{4}7$ ,  $\bar{1}\bar{2}\bar{3}4$ ,  $\bar{1}\bar{2}\bar{3}\bar{5}$ ,  $\bar{1}\bar{2}\bar{3}\bar{6}$ ,  $\bar{1}\bar{2}\bar{4}\bar{5}$ ,  $\bar{1}\bar{2}\bar{4}\bar{6}$ , and 60 more that are obtained by the cyclic permutation  $(1234567)$ . All binary  $x = x_1 \dots x_7$  with weight  $\nu x = 2$  have cost-free flips leading to weight 3, but no such flips to weight 1. Since the only solution has weight 0, Algorithm W loops forever whenever  $\nu x > 1$ . (Is there a smaller example?)

AND NOW  
A SPECIAL REQUEST

WHAT IS YOUR FULL NAME?

....  
Rankin, Robert Alexander, 338, 351, 714.

Rao, Calyampudi Radhakrishna  
(కల్యాంపుడి రాధాకృష్ణ రావు), 518.

Rapaport, Elvira Strasser, 713.

Rashed, Roshdi (= Rashid, Rushdi)  
(رشدی راشد), 493, 812.

Raviv, Josef (יוסף רביב), 677.

....  
Yakubovich, Yuri Vladimirovich (Якубович, Юрий  
Владимирович), 402, 428.

Yan, Catherine Huafei (颜华菲), 768.

Yang Hsiung (揚雄 or 楊雄), 487–488.

Yannakakis, Mihalis (Γιαννακάκης, Μιχάλης), 604.

Yano, Tamaki (矢野環), 503, 504.

Yates, Frank, 289.

Yee, Ae Ja (이애자), 750.

....

[http://www-cs-faculty.stanford.edu/  
~knuth/  
fasc6a.ps.gz](http://www-cs-faculty.stanford.edu/~knuth/fasc6a.ps.gz)

THANKS FOR LISTENING!