



# Sudoku Players' Forums

[FAQ](#) [Search](#) [Memberlist](#) [Usergroups](#)

[Profile](#) [You have no new messages](#) [Log out \[ denis\\_berthier \]](#)

## THE REAL DISTRIBUTION OF MINIMAL PUZZLES

Goto page [Previous](#) [1](#), [2](#), [3](#) ... , [30](#), [31](#), [32](#) [Next](#)



[Sudoku Players' Forums Forum Index -> General/puzzle](#)

[View previous topic](#) :: [View next topic](#)

### Author

### Message

**eleven**

Posted: Sun Oct 04, 2009 12:30 pm Post subject:



Joined: 10 Feb 2008  
Posts: 518

I agree with PIsaacson, that the bb\_solver in any case will speed up the computation and - as said - hopefully B. Turner will allow to use it.

Its not well tested however, which speed factor we can gain from it, because bands 300-416 (only 2 mio grids) may be rather biased (i saw that because of the few U4's in band 299 suexb would be faster without U4 test in this band). To compare U4 test with bb\_solver it would be better to use Allan's low biased grids as input.

But thats not so important to invest much time.

[Back to top](#)



**gsf**

Posted: Sun Oct 04, 2009 4:16 pm Post subject:



Joined: 22 Sep 2005  
Posts: 3890  
Location: NJ USA

#### denis\_berthier wrote:

gsf,  
I've answered your PM but it seems it takes some time to be transmitted. I've corrected the "compression" in my post.

reading of PM delayed due to sleep  
haven't figured a way to code around that yet

[Back to top](#)



**gsf**

Posted: Sun Oct 04, 2009 4:21 pm Post subject:



Joined: 22 Sep 2005  
Posts: 3890  
Location: NJ USA

#### denis\_berthier wrote:

gsf wrote: what band are you up to?

I've finally used 1 to 3 of the 4 cores of my Mac in parallel but interrupted their jobs several times - so that I'll have no estimation of the generation/compression time.  
I now have a total of 2.8 Gb, bands:

1 to 38  
61 to 73  
101 to 121  
201 to 299  
and your 300-416

you might be able to use the generated file modify times within each group to help estimate

[Back to top](#)

[profile](#) [pm](#) [www](#)

**denis\_berthier**

Posted: Mon Oct 05, 2009 6:50 am Post subject:

[quote](#) [edit](#)

Joined: 19 Jun 2007  
Posts: 895  
Location: Paris, France

**eleven wrote:**

To compare U4 test with bb\_solver it would be better to use Allan's low biased grids as input.  
But thats not so important to invest much time.

We shouldn't invest too much time, but at least some time, because it is interesting to know the mean gain in speed brought by each modification.

BTW, after all these modifications, will there stay much of the original suexg code?

[Back to top](#)

[profile](#) [pm](#) [www](#)

**eleven**

Posted: Mon Oct 05, 2009 1:39 pm Post subject:

[quote](#)

Joined: 10 Feb 2008  
Posts: 518

**denis\_berthier wrote:**

We shouldn't invest too much time, but at least some time, because it is interesting to know the mean gain in speed brought by each modification.

For 10 mio grids pre-generated with suexg and an average of 10.36 U4's per grid, the optim46 version with U4 test was 1.73 times faster than without U4 test.

So using the bb-Solver additionally should gain another factor of about 1.35.

If we replace the solver, only a few lines of the original sudogen.c will remain unchanged (like input/output and calculating the order of the cells to delete). But it was a very good starting point.

[Back to top](#)

[profile](#) [pm](#)

**PIsaacson**

Posted: Tue Oct 06, 2009 1:00 am Post subject:

[quote](#)

Joined: 02 Jul 2008  
Posts: 208  
Location: Campbell, CA

**denis\_berthier wrote:**

In ~ 12 hrs, based on only the first 9 gsf files, the controlled-bias generator has produced 994 minimals (a little more than my first expectations, which now rise to ~ 2000 per day), with the following distribution:  
mean= 25.66  
standard-deviation= 1.135

**Code:**

```
nb-clues    nb-instances
21          0
22          4
23          15
24          114
25          320
26          336
27          147
28          51
29          6
30          1
31          0
```

This gives an estimated value for the number-of-clues unbiased average: 26.64

This is consistent with the mean obtained from suexg-cb (~ 0.1 standard deviations above it) considering that:

- the sample is very small,
- as explained by gsf, the subset 0-9.sudz is biased

For comparison using the latest suexg-cb from Eleven with U4 testing and the Mersenne Twister RNG + the bb\_solver (termed suexg-cb.new):

In ~ 5 hours and based on the gsf bands 1.sudz, 10.sudz - 19.sudz (probably not fully complete) the suexg-cb.new generator produced 1580 minimals on my AMD X2 4200+ running at 2.4 GHz. If I did the math right, this works out to about 7584 puzzle/day/core. I'll let it run until it catches up with my generation of \*.sudz files which are somewhere in the 200s bands now. What concerns me is that my average number of clues is only 25.63 with the following distribution curve from loading it into Excel and using the FREQUENCY function:

**Code:**

```
clues      instances
21         0
22         3
23         33
24         186
25         509
26         513
27         267
28         63
29         6
30         0
```

So the mean using the suexg-cb.new is 25.63 instead of 25.66 plus the mean number of grids processed per minimal generated is 193858 which is less than (somewhere stated) 200000+. Is this unexpected? Should I just let this thing run for a day or two to get a larger sample of minimals? I can set the priority of the shell script to a low priority and run copies in both of my cores to crank it to about 15000 puzzles/day if that helps get us somewhere... I guess I either skipped it or didn't see it, but how many of these things do you want/need anyway???

Cheers,  
Paul

[Back to top](#)



**denis\_berthier**

Posted: Tue Oct 06, 2009 4:15 am Post subject:



Joined: 19 Jun 2007  
Posts: 895  
Location: Paris, France

**eleven wrote:**

For 10 mio grids pre-generated with suexg and an average of 10.36 U4's per grid, the optim46 version with U4 test was 1.73 times faster than without U4 test.

After a 1/2 hour test of gsf|optim46-U4 and gsf|optim46, I find a 1.8 to 2.0 ratio, consistent with your 1.73.

Mike and Coloin will be happy: U4 is a very simple and effective improvement.

Replacing optim46 by optim48, a further ratio of 1.15 is obtained.

[Back to top](#)



**denis\_berthier**

Posted: Tue Oct 06, 2009 4:41 am Post subject:



Joined: 19 Jun 2007  
Posts: 895  
Location: Paris, France

**Pisaacson wrote:**

In ~ 5 hours and based on the gsf bands 1.sudz, 10.sudz - 19.sudz (probably not fully complete) the suexg-cb.new generator produced 1580 minimals on my AMD X2 4200+ running at 2.4 GHz. If I did the math right, this works out to about 7584 puzzle/day/core.

If you use piping, the 2 processes, gsf and suexg-cb-optim..., use two cores. On my Mac, yesterday evening, I launched 3 gsf|suexg-cb-... in parallel. The 4 cores were 100% used, with a ratio: 89 for the deleter - 43 for gsf. Still almost one third of the time for the generator (i.e. decompressor) is more than I expected. Take this with caution: it's only a visual observation of the activity manager. It might also be due to the 64-bit mode for the deleter and the 32-bit mode for gsf. **gsf**, do you have an estimation of any 64-bit mode improvement? [Edit: I've now found a better estimate by looking at the running times of each process. The ratio deleter/gsf is 10/4, i.e. 28% of total time for decompression.]

**Pisaacson wrote:**

What concerns me is that my average number of clues is only 25.63 with the following distribution curve from loading it into Excel and using the FREQUENCY function:

...

So the mean using the suexg-cb.new is 25.63 instead of 25.66 plus the mean number of grids processed per minimal generated is 193858 which is less than (somewhere stated) 200000+. Is this unexpected? Should I just let this thing run for a day or two to get a larger sample of minimals?

All this is within normal sampling errors.

**Pisaacson wrote:**

I can set the priority of the shell script to a low priority and run copies in both of my cores to crank it to about 15000 puzzles/day if that helps get us somewhere... I guess I either skipped it or didn't see it, but how many of these things do you want/need anyway???

It depends on your goal.

If your goal is to estimate the gain brought by the new solve function, I think you have enough.

If your goal is to generate minimals and to compare gsf's source of complete grids with suexg-cb or Allan's, *an integer number of full passes on the whole \*.sudz is necessary* in order to avoid bias. Exactly how many passes, I don't know, but a few full passes should be enough (my estimation is that one pass should produce 22,000 minimals). I think Red Ed will be very attentive to find small discrepancies in the results - so we'll be safe if he doesn't detect anything after 1 pass.

When we are sure that the various sources produce minimals with the same distribution, we can assemble the results from all of them and get better precision for the distribution, in particular the 31s.

If your goal is to estimate the proportion of 32s, or even to find ONE 32, I fear that hundreds of passes would be necessary.

I'll soon reach the 9 hrs for the current 3 parallel processes I've launched. Results in my next post.

Last edited by denis\_berthier on Tue Oct 06, 2009 6:11 am; edited 4 times in total

[Back to top](#)



**denis\_berthier**

Posted: Tue Oct 06, 2009 5:13 am Post subject:



Joined: 19 Jun 2007

Posts: 895

Location: Paris, France

### **First results with the gsf collection and optimised suexg-cb**

I've now finished generating all the .sudz files.

The procedure discussed below is gsf|suexg-cb-optim48-U4, i.e. generator: gsf (decompression of the \*.sudz files), in 32-bit mode (**gsf**: would a 64-bit mode be faster?)

piped with:

deletor: suexg-cb-optim48-U4 (Turner's solver is not yet available), compiled with the -O3 option (almost all optimisations for speed, including 64-bit mode)

Yesterday night, I launched 3 instances of this procedure in parallel on my 4-core Mac (with seeds 0, 1, 2 - the RNG is still MRW). The 4 cores have been working 100% of the time (they had nothing else to do while I was sleeping), which means that each procedure has had 1.33 cores.

Exactly 12 hrs later, I get a total of 21,943 minimals - which makes **10,971**

**minimals per core and per day** - 55 times more than with the first suexg-cb version. The good news is that a full pass on all the \*.sudz should take "only" 2 days on a single core.

The following results are not final, as only part of the .sudz files have been used up to now.

**Code:**

```
nb-clues    nb-instances
19          0
20          0
21          0
22          25
23          411
24          2697
25          6822
26          7317
27          3692
28          870
29          105
30          3
31          1
32          0
mean= 25.645
standard-deviation= 1.112
```

This leads to the estimated unbiased average = 26.547, very close to the suexg|suexg-cb result.

As for the number X of complete grids used per minimal:

$E(X) = 233,424$  (within bounds of sampling errors)

$sd(X) = 250,586$

$\min(X) = 23$

$\max(X) = 3,062,530$

Of course, all this must be taken with care, as only part of the .sudz files (estimated: 1/4 of the grids) have been used.

[Edit: results after 12 hrs instead of 9]

[Additional results: sudz decompression takes 29% of the time]

Last edited by denis\_berthier on Tue Oct 06, 2009 7:40 am; edited 3 times in total

[Back to top](#)



**Red Ed**

Posted: Tue Oct 06, 2009 7:16 am Post subject:



Joined: 06 Jun 2005

Posts: 760

Could you please start recording the total number of trials now. Then, for only the second time, we can get an unbiased estimate of the number of c-clue minimals for various c. The modification should be "free" and would obviate the need for me to do an expensive CPU burn on generating and testing \*.sudz files

(which I no longer intend to do).

[Back to top](#)



**denis\_berthier**

Posted: Tue Oct 06, 2009 7:39 am Post subject:



Joined: 19 Jun 2007  
Posts: 895  
Location: Paris, France

**Red Ed wrote:**

Could you please start recording the total number of trials now. Then, for only the second time, we can get an unbiased estimate of the number of  $c$ -clue minimals for various  $c$ . The modification should be "free" and would obviate the need for me to do an expensive CPU burn on generating and testing \*.sudz files (which I no longer intend to do).

I've always recorded this number since eleven implemented this possibility. The last lines in my previous post refer to it.

[Back to top](#)



**eleven**

Posted: Tue Oct 06, 2009 11:23 am Post subject:



Joined: 10 Feb 2008  
Posts: 518

Nice, that you have it running now, Denis.  
This are stats from the 12 h sample, already looks trustable (though the main part of grids is not tried yet):

**Code:**

```
21943 puzzles, avg 233424 tries

      estim/grid      %
22:      25  1.81e+11  0.0036
23:     411  7.66e+12  0.1517
24:    2697  1.21e+14  2.4067
25:    6822  7.00e+14  13.8802
26:    7317  1.62e+15  32.065
27:    3692  1.66e+15  32.9578
28:     870  7.55e+14  14.9779
29:     105  1.67e+14  3.3037
30:        3  8.25e+12  0.1636
31:        1  4.53e+12  0.0897

Sum 5.04398e+15 estim. puzzles/grid, average 26.55
clues/puzzle
```

[Back to top](#)



**Red Ed**

Posted: Tue Oct 06, 2009 12:23 pm Post subject:



Joined: 06 Jun 2005  
Posts: 760

**eleven wrote:**

**Code:**

```
26:    7317  1.62e+15  32.065
27:    3692  1.66e+15  32.9578
```

Interesting. More 27s than 26s, somewhat contrary to my results; and more in total by quite a long way compared to my estimates. This just adds weight to what Denis and I both said: you can't just sample the first few \*.sudz files -- you have to do them all.

Someone should consider adding a sampling probability parameter,  $p$ , to *suexg-cb-with-added-oompf* so that each solution grid is used (otherwise: ignored) with probability  $p$ . Then you could do a very-very-nearly-unbiased swoosh across the whole set of \*.sudz files in less time, which would be handy for posts like the previous one that look to get a quick view of the statistics.

[Back to top](#)



**eleven**

Posted: Tue Oct 06, 2009 1:10 pm Post subject:



Joined: 10 Feb 2008  
Posts: 518

**Red Ed wrote:**

Someone should consider adding a sampling probability parameter,  $p$ , to *suexg-cb-with-added-oompf* so that each solution grid is used (otherwise: ignored) with probability  $p$ . Then you could do a very-very-nearly-unbiased swoosh across the whole set of \*.sudz files in less time, which would be handy for posts like the previous one that look to get a quick view of the statistics.

This would do it with an (optional) 4th parameter, which gives the probability between 0 and 1.

After

**Code:**

```
int nClues;
```

insert

**Code:**

```
double prob_to_take_grid = 2.0;
```

Insert

**Code:**

```
if (argc>4){
    char* tmp;
    prob_to_take_grid = strtod(argv[4], &tmp);
}
```

before

**Code:**

```
for(i=0;i<888;i++){j=1;while(j<=i)j+=j;Two[i]=j-1;}
```

Replace

**Code:**

```
m6:if(argc>3)
    for(i=1;i<=81;i++){
        m6a:A0[i]=fgetc(file)-48;if(feof(file))return(8);
        if(A0[i]==-2)A0[i]=0;if(A0[i]<0 || A0[i]>9)goto
m6a;}
```



by

**Code:**

```

m6:if(argc>3)
  while(1) {
    for(i=1;i<=81;i++){
      m6a:A0[i]=fgetc(file)-48;if(feof(file))return(8);
      if(A0[i]==-2)A0[i]=0;if(A0[i]<0 || A0[i]>9)goto
m6a;}
    if (prob_to_take_grid > 1.0)
      break;
    if (genrand_int32()/4294967295.0 <
prob_to_take_grid)
      break;
  }

```

Note, that then the parameters

0 100 - 0.001

would mean, that each grid is taken with probability 1/1000, but when chosen, it is tried 100 times.

[Back to top](#)**Red Ed**

Posted: Tue Oct 06, 2009 1:19 pm Post subject:



Nice; thanks.

Joined: 06 Jun 2005

Posts: 760

I suppose the bottleneck for small samples is now the decompression step. So I guess the overall speedup coming from this new parameter is a factor of at most 3 on top of what's already been done.

[Back to top](#)

Display posts from previous:



**Sudoku Players'**

**Forums Forum**

**Index ->**

**General/puzzle**

All times are GMT  
Goto page [Previous](#) [1](#), [2](#), [3](#) ... , [30](#), [31](#), [32](#) [Next](#)

**Page 31 of 32**

[Stop watching this topic](#)

Jump to:

You **can** post new topics in this forum  
 You **can** reply to topics in this forum  
 You **can** edit your posts in this forum  
 You **can** delete your posts in this forum  
 You **can** vote in polls in this forum

Powered by phpBB © 2001, 2005 phpBB Group