Sudoku Players' Forums

THE REAL DISTRIBUTION OF MINIMAL PUZZLES

Goto page Previous 1, 2, 3 ... , 31, 32, 33 Next

newtopic postreply **Sudoku Players' Forums Forum Index -> General/puzzle**

View previous topic :: View next topic

Author	Message
eleven Joined: 10 Feb 2008 Posts: 521	Posted: Tue Oct 06, 2009 8:35 am Post subject: quote **Red Ed wrote:** 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. When both programs are run in 32-bit mode, the decompression only needs 7-8 % of the cpu time. But e.g. with p=1/10 not only this raises above 30%, but also the input from the pipe becomes very expensive (with file input you also could use fseek to jump over the grids you dont take). So i guess, that you would need about the double time then (1/5 of the time to get 1/10 of the puzzles). Back to top profile pm
gsf Joined: 21 Sep 2005 Posts: 3896 Location: NJ USA	Posted: Tue Oct 06, 2009 8:38 am Post subject: quote **denis_berthier wrote:** 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?) I did a test on linux.i386-64 comparing a 64 bit and 32 bit executable on 006.sudz looks like a 64 bit executable will run 33% faster but at the moment I don't have access to a 64 bit intel mac later today I'll put together a src package with a standalone decompressor **sudz** Back to top profile pm www
denis_berthier	Posted: Tue Oct 06, 2009 8:45 am Post subject: quote I've had no time for Sudoku today.

Joined: 19 Jun 2007
Posts: 901
Location: Paris, France

Coming home, I find nothing really new with the puzzles produced since my last post.
I deliberately didn't give the "unbiased" distribution because, with only part of the .sudz explored, there's no guarantee of having no systematic bias.

> **Red Ed wrote:**
>
> 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.

True (last sentence) but the 500,000 controlled-bias minimals produced with the original suexg-cb allowed to conclude that there are a little (1.7%) more 27s than 26s in the real distribution. This was true in all the initial sub-samples I tried from 10,000 to 500,000. So this was a very stable result.

> **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.

One can always do this, but is it really worth?
After spending 2 weeks generating the .sudz files, can't we wait 2 more days for the first results with the whole collection?
The results obtained thus far are stable over time and therefore don't show any obvious dependency on the band.

Back to top profile pm www

Red Ed

Posted: Tue Oct 06, 2009 9:26 am Post subject: quote

Joined: 06 Jun 2005
Posts: 763

> **denis_berthier wrote:**
>
> > **Red Ed wrote:**
> >
> > More 27s than 26s, somewhat contrary to my results ...
>
> ... but the 500,000 controlled-bias minimals produced with the original suexg-cb allowed to conclude that there are a little (1.7%) more 27s than 26s in the real distribution.

Gosh, yes, the suexg-cb results do say that. That's very strange. I don't understand why my experiments using suexg as the solution grids source gave counts the other way round. Will have to look into that.

> **Quote:**
>
> > **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.
>
> One can always do this, but is it really worth?
> After spending 2 weeks generating the .sudz files, can't we wait 2 more days for the first results with the whole collection?

It's worth it if you want to be able to do other sorts of tests (other than number of clues, and complexity) relatively quickly. Perhaps you don't want that flexibility, but others might. For another five minutes of coding, what's the issue?

Back to top profile pm

eleven

Posted: Tue Oct 06, 2009 11:20 am Post subject: quote

Joined: 10 Feb 2008
Posts: 521

> **denis_berthier wrote:**
>
> I've had no time for Sudoku today.

Oh yes, you are French. Whats her name ?

Back to top profile pm

denis_berthier

Posted: Tue Oct 06, 2009 9:57 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 901
Location: Paris, France

> **eleven wrote:**
>
> > **denis_berthier wrote:**
> >
> > I've had no time for Sudoku today.
>
> Oh yes, you are French. Whats her name ?

... teaching
Today again.

Nothing new this morning: ~ same distribution of clues.
But mean number of grids per minimal is now higher.

Back to top profile pm www

denis_berthier

Posted: Tue Oct 06, 2009 9:59 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 901
Location: Paris, France

> **gsf wrote:**
>
> later today I'll put together a src package with a standalone decompressor **sudz**

That would be very useful. Thanks.

Back to top profile pm www

denis_berthier

Posted: Tue Oct 06, 2009 10:34 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 901

FINAL RESULTS FOR THE CONTROLLED-BIAS SUEXG-CB GENERATOR

Location: Paris, France

I realise that the last results I've published for suexg-cb (before we started optimising it and changing the source of complete grids) are at the bottom of this page: http://www.sudoku.com/boards/viewtopic.php?t=14615&postdays=0&postorder=asc&start=315 and bear on only 250,000 puzzles.
I was busy with the improvements discussed above but I now have 500,000 minimals generated before them: 180,000 with the original suexg-cb version and the rest with optim46 (the first 180,000 weren't useless, as they justify using optim46, i.e. deleting the first 46 clues without doing any test).

Here are the results for the number-of-clues distribution. These are the final results as we now have a faster generator and I won't continue to generate more puzzles with this version of suexg-cb. But faster doesn't mean better or worse and the new generators don't invalidate the following results.

Code:

```
#clues  raw-dist   unbiased-dist * 1,000,000
19      0          0.0(*)
20      2          0.0158 (*)
21      4          0.0920 (*)
22      615        38.589
23      9848       1585
24      60576      23563
25      154024     136602
26      168070     321050
27      83911      326513
28      20234      151844
29      2566       35193
30      147        3495
31      3          117 (*)
32      0          0 (*)

* values based on few data are not reliable.

raw mean= 25.65
raw standard-deviation= 1.120

unbiased mean= 26.56
unbiased standard-deviation= 1.113
```

The estimated mean SER and NRCZT are unchanged
SER:
raw-average = 4.147 unbiased-average = 4.48
raw-standard-deviation = 2.49 unbiased-standard-deviation = 2.53

NRCZT:
raw-average = 2.135 unbiased-average = 2.308
raw-standard-deviation = 1.340 unbiased-standard-deviation = 1.379

As usual, more detailed results, with html tables, will appear on my web pages,

as soon as I find some time for it.

Back to top profile pm www

Red Ed

Posted: Wed Oct 07, 2009 2:24 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 763

An idea for faster estimation of the number of 31-clue minimals ...

Pick a smaller (say 24-clue) subgrid of a solution grid. We're going to loop over all 31-clue supersets of that and will record all minimal proper puzzles that pop out. The problem is to find a way of doing this quickly.

First, do some preprocessing to check that the subgrid is minimal (i.e. no clue is implied by the others) and to find all +1 and +2 supersets that are also minimal. The latter check gives lists of single clues and pairs of additional clues that are "safe". Now loop over sets of 7 (say, if the subgrid size was 24) clues such that each clue and each clue pair within the set of 7 is safe; and, for each safe set, check first that there's a unique solution; then check that the solution is minimal.

This should be faster than the current method for 31s because you can prune the search space with those "safe" clues *and* certain classes of solver (e.g. template-based) can do a lot of preprocessing in recognition of the constant base of 24 (say) clues, so each subsequent unique-solution test can be that much quicker.

At least that's the theory. I'll try coding it over the next few days. Any comments before I start?

Back to top profile pm

denis_berthier

Posted: Wed Oct 07, 2009 9:01 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 901
Location: Paris, France

Red Ed,
Do you consider multi-sol 24-clue subgrids?

Back to top profile pm www

denis_berthier

Posted: Wed Oct 07, 2009 9:18 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 901
Location: Paris, France

ESTIMATES for the NUMBER of N-CLUE MINIMAL PUZZLES
I'm not really interested in the absolute number of minimals, but rather in their proportions.
Neverteless , the above suexg-cb results give the following estimates:

Code:

```
#clues    #minimals
22        1.35529578042937e+33
23        5.5671277334354e+34
24        8.27561904095325e+35
25        4.79759013393657e+36
26        1.12755990408589e+37
27        1.14674603199864e+37
```

```
28        5.33292926612237e+36
29        1.23600037208473e+36
30        1.22733006549957e+35
31        4.1207263555285e+33 (*)
* values based on a small sub-sample are not reliable
```

Back to top profile pm www

Red Ed

Posted: Wed Oct 07, 2009 10:31 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 763

Re multi-solution 24-clue subgrids -- yes, *only* multi-solution ones.

Re estimates -- thanks -- for direct visual comparison with what's gone before, could you edit (or in future post) the per-grid estimates instead, i.e. divided by ~6.67e21. It just saves the reader a minute or so with a calculator. Also, can you give the standard deviation of each estimate. Ta.

Back to top profile pm

David P Bird

Posted: Thu Oct 08, 2009 12:39 am Post subject: quote

Joined: 16 Sep 2008
Posts: 165
Location: Middle England

Withdrawn - way off beam

Last edited by David P Bird on Thu Oct 08, 2009 4:12 am; edited 1 time in total

Back to top profile pm

JPF

Posted: Thu Oct 08, 2009 1:17 am Post subject: quote

Joined: 06 Dec 2005
Posts: 2921
Location: Paris, France

denis_berthier wrote:

```
#clues #minimals
22 1.35529578042937e+33
23 5.5671277334354e+34
24 8.27561904095325e+35
25 4.79759013393657e+36
26 1.12755990408589e+37
27 1.14674603199864e+37
28 5.33292926612237e+36
29 1.23600037208473e+36
30 1.22733006549957e+35
31 4.1207263555285e+33 (*)
* values based on a small sub-sample are not reliable
```

Whoa, 15 significant digits ...
What does it mean ?

JPF

Back to top profile pm

eleven

Posted: Thu Oct 08, 2009 2:42 am Post subject: quote

Red Ed wrote:

Joined: 10 Feb 2008
Posts: 521

> First, do some preprocessing to check that the subgrid is minimal (i.e. no clue is implied by the others) ...

I would not know how to do this quickly. suexk needed 53 sec to calculate the numbers of solutions for 240 multisolution 24's on my PC, so one test would need about 5 secs.

Back to top profile pm

Display posts from previous: All Posts Oldest First Go

newtopic postreply Sudoku Players' Forums Forum Index -> General/puzzle

All times are GMT - 8 Hours
Goto page Previous 1, 2, 3 ... , 31, 32, 33 Next

Page 32 of 33

Jump to: General/puzzle Go

You **cannot** post new topics in this forum
You **cannot** reply to topics in this forum
You **cannot** edit your posts in this forum
You **cannot** delete your posts in this forum
You **cannot** vote in polls in this forum