Sudoku Players' Forums

THE REAL DISTRIBUTION OF MINIMAL PUZZLES

Goto page Previous 1, 2, 3 ... , 35, 36, 37 Next

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

View previous topic :: View next topic

Author	Message
PIsaacson Joined: 02 Jul 2008 Posts: 216 Location: Campbell, CA	Posted: Mon Oct 12, 2009 7:02 pm Post subject: quote Brian Turner has released his code on the Sudoku Programmers' Forum http://www.setbb.com/phpbb/viewtopic.php?t=1824 I am having difficulty uploading my modified code to my hosting site FileAve.com. The login page is timing out for some reason, so if anyone wants the modified library code to adapt their various versions of suexg etc. if you PM me with an e-mail address, I can send you a zip file containing the modified bb_solver library, my modified suexg-cb, the makefile (using g++) and my comments on the modifications. Otherwise, as soon as FileAve is reachable, I'll upload the same zip file and post the location here for downloading. Cheers, Paul Okay - it's now uploaded and available at http://pisaacson.fileave.com/Sudoku/suexg-cb.zip
Back to top	profile pm

Red Ed
 	
Joined: 06 Jun 2005
Posts: 790
	Posted: Tue Oct 13, 2009 7:07 am Post subject: quote
	
	Results from 10.5 hours overnight on a single 1.4GHz core:
	
	Code:
	`jct[ 0 ] = 21208`
	`jct[ 1 ] = 40`
	`jct[ 2 ] = 24`
	`jct[ 3 ] = 11`
	`jct[ 4 ] = 9`
	`jct[ 6 ] = 2`
	`jct[ 7 ] = 5`
	`jct[ 8 ] = 1`
	`jct[ 9 ] = 2`
	`jct[ 11 ] = 2`
	`jct[100+] = 0`

(jct[n] = number of 24s that yielded n 31s)
The mean is about 0.012 31s per 24, so estimated number of 31s per grid = 0.01182876*choose(81,24)/choose(31,24) = about 1e12. Denis' results after 31 full scans of gsf's collection say ~0.55e12. I've not evaluated the p-value of the discrepancy, but I doubt it's significant.

The standard deviation of my estimator on a sample of this size is about 2e10, or about 2%. The top-down controlled-bias method would need to find about 2500 31s to achieve the same relative sample error, so about 330 full scans of gsf's collection.

In other words, if I've got my stats right (done in a rush ...), for the purposes of estimating the number of proper minimal 31s per grid (and only that part of the distribution), **the supersets method runs at a rate equivalent to about 750 controlled-bias scans of gsf's collection per day**.

Last edited by Red Ed on Tue Oct 13, 2009 7:08 am; edited 1 time in total

Back to top profile pm

denis_berthier

Posted: Tue Oct 13, 2009 7:08 am Post subject: quote edit

Joined: 19 Jun 2007
Posts: 925
Location: Paris, France

CONTROLLED-BIAS GENERATOR: WHAT ARE THE NEXT STEPS?

In a recent post, only 10 days ago,
http://www.sudoku.com/boards/viewtopic.php?t=14615&postdays=0&postorder=asc&start=446
I tried to define my roadmap with the controlled-bias generator.
In these 10 days, much has been done.
I've used extensively:
- gsf's list and decompressor as the source of complete grids,
- suexg-cb-optim48-U4 as the deletor.
This combination is relatively fast (very fast compared with the first version of suexg-cb), as it can produce about 16,000 minimals per day on a 2.66 GHz CPU.
Further optimisations may be expected if Brian Turner's faster solver becomes available and still larger collections of minimals can be computed in a few days.
I now have 750,000 minimals and I'll soon have 1,250,000.
The question now is (at least for me): how far should we (I) go?

From the results I already published (http://www.sudoku.com/boards/viewtopic.php?t=14615&start=497), based on 650,000 minimals, it appears that **what's important for most statistics is the proportions of puzzles with a number of clues in the interval [23, 30]**.
Below or above this interval, the (estimated unbiased) proportion of minimals is so small that it can hardly have any impact on statistics - except, of course, if you're interested in statistics of extremes - and in particular it can have no impact on what I'm interested in, i.e. statistics on the complexity of puzzles:
- for 22 clues: 0.0033% +/- 0.00125%
- for 31 clues: 0.012% +/- 0.006%

It can also be seen that the boundary values of this interval can have some impact, but their precision is already fine with the current sample:
n=23: 0.149% +/- 0.00134%
n=30: 0.40% +/- 0.0267%.

It turns out that the smaller precision is
- on 27s: 32.77% +/- 0.098%
- on 28s: 15.40% +/- 0.093%
- on 26s: 31.98% +/- 0.068%
(Well, that's already a rather good precision: 1 in 1000)

As a result, no specific method for looking preferentially for 30s (or above) is needed, the best strategy for improving the global precision is still to run the current gsf/suexg-cb combination (with suexg-cb optimisations as they become available). I'll let it run a little more, but the end of this all is clearly approaching.

Back to top profile pm www

Red Ed

Posted: Tue Oct 13, 2009 7:11 am Post subject: quote

Ha ha, timely cross-post.

Joined: 06 Jun 2005
Posts: 790

Let me make clear that my estimate of the number of 31s is not produced in order to make the overall estimate of the mean number of clues better; it's done to test the right tail of the distribution and for the academic pleasure of developing a method that can reach out further, faster, than any before it.

Back to top profile pm

denis_berthier

Posted: Tue Oct 13, 2009 7:27 am Post subject: quote edit

Joined: 19 Jun 2007
Posts: 925
Location: Paris, France

> **Red Ed wrote:**
>
> if I've got my stats right (done in a rush ...), for the purposes of estimating the number of proper minimal 31s per grid (and only that part of the distribution), **the supersets method runs at a rate equivalent to about 750 controlled-bias scans of gsf's collection per day**.

Isn't this a little hasty conclusion (in bold), considering how "unlucky" you were with the 27s?
The controlled-bias generator outputs minimals of any number of clues. With your method, you have to restart your generator for each number of clues. The comparison of computation times is therefore not meaningful.

> **Red Ed wrote:**
>
> Let me make clear that my estimate of the number of 31s is not produced in order to make the overall estimate of the mean number of clues better

As shown in my previous post, that'd be mostly useless, and not only for the mean number of clues, but for any random variable that is not concentrated on

extreme values.

> **Red Ed wrote:**
>
> it's done to test the right tail of the distribution and for the academic pleasure of developing a method that can reach out further, faster, than any before it.

Provided that it is unbiased, which remains to be proved.

How many complete grids did you use?
How many 24-clue subgrids did you use for each of them?

Back to top profile pm www

denis_berthier

Joined: 19 Jun 2007
Posts: 925
Location: Paris, France

Posted: Tue Oct 13, 2009 7:43 am Post subject: quote edit

> **PIsaacson wrote:**
>
> it's now uploaded and available at
> http://pisaacson.fileave.com/Sudoku/suexg-cb.zip

Thanks for your work.
Is there any way of having a stream as input, instead of a file of complete grids?

Back to top profile pm www

Red Ed

Joined: 06 Jun 2005
Posts: 790

Posted: Tue Oct 13, 2009 7:45 am Post subject: quote

No time to respond now, except to your "useless" comment, which comes across as just sour grapes. My interest in this is mostly academic and it's not important to me whether my interests overlap yours. As it happens, there are people out there interested in the tail of the distribution -- why else would there be a maximum number of clues thread? But that's incidental to my main point: you are not the sudoku community's arbiter of taste, Denis, so please think twice before grumbling that something's "useless" in future. Ta ta.

Back to top profile pm

denis_berthier

Joined: 19 Jun 2007
Posts: 925
Location: Paris, France

Posted: Tue Oct 13, 2009 7:51 am Post subject: quote edit

> **Red Ed wrote:**
>
> No time to respond now, except to your "useless" comment, which comes across as just sour grapes. My interest in this is mostly academic and it's not important to me whether my interests overlap yours. As it happens, there are people out there interested in the tail of the distribution -- why else would there be a maximum number of clues thread? But that's incidental to my main point: you are not the sudoku community's arbiter of taste, Denis, so please think twice before grumbling that something's "useless" in future.

Ah, Red Ed, you should read twice what's written before reacting so violently!
I didn't write "useless", but "mostly useless, and not only for the mean number of clues, but for any random variable that is not concentrated on extreme

values".
So, where does that contradict what you're saying? Is "extreme values" different from "tail of the distribution"?

Back to top profile pm www

denis_berthier

Joined: 19 Jun 2007
Posts: 925
Location: Paris, France

Posted: Tue Oct 13, 2009 8:13 am Post subject: quote edit

Paul,
How can we build on Unix?
I tried your "make -f suexg-cb.mk". But I get the following errors:

Code:

```
g++ -c -g -O3 -MD -enable-auto-import -DWIN32 -I.
suexg-cb.cpp
suexg-cb.cpp: In function 'int main(int, char**)':
suexg-cb.cpp:251: warning: format '%Lf' expects type
'long double*', but argument 3 has type 'double*'
g++ -c -g -O3 -MD -enable-auto-import -DWIN32 -I.
bb_sudoku_solver.cpp
ar -r libbbsolver.a bb_sudoku_solver.o
ar: creating archive libbbsolver.a
g++ -g -O3 -MD -enable-auto-import -DWIN32 -I. -
osuexg-cb.exe suexg-cb.o libbbsolver.a
ld: unknown option: -enable-auto-import
collect2: ld returned 1 exit status
make: *** [suexg-cb.exe] Error 1
```

Back to top profile pm www

PIsaacson

Joined: 02 Jul 2008
Posts: 216
Location: Campbell, CA

Posted: Tue Oct 13, 2009 8:38 am Post subject: quote

Denis,

For compiling on Unix/Linux/Mac OS, I think the following makefile should work:

Code:

```
CFLAGS = -g -O3 -MD
IFLAGS = -I.

.SUFFIXES: .o .cpp

CSRC = \
  suexg-cb.cpp bb_sudoku_solver.cpp

COBJ = $(CSRC:.cpp=.o)

CPP = g++
AR = ar

all : clean suexg-cb.exe

clean ::
        @rm -fr $(COBJ) suexg-cb.exe libbbsolver.a

libbbsolver.a : bb_sudoku_solver.o
        $(AR) -r $@ $^

suexg-cb.exe : suexg-cb.o libbbsolver.a
        $(CPP) $(CFLAGS) $(DFLAGS) $(IFLAGS) -o$@ $^
```

```
.cpp.o ::
        $(CPP) -c $(CFLAGS) $(DFLAGS) $(IFLAGS) $<
```

Makefile(s) require a tab following a dependency, so make sure you either carefully enter and hit tab after entering the clean/libbbsolver/suexg-cb.exe/.cpp.o lines for the start of the dependency action @rm/$(AR)/$(CPP)..., or use an editor that doesn't replace hard tabs with spaces.

Alternatively, you could force a compile/link with the single command:

Code:

```
g++ -O3 -I. suexg-cb.cpp bb_sudoku_solver.cpp -o suexg-cb.exe
```

The suexg-cb.cpp file contains a format statement for the 4th runtime argument (Red Ed's request for a probability hit/skip), so a quick fix would be to simply comment out the sscanf (argv[4]...) line. I'm unfamiliar with the formats other than GNU, which uses %Lf for doubles. Based on the error, if you want to retain the probability mod you could try changing the format to just use "%f" and see if that compiles cleanly.

The stream option is already there. The code is based on what I think is the latest from Eleven, but he should look it over and decide on the actual content. Much can be eliminated since the DLX code is no longer there, so there's no need for the various DLX support variables/tables/initialization...

Hope this works...
Paul

Back to top profile pm

Red Ed

Posted: Tue Oct 13, 2009 4:23 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 790

More "mostly useless" results for this whole "mostly useless" endeavour

In just over 8 hours, I got an estimated number of 30s equal to about 2.1e13 with standard error about 10% of that (which is consistent with Denis' estimate). For 19-20 hours work, I could get the standard error down to around 6%, the same as Denis' uncertainty obtained after 31 full scans of gsf's collection. That's an equivalent rate for estimation of the number of proper minimal 30s of around **35-40 full scans per day** on a single 1.4GHz CPU.

It may be possible to search for all 30+ clue minimals concurrently without the "+" part adding much to the runtime.

Back to top profile pm

denis_berthier

Posted: Tue Oct 13, 2009 4:33 pm Post subject: quote edit

Joined: 19 Jun 2007
Posts: 925

Red Ed,
But you carefully avoided to answer my questions related to bias:
how many complete grids?

Location: Paris, France

how many 24-clue subgrids for each?

For the rest of your claims, nothing of them can be checked by anybody. So I don't really care if you answer or not.

Back to top profile pm www

denis_berthier

Posted: Tue Oct 13, 2009 4:35 pm Post subject: quote edit

Joined: 19 Jun 2007
Posts: 925
Location: Paris, France

Paul,
I could compile your code (after deleting the argv[4]) and run it in conjunction with gsf complete grids. I also changed optim46 into optim48. Seems to be quite fast.
Tomorrow, I'll do a more extensive test for speed.

Back to top profile pm www

Red Ed

Posted: Tue Oct 13, 2009 7:50 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 790

> **denis_berthier wrote:**
> how many complete grids?
> how many 24-clue subgrids for each?

One 24 per complete grid; 21208 yielded no 31s, 40 yielded one 31, etc.

> **denis_berthier wrote:**
> For the rest of your claims, nothing of them can be checked by anybody. So I don't really care if you answer or not.

What do you want to check? That my code is producing an estimate consistent with yours? That my code is as fast as I claim? That the implementation doesn't contain a subtle bug? All that I recall **you** asking for is proof that the algorithm is unbiased: and quite why you can't understand that it is, given the existing descriptions (1 and 2), I don't know -- would you like me to elaborate on a part of it?

I intend to post code, if that's what you're getting at. But I need to do more testing first. It's fairly complex and I've already caught a couple of bugs (one affecting only speed, the other possibly affecting accuracy) since posting the two sets of interim results (for 30s and for 31s). The corrected code is running now: by the morning, I should have a good estimate of the number of 30s for comparison with your counts.

btw, well done for indicating the (informal) 1-sigma standard error in your recent post.

Back to top profile pm

denis_berthier

Posted: Wed Oct 14, 2009 5:43 am Post subject: quote edit

Joined: 19 Jun 2007
Posts: 925

> **Red Ed wrote:**
> > **denis_berthier wrote:**

Location: Paris, France

> how many complete grids?
> how many 24-clue subgrids for each?

One 24 per complete grid

First possible source of bias: one 24-clue per complete grid. We know that the suexg algorithm is strongly biased because it produces 1 minimal per complete grid.
Why choosing one 24 (not minimal) per complete grid would not introduce a bias?

Second possible source of bias: the source of complete grids; they don't all have the same numbers of 24s and all their 24s don't have the same number of 31 minimals. So, same question: how many complete grids?
[Edit: same as #24s: 22,000 - a very small number]

Red Ed wrote:

> What do you want to check? That my code is producing an estimate consistent with yours? That my code is as fast as I claim? That the implementation doesn't contain a subtle bug

I'm totally unable to check your C implementation, but others would do it.
For the estimate consistent with mine, it can be interesting for you, but it couldn't bring much to me now that I have more than a million minimals.
Concerning speed, yes, I'd like to know the conditions:
- does your estimate of the 30s relies on an estimate of the 24s?
- what if you had to produce the full number-of-clues distribution?
- how does it scale with precision?

Red Ed wrote:

> well done for indicating the (informal) 1-sigma standard error in your recent post.

That was your suggestion. When you are constructive, I can hear you.

Last edited by denis_berthier on Wed Oct 14, 2009 6:18 am; edited 1 time in total

Back to top profile pm www

Display posts from previous: All Posts Oldest First Go

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

All times are GMT
Goto page Previous 1, 2, 3 ... , 35, 36, 37 Next

Page 36 of 37

Stop watching this topic

Jump to: General/puzzle Go

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