Sudoku Players' Forums

THE REAL DISTRIBUTION OF MINIMAL PUZZLES

Goto page Previous 1, 2, 3 ... , 14, 15, 16 Next

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

View previous topic :: View next topic

Author	Message
denis_berthier Joined: 19 Jun 2007 Posts: 799 Location: Paris, France	Posted: Thu Sep 17, 2009 11:31 pm Post subject: quote **David P Bird wrote:** It's a question of view point isn't it? Right. Either you only want a numerical analysis or you also want a picture of it. My proof (either the first or the modified one) introduces an optional visual probabilistic picture of the algorithm. Once more, let's not be distracted from the fundamental result: the controlled-bias generator is a reality, it (slowly) produces puzzles, from which unbiased statistics can be computed, and there's currently no better way of doing so. Back to top profile pm www
Red Ed Joined: 06 Jun 2005 Posts: 703	Posted: Fri Sep 18, 2009 10:08 am Post subject: quote When viewed my way, you can think of the algorithm as: 1. generate solution grid 2. generate all subgrids along a path from 81 clues to 0 clues 3. find the biggest subgrid that has a single solution 4. check for minimality Clearly steps 3 and 4 are the expensive ones. Step 3 involves finding n such that subgrids S(n) and S(n+1), of n and n+1 clues respectively, have solns(S(n))>1 and solns(S(n+1))=1. The current implementation tries n=80, 79, 78, ... in that order. An implementation improvement may therefore be to try more likely values of n first. An alternative implementation improvement in step 3 might be to try n=0,1,2,3,... but to save state after finding the first solution in each case. So the solution finding always moves "forwards" in its search, but the search space gets reduced at each step. In theory, suexg can be modified to do that (but who

wants to take on the challenge!).

There are probably some good short-cuts for step 4, too. Most of the time, the puzzle isn't minimal: so if we could find a good way of spotting which clues are most likely to be redundant then we should test those first.

Back to top profile pm

denis_berthier

Posted: Fri Sep 18, 2009 12:33 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 799
Location: Paris, France

Red Ed,
This is not the way suexg-cb works.
The path is not really generated from 81 down to 0 clue. It stops when a minimal or a multi-clue is encountered. What would happen after that if the algorithm wasn't stopped is purely virtual.
See my web page for the full description of the real and the virtual algorithms that have the same output.

Your "improvement" would not guarantee the controlled bias formula. Starting from a complete grid and going downwards step by step is essential for that.

Back to top profile pm www

Red Ed

Posted: Fri Sep 18, 2009 1:02 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 703

denis_berthier wrote:

> Your "improvement" would not guarantee the controlled bias formula. Starting from a complete grid and going downwards step by step is essential for that.

Maybe my explanation wasn't great, but I'm quite sure that my proposed implementation improvements all preserve the "controlled-bias" ratio of probabilities. The proof that they do so is exactly that given by me a couple of (my) posts ago.

To clear up one possible ambiguity: the "path" referred to in step 2 of my previous post is a random sequence of deletions from 0 clues deleted (full grid) to 81 clues deleted (empty grid). It's chosen in advance of step 3, with no guarantee that the smallest single-solution subgrid in the path should be minimal: if it's *not* minimal then no output is generated.

Presumably you only object to the optimisation of step 3, not the optimisation of step 4?

Back to top profile pm

coloin

Posted: Fri Sep 18, 2009 2:00 pm Post subject: quote

Joined: 05 May 2005
Posts: 1070
Location: Devon UK

Apologies if this is a "spanner" in the works....

We would expect the randomly generated/controlled bias puzzles to appear to come from random grids wouldnt we ?

I wonder if this is actually true.....or rather I hope it is.

Denis....have you published a collection of these puzzles ? I had a look and couldnt find them.

C

Back to top profile pm

denis_berthier

Posted: Fri Sep 18, 2009 7:53 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 799
Location: Paris, France

> **Red Ed wrote:**
>
> To clear up one possible ambiguity: the "path" referred to in step 2 of my previous post is a random sequence of deletions from 0 clues deleted (full grid) to 81 clues deleted (empty grid). It's chosen in advance of step 3, with no guarantee that the smallest single-solution subgrid in the path should be minimal: if it's *not* minimal then no output is generated.

Oh, so you are reverting to keeping track of clue deletions, as in my approach, instead of grids. (This is equivalent, of course.)
In the real algorithm, we don't have to choose a full path down to 0 clue. As the raw mean number of clues is 25, this spares 25 useless choices in the mean. On this point, your "improvement" would be regressive.
Step 3 is already done in suexg-cb, without going down to 0 clue.

> **Red Ed wrote:**
>
> Presumably you only object to the optimisation of step 3, not the optimisation of step 4?

I don't see you propose any optimisation in step 4.

Of course any real optimisation would be welcome.

Back to top profile pm www

denis_berthier

Posted: Fri Sep 18, 2009 8:24 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 799
Location: Paris, France

> **coloin wrote:**
>
> We would expect the randomly generated/controlled bias puzzles to appear to come from random grids wouldnt we ?

I guess, by "random", you mean "unbiased".

> **coloin wrote:**
>
> I wonder if this is actually true.....or rather I hope it is.

This point has already been discussed and I've given several complementary answers. To summarise:

1) NO. One should not expect the solution grids of an unbiased collection of

minimal puzzles to be unbiased. If it was unbiased, this could mean that all the complete grids should have ~ the same number of minimal puzzles, which is known to be false.

2) I've shown that the number of clues of minimal puzzles is uncorrelated with Red Ed's 3322 unbiasedness tests of their solution grids.

3) The complete grid generator part of the algorithm has little impact on the generation of puzzles.
3a) I've shown this for the standard top-down algorithm: suexg's or Allan's (which are based on very different principles) give the same results (see my web pages).
3b) Red Ed has been unable to show any significant difference in the distribution of clues when he changed the (supposedly unbiased) generator part of his own algorithm (which is more complex and doesn't output puzzles) for suexg's.

4) For the controlled-bias generator, only 1 in 225,000 complete grids produces a minimal puzzle. The algorithm acts as a second randomisation process.

> **coloin wrote:**
> Denis....have you published a collection of these puzzles ? I had a look and couldnt find them.

Don't worry. I'll publish everything in due time, as I've always done (contrary to some). But please, be patient, this is work in progress, I haven't even finished generating all the puzzles I need or computing all the ratings I need. What I've given until now is only preliminary results.
If you're mainly interested in the generation part of suexg-cb, it is the same as that of suexg, for which you already have the 1,000,000 grids of the sudogen0_1M collection.
Finally, suexg-cb has been available on my web pages ever since eleven has implemented the modifications I had specified to the standard top-down suexg. Anyone can use it to generate unbiased minimal puzzles.

Back to top profile pm www

Red Ed

Posted: Fri Sep 18, 2009 10:21 pm Post subject: quote

Joined: 06 Jun 2005
Posts: 703

> **denis_berthier wrote:**
> 3) The complete grid generator part of the algorithm has little impact on the generation of puzzles.

Yes, but ...

> **Quote:**
> 3b) Red Ed has been unable to show any significant difference in the distribution of clues when he changed the (supposedly unbiased) generator part of his own algorithm (which is more complex and doesn't output puzzles) for suexg's.

... this is potentially misleading. I was fully able to show what I set out to show: that the difference is real and not an artefact of experimental error (**statistically** significant), but is small (probably not **practically** significant).

Back to top profile pm

denis_berthier

Joined: 19 Jun 2007
Posts: 799
Location: Paris, France

Posted: Fri Sep 18, 2009 10:26 pm Post subject: quote

Red Ed wrote:

> the difference is real and not an artefact of experimental error (**statistically** significant), but is small (probably not **practically** significant).

Oh boy, early morning and already nitpicking ?

Back to top profile pm www

Red Ed

Joined: 06 Jun 2005
Posts: 703

Posted: Fri Sep 18, 2009 10:36 pm Post subject: quote

Denis, try to take these optimisations in the helpful spirit they're intended. You're coming across as somewhat ungracious.
(what's with the "unable to show" / "supposedly unbiased" / "reverting" / "nitpicking" tone ... c'mon, play nicely)

denis_berthier wrote:

> In the real algorithm, we don't have to choose a full path down to 0 clue. As the raw mean number of clues is 25, this spares 25 useless choices in the mean. On this point, your "improvement" would be regressive.

Generating the subgrids is essentially free; it's testing them that's expensive. I'm saying that you may not need to do all the testing you're currently doing.

Let P80, P79, P78, ... be the subgrids of some solution grid on a path from 80 to 0 clues. Suppose that P40 is the first one with multiple solutions. Then, excusing 1-off errors in my description, suexg does this:

```
count solutions for P80 : answer=1
count solutions for P79 : answer=1
count solutions for P78 : answer=1
...
count solutions for P41 : answer=1
count solutions for P40 : answer>1
```

... which is 41 expensive calls to the solver.

A better strategy (one of many possible better strategies) would be to step 2 grids at a time:

```
count solutions for P80 : answer=1
count solutions for P78 : answer=1
count solutions for P76 : answer=1
...
count solutions for P40 : answer>1
count solutions for P41 : answer=1
```

... which is only 22 calls to the solver. Do you understand the principle now?

Quote:

> **Red Ed wrote:**
>
> > Presumably you only object to the optimisation of step 3, not the optimisation of step 4?
>
> I don't see you propose any optimisation in step 4.

I thought you might be grateful for the suggestion in general terms, and interested enough to investigate specifics yourself. But if you want a concrete proposal: how about prioritising removal of clues that have the largest number of rows/columns/boxes/clues-of-same-value in common.

Back to top profile pm

denis_berthier

Posted: Fri Sep 18, 2009 10:54 pm Post subject: quote

Joined: 19 Jun 2007
Posts: 799
Location: Paris, France

> **Red Ed wrote:**
>
> > You're coming across as somewhat ungracious.

The "nitpicking" was your own suggestion, remember ?

> **Red Ed wrote:**
>
> > maybe this is just nit-picking again.

Apologies if I upset you. This was only an early morning friendly joke.

As I said, any real improvement to the implementation of the controlled-bias generator would be welcome. But you know I'm not a C programmer (my "natural" language is Lisp)*. I thought you were one (among other things). If I'm right, why don't you implement and test your ideas?

(* and when I see outer calls to the inside of a function, as in suexg, that makes me very uneasy.)

Back to top profile pm www

Red Ed

Posted: Sat Sep 19, 2009 1:36 am Post subject: quote

Joined: 06 Jun 2005
Posts: 703

Unless I made a mistake, data from a few thousand paths generated by *suexg-cb* suggest that using the optimal* strategy I can get step 3 down from ~37.6 solver calls to ~5.0 solver calls on average:
(* = among those strategies that are allowed only to check for 0/1/2+ solutions, not to probe the structure of the grid e.g. unavoidable sets)

Code:

```
Distribution of number of clues in smallest single-
solution subgrid:
81      0
80      0
79      0
78      3
77      4
76     12
75     38
74     69
73     96
72    147
71    217
```

```
70    338
69    452
68    544
67    696
66    887
65   1109
64   1388
63   1618
62   1933
61   2377
60   2766
59   3241
58   3675
57   4237
56   4818
55   5569
54   6167
53   6772
52   7755
51   8520
50   9365
49  10290
48  11001
47  11861
46  12547
45  13434
44  14230
43  14758
42  15103
41  15132
40  15186
39  14832
38  14132
37  13403
36  11952
35  10397
34   8527
33   6787
32   4913
31   3213
30   1856
29   1021
28    419
27    155
26     27
25     11
24      0
23      0
22      0
21      0
20      0
19      0
18      0
17      0
Average number of calls to the solver  = 37.58
Average number of calls when optimised = 5.04
```

You may wish to check the calculations: everything you need is in the table above.

Back to top profile pm

coloin Posted: Sat Sep 19, 2009 2:24 am Post subject: quote

Thanks **denis**, i will ponder on what you said.

Joined: 05 May 2005
Posts: 1070
Location: Devon UK

Ed - you can remove 27 [3*9] clues with only a little risk of >1 sol.
A minimal puzzle cant have more than 6 occurances of one clue value !

C

Back to top profile pm

Red Ed

Posted: Sat Sep 19, 2009 2:30 am Post subject: quote

Joined: 06 Jun 2005
Posts: 703

Maybe so, coloin, but I doubt it's of much help in step 3 (as I describe it) of the algorithm because the subgrids/clue-deletion path is predetermined: you only have flexibility in the order in which you visit those subgrids.

Back to top profile pm

Red Ed

Posted: Sat Sep 19, 2009 2:37 am Post subject: quote

Joined: 06 Jun 2005
Posts: 703

Now on to optimisation of step 4 (as I describe it). In short: there's no point. The average number of calls to the solver in the minimality test is 1.18; we obviously can't get this lower than 1.0, so the saving in optimising step 4 is not worth the bother.

So it looks like the best bet is to implement my improvement to step 3, whilst doing nothing to step 4, in the anticipation of a factor of six speed-up overall. If the theory holds, that is.

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 ... , 14, **15**, 16 Next

Page 15 of 16

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