Short, difficult model counting benchmarks

Ivor Spence*

Queen's University Belfast

March 11, 2020

1 Background

The purpose of the work described here is to create model counting benchmarks which are as difficult as possible with respect to their size (number of literals). The benchmarks do not implement any particular real work problem, but rather they are intended to encourage the development of efficient solvers.

The benchmarks derive from the author's related work on creating short, difficult satisfiability benchmarks[1], in particular the balanced-sat series[2].

2 Structure of a benchmark

The number of variables (n) is always a multiple of 3. Each clause has 3 literals and the clauses are generated in a series of rounds with each round containing every variable exactly once. In successive rounds the allocation of variables to clauses is re-arranged, avoiding as much as possible the situation where the same two variables share a clause more than once.

The sign of the first occurrence of any variable is chosen at random; the signs of subsequent occurrences alternate. For example, in a benchmark with 9 variables, each round consists of three clauses and the first could be

```
(1 2 3)(-4 7 9)(-5 -6 -8)
```

In the second round the variables are grouped differently and the signs reversed, giving possibly

```
(-1 4 5)(-2 -7 6)(-3 -9 8)
```

To generate the most difficult instances for SAT-solvers, clauses are added until the benchmark is unsatisfiable, which means that the final round of clauses may be incomplete[2]. To generate satisfiable model counting benchmarks, the approach is simply to reduce the number of clauses added.

*i.spence@qub.ac.uk

3 Chosen parameter values

The goal is to find the most difficult benchmark for a given number number of literals. Here, the number of literals (l) is 3 times the number of clauses (m). If $2 * n > 3 * m$ then at least one variable would occur not at all or with only one sign, which is undesirable. In fact it has been found empirically that, for a given value of m, the most difficult benchmarks are generated by letting n be equal to or slightly less than m. Therefore benchmarks have been submitted with the following parameters:

n (variables)	m (clauses)	l(literals)	model count
39	42	126	449260070
42	42	126	3691427579
51	54	162	230641519808
54	54	162	2433432443837
63	66	198	134746112245856
66	66	198	1115259056499565
75	78	234	78358692442136867
78	78	234	671557784116945527
87	90	270	unknown at date of submission
90	90	270	unknown at date of submission
99	102	306	unknown at date of submission
102	102	306	unknown at date of submission

Viewed as SAT benchmarks these are all very easy, but they seem to be challenging for model counting - not just because of the number of models.

4 Filenames

The filename for each benchmark reflects the number of variables, the number of clauses and the chosen random seed. For the benchmarks submitted to MC2020 the random seed is always 1. For example, the benchmark `bsat-39-42-1.cnf` has 39 variables, 42 clauses, and was generated using a random seed of 1.

References

[1] Ivor Spence. Weakening cardinality constraints creates harder satisfiability benchmarks. *J. Exp. Algorithmics*, 20, May 2015.

[2] Ivor Spence. Balanced random sat benchmarks. In Tomáš Balyo, Marijn J.H. Heule, and Matti Järvisalo, editors, *Proceedings of SAT Competition 2017: Solver and Benchmark Descriptions*, volume B-2017-1 of *Series of Publications B*, Finland, 2020. Department of Computer Science, University of Helsinki.