



Vampire Getting Noisy: Will Random Bits Help Conquer Chaos? (System Description)

Martin Suda^(✉) 

Czech Technical University in Prague, Prague, Czech Republic
`martin.suda@cvut.cz`

Abstract. Treating a saturation-based automatic theorem prover (ATP) as a Las Vegas randomized algorithm is a way to illuminate the chaotic nature of proof search and make it amenable to study by probabilistic tools. On a series of experiments with the ATP Vampire, the paper showcases some implications of this perspective for prover evaluation.

Keywords: Saturation-based proving · Evaluation · Randomization

1 Introduction

Saturation-based proof search is known to be fragile. Even seemingly insignificant changes in the search procedure, such as shuffling the order in which input formulas are presented to the prover, can have a huge impact on the prover's running time and thus on the ability to find a proof within a given time limit.

This *chaotic* aspect of the prover behaviour is relatively poorly understood, yet has obvious consequences for evaluation. A typical experimental evaluation of a new technique T compares the number of problems solved by a baseline run with a run enhanced by T (over an established benchmark and with a fixed timeout). While a higher number of problems solved by the run enhanced by T indicates a benefit of the new technique, it is hard to claim that a certain problem P is getting solved *thanks* to T . It might be that T just helps the prover get lucky on P by a complicated chain of cause and effect not related to the technique T —and the original idea behind it—in any reasonable sense.

We propose to expose and counter the effect of chaotic behaviours by deliberately *injecting randomness* into the prover and observing the results of many independently seeded runs. Although computationally more costly than standard evaluation, such an approach promises to bring new insights. We gain the ability to apply the tools of probability theory and statistics to analyze the results, assign confidences, and single out those problems that *robustly* benefit

This work was supported by the Czech Science Foundation project 20-06390Y and the project RICAIP no. 857306 under the EU-H2020 programme.

© The Author(s) 2022

J. Blanchette et al. (Eds.): IJCAR 2022, LNAI 13385, pp. 659–667, 2022.

https://doi.org/10.1007/978-3-031-10769-6_38

from the evaluated technique. At the same time, by observing the changes in the corresponding runtime distributions we can even meaningfully establish the effect of the new technique on a single problem in isolation, something that is normally inconclusive due to the threat of chaotic fluctuations.

In this paper, we report on several experiments with a randomized version of the ATP Vampire [9]. After explaining the method in more detail (Sect. 2), we first demonstrate the extent in which the success of a typical Vampire proof search strategy can be ascribed to chance (Sect. 3). Next, we use the collected data to highlight the specifics of comparing two strategies probabilistically (Sect. 4). Finally, we focus on a single problem to see a chaotic behaviour smoothed into a distribution with a high variance (Sect. 5). The paper ends with an overview of related work (Sect. 6) and a discussion (Sect. 7).

2 Randomizing Out Chaos

Any developer of a saturation-based prover will confirm that the behaviour of a specific proving strategy on a specific problem is extremely hard to predict, that a typical experimental evaluation of a new technique (such as the one described earlier) invariably leads to both gains and losses in terms of the solved problems, and that a closer look at any of the “lost” problems often reveals just a complicated chain of cause and effect that steers the prover away from the original path (rather than a simple opportunity to improve the technique further).

These observations bring indirect evidence that the prover’s behaviour is chaotic: A specific prover run can be likened to a single bead falling down through the pegs of the famous Galton board¹. The bead follows a deterministic trajectory, but only because the code fixes every single detail of the execution, including many which the programmer did not care about and which were left as they are merely out of coincidence. We put forward here that any such fixed detail (which does not contribute to an officially implemented heuristic) represents a candidate location for randomization, since a different programmer could have fixed the detail differently and we would still call the code essentially the same.

Implementation: We implemented randomization on top of Vampire version 4.6.1; the code is available as a separate git branch². We divided the randomization opportunities into three groups (governed by three new Vampire options).

Shuffling the input (`-si on`) randomly reorders the input formulas and, recursively, sub-formulas under commutative logical operations. This is done several times throughout the preprocessing pipeline, at the end of which a finished clause normal form is produced. Randomizing traversals (`-rtra on`) happens during saturation and consists of several randomized reorderings including: reordering literals in a newly generated clause and in each given clause before activation, and shuffling the order in which generated clauses are put into the

¹ https://en.wikipedia.org/wiki/Galton_board.

² <https://github.com/vprover/vampire/tree/randire>.




Fig. 1. Blue: first-order TPTP problems ordered by the decreasing probability of being solved by the `dis10` strategy within 50 billion instruction limit. Red: a cactus plot for the same strategy, showing the dependence between a given instruction budget (y -axis) and the number of problems on average solved within that budget (x -axis). (Color figure online)

passive set. It also (partially) randomizes term ids, which are used as tiebreakers in various term indexing operations and determine the default orientation of equational literals in the term sharing structure. Finally, “randomized age-weight ratio” (`-rawr on`) swaps the default, deterministic mechanism for choosing the next queue to select the given clause from [13] for a randomized one (which only respects the age-weight ratio probabilistically).

All the three options were active by default during our experiments.

3 Experiment 1: A Single-Strategy View

First, we set out to establish to what degree the performance of a Vampire strategy can be affected by randomization. We chose the default strategy of the prover except for the saturation algorithm, which we set to Discount, and the age-weight ratio, set to 1:10 (calling the strategy `dis10`). We ran our experiment on the first-order problems from the TPTP library [15] version 7.5.0³.

To collect our data, we repeatedly (with different seeds) ran the prover on the problems, performing full randomization. We measured the executed instructions⁴ needed to successfully solve a problem and used a limit of 50 billion instructions (which roughly corresponds to 15 s of running time on our machine⁵) after which a run was declared unsuccessful. We ran the prover 10 times on each problem and additionally as many times as required to observe the instruction count average (over both successful and unsuccessful runs) stabilize within 1% from any of its 10 previously recorded values⁶.

A summary view of the experiment is given by Fig. 1. The most important to notice is the shaded region there, which spans 965 problems that were solved by

³ Materials accompanying the experiments can be found at <https://bit.ly/3JDCwea>.

⁴ As measured via the `perf_event_open` Linux performance monitoring feature.

⁵ A server with Intel(R) Xeon(R) Gold 6140 CPUs @ 2.3 GHz and 500 GB RAM.

⁶ Utilizing all the 72 cores of our machine, such data collection took roughly 12 h.




Fig. 2. The effect of turning AVATAR off in the `dis10` strategy (cf. Figure 1).

`dis10` at least once but not by every run. In other words, these problems have probability p of being solved between $0 < p < 1$. This is a relatively large number and can be compared to the 8720 “easy” problems solved by every run. The collected data implies that 9319.1 problems are being solved on average (marked by the left-most dashed line in Fig. 1) with a standard deviation $\sigma = 11.7$. The latter should be an interesting indicator for prover developers: beating a baseline by only 12 TPTP problems can easily be ascribed just to chance.

Figure 1 also contains the obligatory “cactus plot” (explained in the caption), which—thanks to the collected data—can be constructed with the “on average” qualifier. By definition, the plot reaches the left-most dashed line for the full instruction budget of 50 billion. The subsequent dashed lines mark the number of problems we would on average expect to solve by running the prover (independently) on each problem twice, three, four and five times. This is an information relevant for strategy scheduling: e.g., one can expect to solve whole additional 137 problems by running randomized `dis10` for a second time.

Not every strategy exhibits the same degree of variability under randomization. Observe Fig. 2 with a plot analogous to Fig. 1, but for `dis10` in which the AVATAR [16] has been turned off. The shaded area there is now much smaller (and only spans 448 problems). The powerful AVATAR architecture is getting convicted of making proof search more fragile and the prover less robust⁷.

Remark. Randomization incurs a small but measurable computational overhead. On a single run of `dis10` over the first-order TPTP (filtering out cases that took less than 1s to finish, to prevent distortion by rounding errors) the observed median relative time spent randomizing on a single problem was 0.47%, the average 0.59%, and the worse⁸ 13.86%. Without randomization, the `dis10` strategy solved 9335 TPTP problems under the 50 billion instruction limit, i.e., 16 problems more than the average reported above. Such is the price we pay for turning our prover into a Las Vegas randomized algorithm.

⁷ Another example of a strong but fragile heuristic is the lookahead literal selection [5], which selects literals in a clause based on the current content of the active set: `dis10` enhanced with lookahead solves 9512.4 (± 13.8) TPTP problems on average, 8672 problems with $p = 1$ and additional 1382 (!) problems with $0 < p < 1$.

⁸ On the hard-to-parse, trivial-to-solve HWV094-1 with 361 199 clauses.




Fig. 3. Scatter plots comparing probabilities of solving a TPTP problem by the baseline `dis10` strategy and 1) `dis10` with AVATAR turned off (left), and 2) `dis10` with blocked clause elimination turned on (right). On problems marked red the respective technique could not be applied (no splittable clauses derived / no blocked clauses eliminated).

4 Experiment 2: Comparing Two Strategies

Once randomized performance profiles of multiple strategies are collected, it is interesting to look at two at a time. Figure 3 shows two very different scatter plots, each comparing our baseline `dis10` to its modified version in terms of the probabilities of solving individual problems.

On the left we see the effect of turning AVATAR off. The technique affects the proving landscape quite a lot and most problems have their mark along the edges of the plot where at least one of the two probabilities has the extreme value of either 0 or 1. What the plot does not show well, is how many marks end up at the extreme corners. These are: 7896 problems easy for both, 661 easy for AVATAR and hard without, 135 hard for AVATAR and easy without.

Such “purified”, one-sided gains and losses constitute a new interesting indicator of the impact of a given technique. They should be the first to look at, e.g., during debugging, as they represent the most extreme but robust examples of how the new technique changes the capabilities of the prover.

The right plot is an analogous view, but now at the effect of turning on *blocked clause elimination* (BCE). This is a preprocessing technique coming from the context of propositional satisfiability [7] extended to first-order logic [8]. We see that here most of the visible problems show up as marks along the plot’s main diagonal, suggesting a (mostly) negligible effect of the technique. The extreme corners hide: 8648 problems easy for both, 17 easy with BCE (11 satisfiable and 6 unsatisfiable), and 2 easy without BCE (1 satisfiable and 1 unsatisfiable).




Fig. 4. 2D-histograms for the relative frequencies (color-scale) of how often, given a specific awr (x -axis), solving PR0017+2 required the shown number of instructions (y -axis). The curves in pink highlight the mean y -value for every x . The performance of `dis10` (left) and the same strategy enhanced by a goal-directed heuristic (right). (Color figure online)

5 Experiment 3: Looking at One Problem at a Time

In their paper on age/weight shapes [13, Fig. 2], Rawson and Reger plot the number of given-clause loops required by Vampire to solve the TPTP problem PR0017+2 as a function of age/weight ratio (awr), a ratio specifying how often the prover selects the next clause to activate from its age-ordered and weight-ordered queues, respectively. The curve they obtain is quite “jiggly”, indicating a fragile (discontinuous) dependence. Randomization allows us to smoothen the picture and reveal new, until now hidden, (probabilistic) patterns.

The 2D-histogram in Fig. 4 (left) was obtained from 100 independently seeded runs for each of 1200 distinct values of awr from between $1:1024 = 2^{-10}$ and $4:1 = 2^2$. We can confirm Rawson and Reger’s observation of the best awr for PR0017+2 lying at around 1:2. However, we can now also attempt to explain the “jiggly-ness” of their curve: With a fragile proof search, even a slight change in awr effectively corresponds to an independent sample from the prover’s execution resource⁹ distribution, which—although changing continuously with awr —is of a high variance for our problem (note the log-scale of the y -axis)¹⁰.

The distribution has another interesting property: At least for certain values of awr it is distinctly multi-modal. As if the prover can either find a proof quickly (after a lucky event?) or only after much harder effort later and almost nothing in between. Shedding more light on this phenomenon is left for further research.

It is also very interesting to observe the change of such a 2D-histogram when we modify the proof search strategy. Figure 4 (right) shows the effect of turning on SInE-level split queues [3], a goal directed clause selection heuristic

⁹ Rawson and Reger [13] counted given-clause loops, we measure instructions.

¹⁰ Even with 100 samples for each value of awr , the mean instruction count (rendered in pink in Fig. 4) looks jiggly towards the weight-heavy end of the plot.

(Vampire option `-slsq on`). We can see that the mean instruction count gets worse (for every tried *awr* value) and also the variance of the distribution distinctly increases. A curious effect of this is that we observe the shortest successful runs with `-slsq on`, while we still could not recommend (in the case of PR0017+2) this heuristic to the user. The probabilistic view makes us realize that there are competing criteria of prover performance for which one might want to optimize.

6 Related Work

The idea of randomizing a theorem prover is not new. Ertel [2] studied the speedup potential of running independently seeded instances of the connection prover SETHEO [10]. The dashed lines in our Figs. 1 and 2 capture an analogous notion in terms of “additional problems covered” for levels of parallelism 1–5. `randoCoP` [12] is a randomized version of another connection prover, `leanCoP` 2.0 [11]: especially in its incomplete setup, several restarts with different seeds helped `randoCoP` improve over `leanCoP` in terms of the number of solved problems.

Gomes et al. [4] notice that randomized complete backtracking algorithms for propositional satisfiability (SAT) lead to heavy-tailed runtime distributions on satisfiable instances. While we have not yet analyzed the runtime distributions coming from saturation-based first-order proof search in detail, we definitely observed high variance also for unsatisfiable problems. Also in the domain of SAT, Brglez et al. [1] proposed input shuffling as a way of turning solver’s runtime into a random variable and studied the corresponding distributions.

An interesting view on the trade-offs between expected performance of a randomized solver and the risk associated with waiting for an especially long run to finish is given by Huberman et al. [6]. This is related to the last remark of the previous section.

Finally, in the satisfiability modulo theories (SMT) community, input shuffling, or scrambling, has been discussed as an obfuscation measure in competitions [17], where it should prevent the solvers to simply look up a precomputed answer upon recognising a previously seen problem. Notable is also the use of randomization in solver debugging via fuzz testing [14, 18].

7 Discussion

As we have seen, the behaviour of a state-of-the-art saturation-based theorem prover is to a considerable degree chaotic and on many problems a mere perturbation of seemingly unimportant execution details decides about the success or the failure of the corresponding run. While this may be seen as a sign of our as-of-yet imperfect grasp of the technology, the author believes that an equally plausible view is that some form of chaos is inherent and originates from the complexity of the theorem proving task itself. (A higher-order logic proof search is expected to exhibit an even higher degree of fragility.)

This paper has proposed randomization as a key ingredient to a prover evaluation method that takes the chaotic nature of proof search into account. The

extra cost required by the repeated runs, in itself not unreasonable to pay on contemporary parallel hardware, seems more than compensated by the new insights coming from the probabilistic picture that emerges. Moreover, other uses of randomization are easy to imagine, such as data augmentation for machine learning approaches or the construction of more robust strategy schedules. It feels that we only scratched the surface of the opened-up possibilities. More research will be needed to fully harness the potential of this perspective.

References

1. Brglez, F., Li, X.Y., Stallmann, M.F.M.: On SAT instance classes and a method for reliable performance experiments with SAT solvers. *Ann. Math. Artif. Intell.* **43**(1), 1–34 (2005). <https://doi.org/10.1007/s10472-005-0417-5>
2. Ertel, W.: OR-parallel theorem proving with random competition. In: Voronkov, A. (ed.) *LPAR 1992*. LNCS, vol. 624, pp. 226–237. Springer, Heidelberg (1992). <https://doi.org/10.1007/BFb0013064>
3. Gleiss, B., Suda, M.: Layered clause selection for saturation-based theorem proving. In: Fontaine, P., Korovin, K., Kotsireas, I.S., Rümmer, P., Tourret, S. (eds.) *PAAR 7*, Paris, France, June–July 2020. *CEUR Workshop Proceedings*, vol. 2752, pp. 34–52. CEUR-WS.org (2020). <http://ceur-ws.org/Vol-2752/paper3.pdf>
4. Gomes, C.P., Selman, B., Crato, N., Kautz, H.A.: Heavy-tailed phenomena in satisfiability and constraint satisfaction problems. *J. Autom. Reason.* **24**(1/2), 67–100 (2000). <https://doi.org/10.1023/A:1006314320276>
5. Hoder, K., Regeer, G., Suda, M., Voronkov, A.: Selecting the selection. In: Olivetti, N., Tiwari, A. (eds.) *IJCAR 2016*. LNCS (LNAI), vol. 9706, pp. 313–329. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_22
6. Huberman, B., Lukose, R., Hogg, T.: An economics approach to hard computational problems. *Science* **275**, 51–4 (1997)
7. Järvisalo, M., Biere, A., Heule, M.: Blocked clause elimination. In: Esparza, J., Majumdar, R. (eds.) *TACAS 2010*. LNCS, vol. 6015, pp. 129–144. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12002-2_10
8. Kiesl, B., Suda, M., Seidl, M., Tompits, H., Biere, A.: Blocked clauses in first-order logic. In: Eiter, T., Sands, D. (eds.) *LPAR-21*, Maun, Botswana, 7–12 May 2017. *EPiC Series in Computing*, vol. 46, pp. 31–48. EasyChair (2017)
9. Kovács, L., Voronkov, A.: First-order theorem proving and VAMPIRE. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 1–35. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_1
10. Letz, R., Schumann, J., Bayerl, S., Bibel, W.: SETHEO: a high-performance theorem prover. *J. Autom. Reason.* **8**(2), 183–212 (1992). <https://doi.org/10.1007/BF00244282>
11. Otten, J.: leanCoP 2.0 and ileanCoP 1.2: high performance lean theorem proving in classical and intuitionistic logic (system descriptions). In: Armando, A., Baumgartner, P., Dowek, G. (eds.) *IJCAR 2008*. LNCS (LNAI), vol. 5195, pp. 283–291. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-71070-7_23
12. Raths, T., Otten, J.: randoCoP: randomizing the proof search order in the connection calculus. In: Konev, B., Schmidt, R.A., Schulz, S. (eds.) *PAAR 1*, Sydney, Australia, 10–11 August 2008. *CEUR Workshop Proceedings*, vol. 373. CEUR-WS.org (2008). <http://ceur-ws.org/Vol-373/paper-08.pdf>

13. Rawson, M., Regeer, G.: Old or heavy? decaying gracefully with age/weight shapes. In: Fontaine, P. (ed.) CADE 2019. LNCS (LNAI), vol. 11716, pp. 462–476. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29436-6_27
14. Scott, J., Sudula, T., Rehman, H., Mora, F., Ganesh, V.: BanditFuzz: fuzzing SMT solvers with multi-agent reinforcement learning. In: Huisman, M., Păsăreanu, C., Zhan, N. (eds.) FM 2021. LNCS, vol. 13047, pp. 103–121. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-90870-6_6
15. Sutcliffe, G.: The TPTP problem library and associated infrastructure. From CNF to TH0, TPTP v.6.4.0. *J. Autom. Reason.* **59**(4), 483–502 (2017). <https://doi.org/10.1007/s10817-017-9407-7>
16. Voronkov, A.: AVATAR: the architecture for first-order theorem provers. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 696–710. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_46
17. Weber, T.: Scrambling and descrambling SMT-LIB benchmarks. In: King, T., Piskac, R. (eds.) SMT@IJCAR 2016, Coimbra, Portugal, 1–2 July 2016. CEUR Workshop Proceedings, vol. 1617, pp. 31–40. CEUR-WS.org (2016). <http://ceur-ws.org/Vol-1617/paper3.pdf>
18. Winterer, D., Zhang, C., Su, Z.: Validating SMT solvers via semantic fusion. In: Donaldson, A.F., Torlak, E. (eds.) PLDI 2020, London, UK, 15–20 June 2020, pp. 718–730. ACM (2020)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

