Heavy-tailed Runtime Distributions: Heuristics, Models and Optimal Refutations

Tudor Hulubei and Barry O'Sullivan
Cork Constraint Computation Centre, Ireland
tudor@hulubei.net, b.osullivan@4c.ucc.ie

1 Introduction

We perform an in-depth empirical study of runtime distributions associated with a continuum of problem formulations for QWH-10 with 90% holes defined between the points where a formulation is entirely specified in terms of binary inequality constraints to hybrid models specified using an increasing number of ALLDIFFERENT global constraints [4]. For each model we study a variety of variable and value ordering heuristics. We compare their runtime distributions against runtime distributions [1] where any mistakes made in search are refuted optimally [2].

2 Experiments

Our experiments were performed on satisfiable QWH-10 problem instances with 90% random balanced holes, and included 4 variable ordering heuristics: min-domain, min-dom/ddeg, brelaz and min-dom/wdeg, and 3 value ordering heuristics: random, min-conflicts and max-conflicts, totalling 12 algorithms (we broke ties randomly).

We refer to a *mistake point* [2] as an assignment that cannot lead to a solution even though one existed before that assignment was made. An *actual refutation* is the search tree corresponding to a mistake, as obtained by some algorithm, with the *optimal refutation* for that mistake corresponding to a search tree of minimum size. Finally, the *quasi-optimal refutation* is the smallest refutation whose height does not exceed that of the actual refutation.

3 Summary of Results

1. For the problems considered, *variations in the heuristics used have a far more significant effect on hybrid models* (i.e. models using both binary and global constraints) than they do on purely binary models (Figure 1).
2. While algorithms tend to perform better on hybrid models, a straight line can still be observed in a log-log plot of their runtime distributions, even when mistakes are refuted optimally (Figure 2(b)). In other words, *runtime distributions of hybrid models can remain inherently heavy-tailed* [3].
3. *Models using global constraints are not always better than purely binary models.* We encountered configurations where increasing the number of global constraints used to enforce distinct values on rows and columns (and removing the corresponding sets of binary constraints) does not lead to a monotonic decrease in search effort (Figure 2(a)). *The discrepancy all but disappeared when we looked at the corresponding (quasi-)optimal refutations* for the exact same configuration.
4. With the exception of a few unusual cases, using global constraints did improve search performance and, when that occured, the *refutations encountered for hybrid models were much closer to their corresponding optimal* than for the binary model (Figure 3).
5. For the problems considered, the variables used in refuting a mistake usually represent only a small fraction of the variables still uninstantiated at the time the mistake was made, *yet this small subset of variables can most of the time be re-ordered to refute the mistake optimally* (Figure 4).

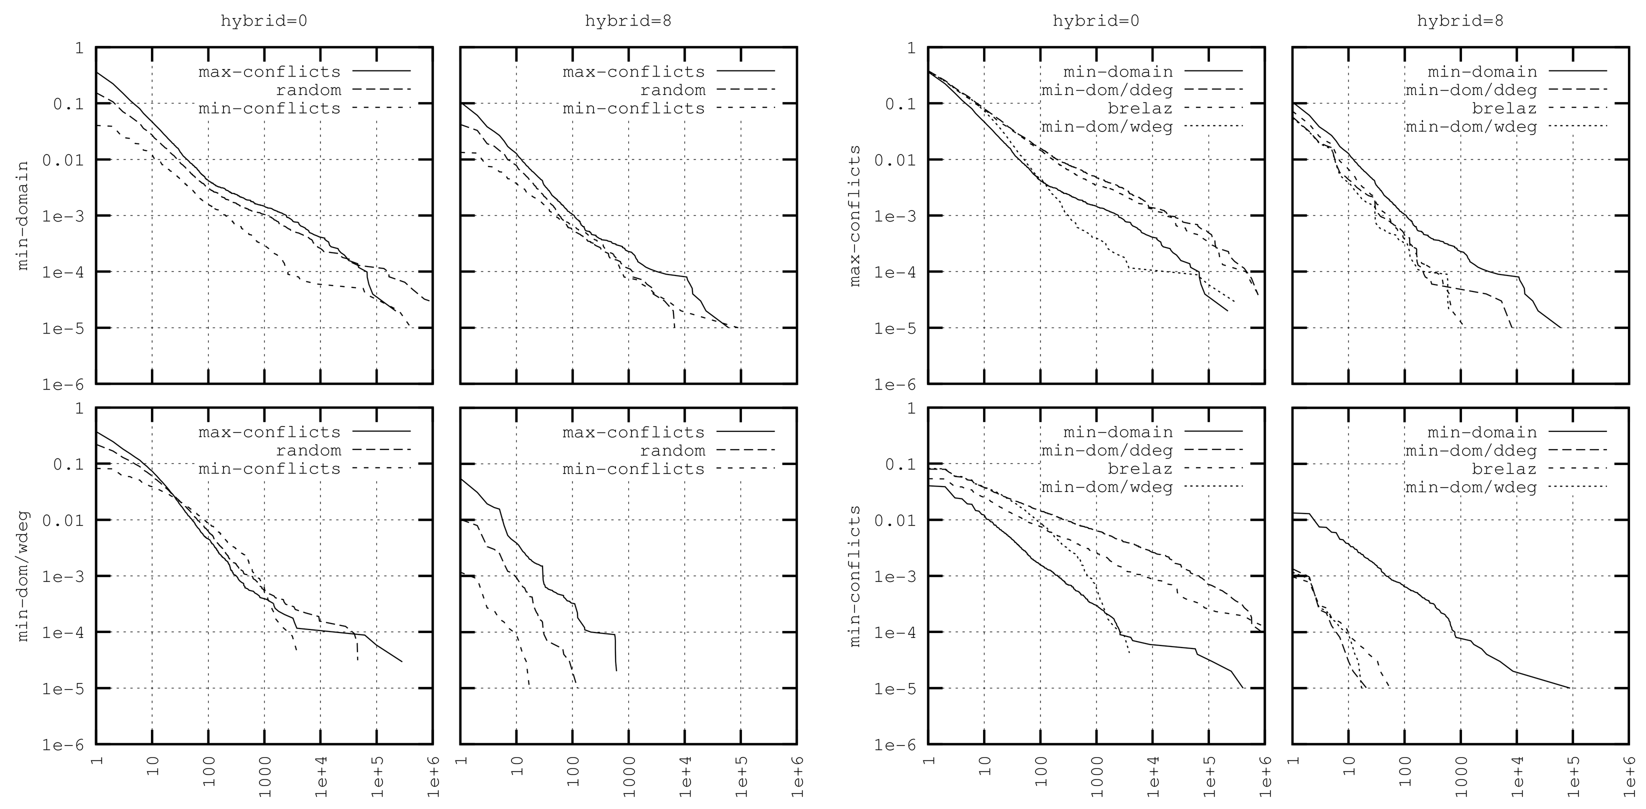

(a) Binary and hybrid=8 models using a poor (min-domain) and a good (min-dom/wdeg) variable ordering (various value orderings).

(b) Binary and hybrid=8 models using a poor (max-conflicts) and a good (min-conflicts) value ordering (various variable orderings).

Figure 1: Complement of the CDF (y-axis) of the cumulative *actual* effort (x-axis).

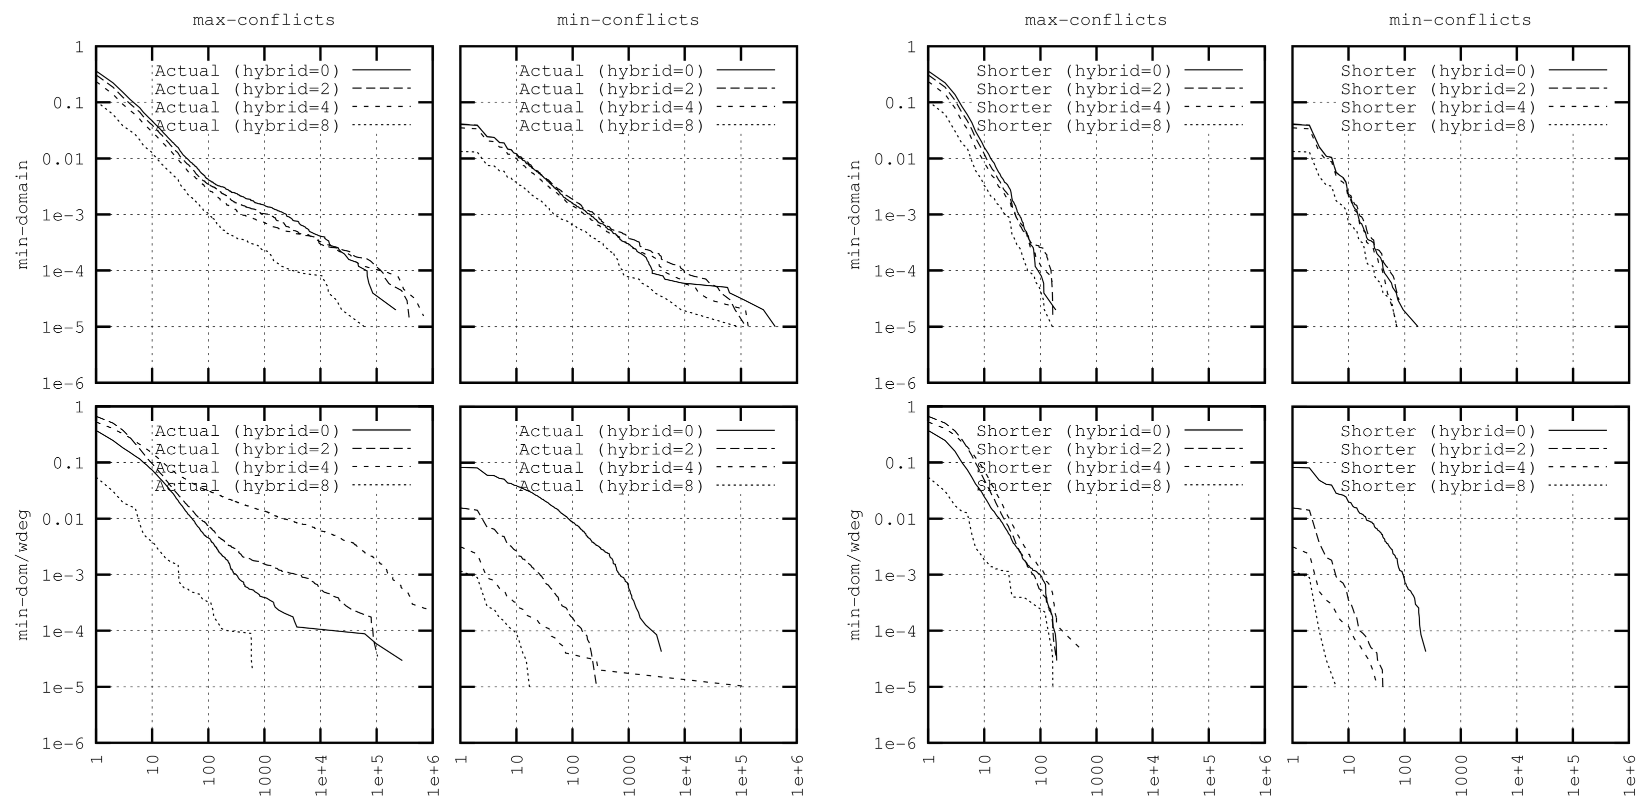

(a) *Actual* refutation effort.

(b) *Shorter* refutation effort.

Figure 2: Complement of the CDF (y-axis) of the actual (left) and shorter (right) effort (x-axis). We vary the value ordering across columns and variable ordering across rows.

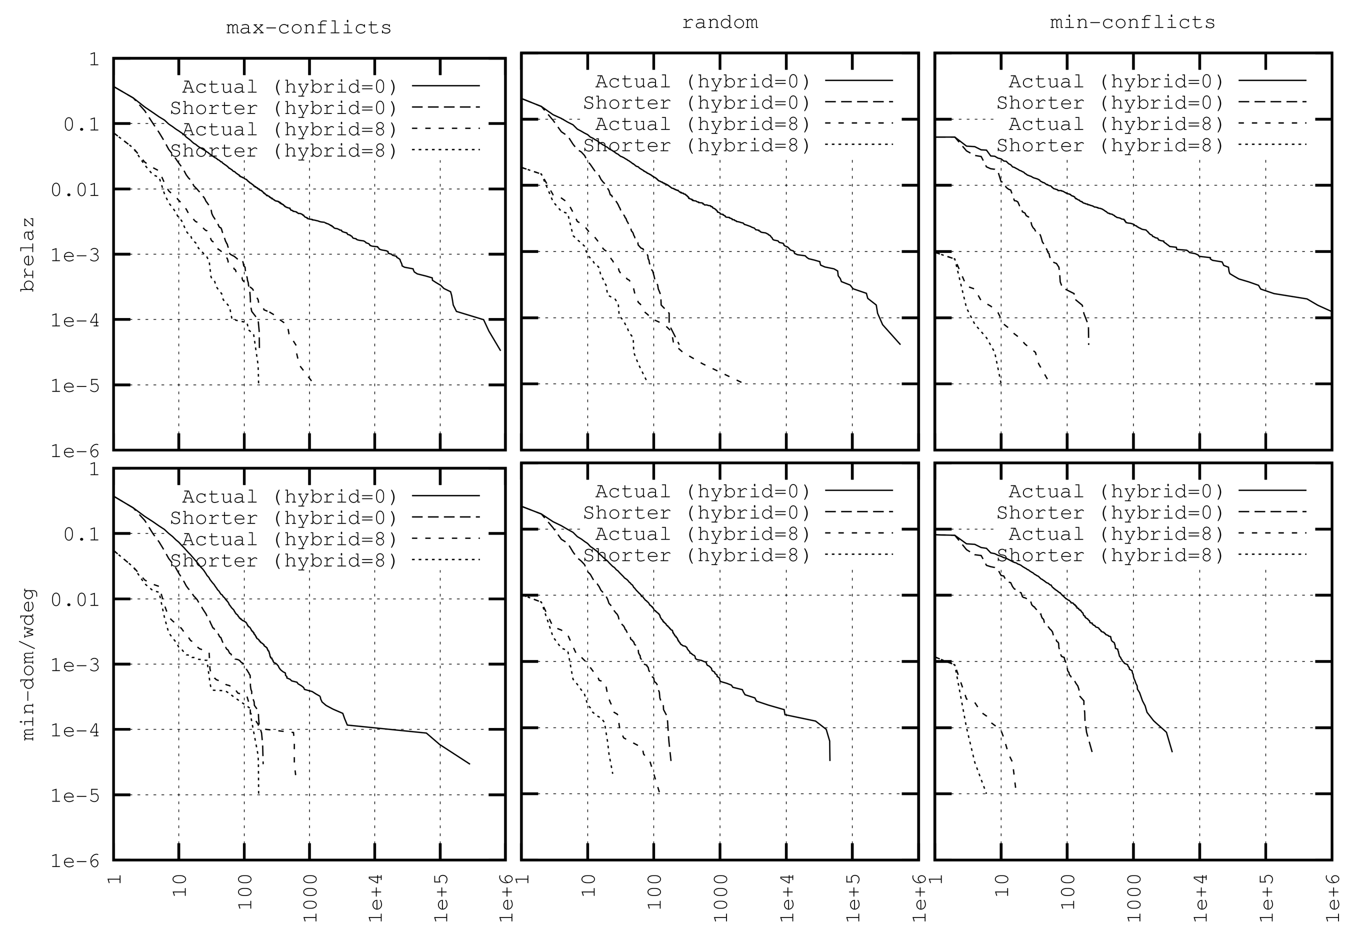

Figure 3: Comparison of the complement of the CDF (y-axis) for the cumulative *actual* and *shorter* effort (x-axis) for the binary and hybrid=8 models.

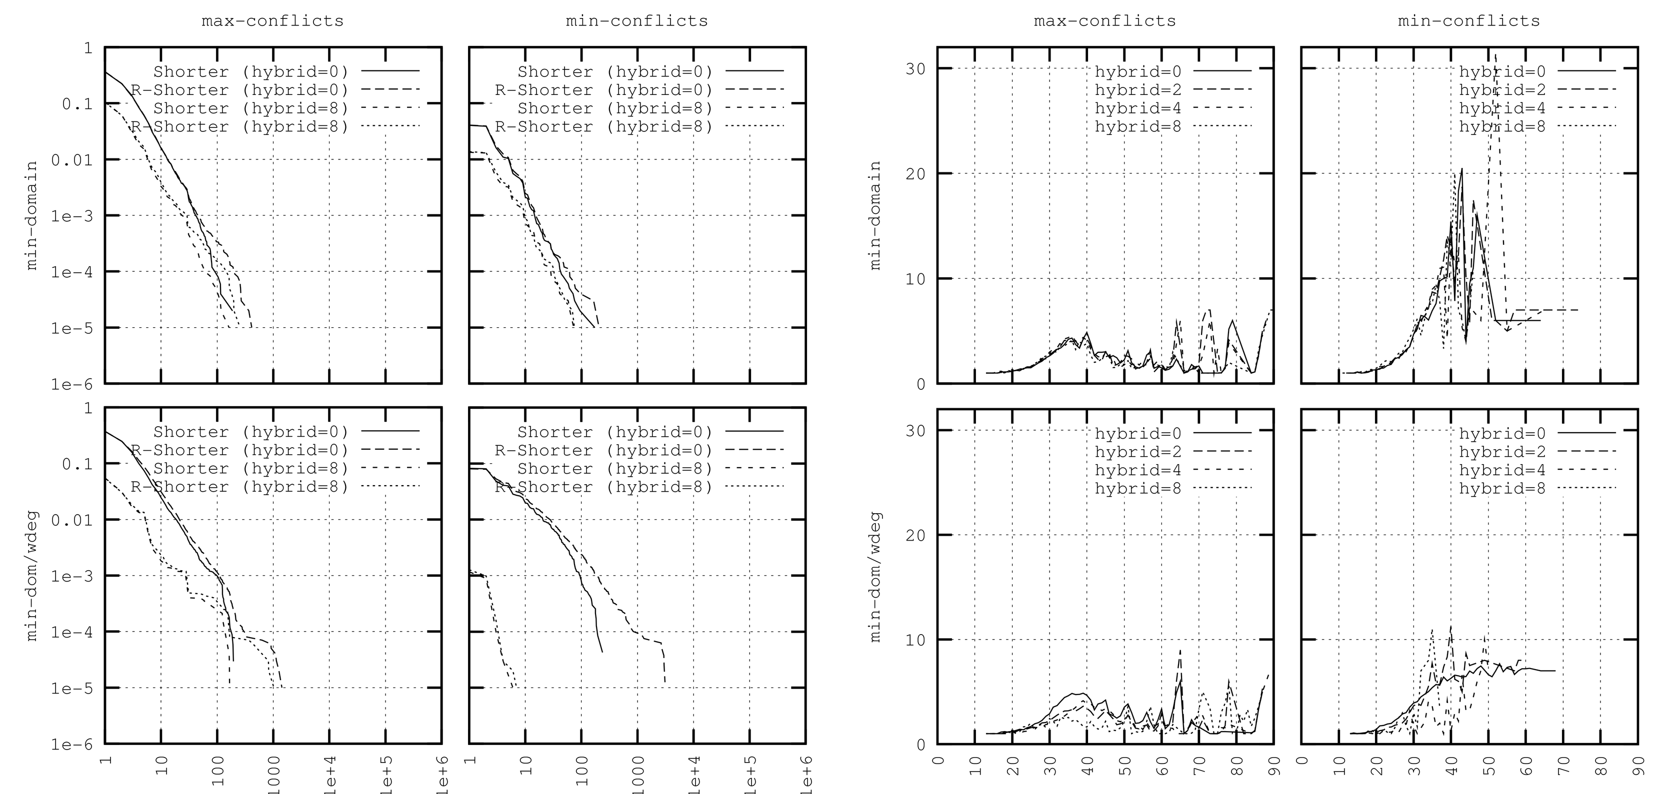

(a) Comparison of the complement of the CDF (y-axis) for the cumulative *shorter* and *restricted shorter (R-Shorter in the plots)* (x-axis) refutations for the binary and hybrid=8 models.

(b) Average number of variables involved in the *actual* refutations (y-axis) as a function of the number of variables still uninstantiated at the time a mistake was made (x-axis).

Figure 4: Statistics on restricted refutations.

4 Conclusions

We have shown empirically that for QWH-10, variations in heuristics have a greater effect on formulations involving a mix of binary and global constraints than on purely binary models. Models using global constraints are not always better than purely binary models. We have also shown that the small subset of variables used by a heuristic to refute a mistake can be *re-ordered* to obtain an almost optimal refutation. This raises the question of why heuristics select the right variables, but fail to find better refutations.

Acknowledgments

This work was supported by Science Foundation Ireland (Grant 00/PI.1/C075). We thank the Boole Centre For Research in Informatics for providing access to their Beowulf cluster.

References

[1] C.P. Gomes, B. Selman, N. Crato, and H. Kautz. Heavy-tailed phenomena in satisfiability and constraint satisfaction problems. *Automated Reasoning*, 24(1/2):67–100, 2000.

[2] T. Hulubei and B. O'Sullivan. Optimal refutations for constraint satisfaction problems. In *Proceedings of IJCAI-2005*, pages 163–168, 2005.

[3] T. Hulubei and B. O'Sullivan. The impact of search heuristics on heavy-tailed behaviour. *Constraints*, 11(2–3):157–176, 2006.

[4] J. C. Regin. A filtering algorithm for constraints of difference in CSPs. In *Proceedings of AAAI*, pages 362–367, 1994.