On the efficiency of impact based heuristics

Marco Correia and Pedro Barahona

Centro de Inteligência Artificial, Departamento de Informática,
Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
{mvc,pb}@di.fct.unl.pt

Abstract. Recently proposed impact based heuristics have been shown to outperform other instances of the first-fail policy such as the common dom and dom/deg heuristics. This paper compares the behaviour of a constraint and a variable centered impact based heuristic and relates it to the amount of constraint propagation inherent to the model of the problem. Additionally, it presents results which suggest that a look-ahead impact heuristic we recently proposed might be the best choice for problems with low locality and where constraint propagation plays an important role.

1 Introduction

Despite the success of Constraint Programming in addressing combinatorial NP problems non trivial instances still require appropriate search strategies to find solutions efficiently. The most general such strategy is possibly the first-fail principle: solve first the most difficult sub-problems. This principle is usually implemented in the variable selection heuristics, by starting enumeration on the variables with domains of least cardinality. This heuristics does not take into account the "structure" of the problem, which is more likely to be accounted for by selecting the variable that participates in more constraints (hence, harder to enumerate). The combination of both ideas results in the popular "dom/deg" heuristic, which performs reasonably well in a number of problems.

Still, other heuristics have been recently proposed that implement the first fail principle more efficiently, by measuring during search the impact of the decisions made, either focusing on the constraints violated, as the wdeg heuristics [2] or the impact on the search space as the variable centered heuristics presented in [12]. However we have found that the singleton arc consistency heuristic we proposed [4] clearly outperforms the dom/wdeg heuristics in a class of CSP problems with some identified features.

In this paper, we firstly describe in more detail the above heuristics. Then we show their results in a number of benchmark problems, for which we study their features, namely the propagation that is achieved, and the correlation that exists between these features and the efficiency shown by the heuristics under study. We conclude with some general comments on this preliminary study and some suggestions for further work.

2 Preliminaries

A constraint network consists of a set of variables $\mathcal{X}$, a set of domains $\mathcal{D}$, and a set of constraints $\mathcal{C}$. Every variable $x \in \mathcal{X}$ has an associated domain $D(x)$ denoting its possible values. Every k-ary constraint $c \in \mathcal{C}$ is defined over a set of k variables $(x_1, \ldots, x_k)$ by the subset of the Cartesian product $D(x_1) \times \ldots \times D(x_k)$ which are consistent values. The constraint satisfaction problem (CSP) consists in finding an assignment of values to variables such that all constraints are satisfied.

A CSP is arc-consistent iff it has non-empty domains and every consistent instantiation of a variable can be extended to a consistent instantiation involving an additional variable. A problem is generalized arc-consistent (GAC) iff for every value in each variable of a constraint there exist compatible values for all the other variables in the constraint.

A CSP P is singleton θ-consistent (SC), iff it has non-empty domains and for any value $a \in dom\,(x)$ of every variable $x \in \mathcal{X}$, the resulting subproblem $P|_{x=a}$ can be made θ-consistent. SC time complexity is in $\mathrm{O}(n^2 d^2 \Theta)$, Θ being the time complexity of the algorithm that achieves θ-consistency on the constraint network. Restricted singleton consistency (RSC) [11] considers each variable only once and has runtime complexity $\mathrm{O}(nd\Theta)$.

3 Impact based heuristics

Impact based heuristics use information collected dynamically during search to score each variable according to its first-failness. We now revise three impact based strategies which use distinct sources of information.

The method presented in [2] associates a counter with each constraint expressing the number of times it was violated since the beginning of search. This information is projected to each variable by summing the counters for all constraints where the variable participates (referred to as the *weighted degree* of the variable). Heuristics then select the variable with largest degree (wdeg), or with smallest ratio between the size of the domain and the degree (dom/wdeg).

The impact based heuristic introduced in [12] measures the size of the search space, given by a function $\sigma\,(P)$, before and after the enumeration of a variable. The impact of a variable x, given by $(\sigma\,(P) - \sigma\,(P|_{x=a}))\,/\sigma\,(P)$, is averaged over all its previous enumerations and the highest impact so far indicates the next variable to instantiate.

In [4] we extended this idea by explicitly computing the actual impact of every variable before each enumeration, which is available if simultaneously maintaining (restricted) singleton consistency. In this case we give preference for the variable with a smaller sum of search space size for every possible instantiation, formally

$$\mathrm{var}_{\mathrm{LA}}\,(P) = \arg \min_{x \in \mathcal{X}(P)} \left(\sum_{a \in D(x)} \sigma\,(P|_{x=a}) \right)$$

which also corresponds to the minimize promise variable heuristic defined in [7]. As said, this heuristic requires lookahead computations which may be inefficient, however it is expected that the reduction in dead end failures compensates such inefficiency.

4 Experiments and discussion

In this section we compare the heuristics described previously on a set of three well known benchmarks. The dom/dweg heuristic was used while maintaining arc-consistency (dom/wdeg). We also tested the dom/wdeg heuristic while maintaining restricted singleton arc consistency, and the dom and wdeg heuristics independently but their performance were consistently worse (results are omitted for space reasons). The var_{LA} heuristic was used with restricted singleton arc consistency (la) further restricted to the subset of variables with domain size $d = 2$ as described in [4], to avoid expensive lookahead computation with variables unlikely to be selected.

Firstly, 200 random CSP instances were generated using model C [9] (generalized to 5-ary CSPs) with 25 variables and domains of size 4. Instances were created by varying the looseness of the constraints in $[0.1 \ldots 0.8]$ and setting the density such that the constrainedness of the instances is $\kappa = 0.95$ [8] (phase transition occurs typically at $\kappa \approx 1$). For more details see [4]. Solutions were stored as positive table constraints and GAC-Schema [1] was used for filtering.

Graph coloring assigns n colors to m nodes of a given graph such that no pair of connected nodes have the same color. We used a k-colorable graph generator [5] to obtain 200 instances of 10-colorable graphs with 65 nodes at the phase transition by varying the average node degree d uniformly around 0.6 [3]. Difference binary constraints were posted for every pair of connected nodes.

The Latin squares problem places N colors in a $N \times N$ grid, such that each color occurs exactly once on the same row or column. A partial Latin squares problem has several preassigned cells, and the goal is to complete the puzzle. We generated 200 instances of satisfiable partial Latin squares of size 35, with 396 cells preassigned, using lsencode-v1.1 [10]. The problem was modelled using alldifferent (GAC) constraints[1].

The results in fig. 1 (obtained using CaSPER[2] on a Pentium4, 1.7GHz with 512Mb RAM, timeout set to 1800 seconds, and time given in seconds) show that there is no clear winner across all problems.

As possibly expected, for problems where propagation achieves less pruning, impact based heuristics are less effective. This is the case of both the random and the graph coloring problems, as illustrated in fig. 2, that shows, in log scale, the reduction of the size of the search space, subsequent to each enumeration. In random problems the propagation is poor given the lack of structure of the problem.

[1] The dual encoding model, as proposed in [6], was also considered but never improved over the direct model on the presented instances.

[2] Code available from http://proteina.di.fct.unl.pt/casper, revision 333.

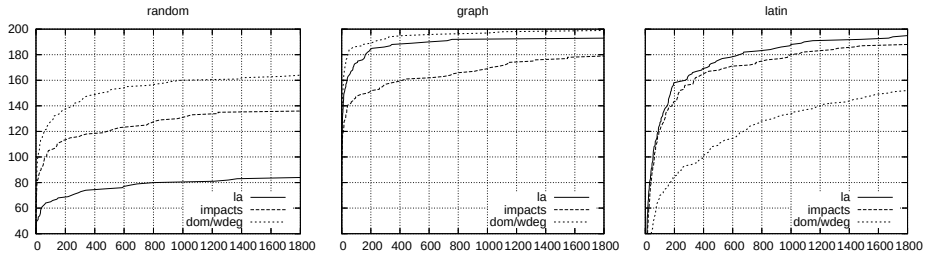

Fig. 1. Number of problems solved (vertical axis) versus time (horizontal axis).

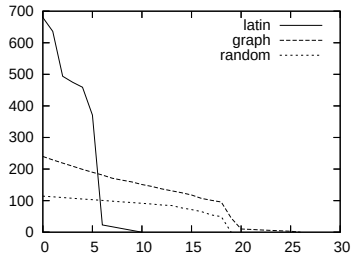

Fig. 2. Search space size during solving of a typical instance in each problem.

The networks for the graph coloring problem exhibit some locality, and propagation mostly affect variables in the same cluster of the variable being assigned, with limited propagation to variables far away in the network. Apparently, the dom/wdeg heuristics, by reasoning at constraint level, is able to "infer" such locality and take advantage of it for variable selection.

In contrast, in the latin square instances modeled by means of all different global constraints locality is not so marked (two variables often share the same constraint, if in the same row or column, or are separated by two constraints, one row and one column, and seldom by more than that, when both row and column pivot elements are already ground). Moreover, generalized arc consistency propagation virtually affects all variables after variable enumeration. The greater impact achieved in these problems, together with the lack of locality to be exploited by a constraint centered impact heuristic such as dom/wdeg, makes variable centered impact heuristics more adequate in this problem. We tried both heuristics in a set of smaller latin square instances modelled with pairwise distinct constraints while maintaining (singleton) node consistency and observed a different ranking, which confirms our thesis (see table 1).

strategy	#timeouts	avgfails	stddevfails	avgtime	stddevtime
dom/wdeg	**2**	63549	422439	**13.7**	94.6
la	42	**8022**	30312	208.3	366.4

Table 1. Results for finding the first solution to latin-15 with a selected strategy.

5 Conclusion

This paper focused on a class of heuristics exploiting some form of impact that decisions may have had in the past (dom/wdeg and impact) or will have in the near future (la). We have shown that the performance of these heuristics is correlated with the model used for the problem, more specifically with the amount and locality of constraint propagation.

The experiments reported in the paper show that work on impact based heuristics is far from over. The fact that there is no clear winner heuristic suggests combinations of these and other impact measures to obtain an efficient compromise between past and future impact information, and/or variable and constraint centered impacts.

Notwithstanding this future work, the new lookahead heuristic that we present in this paper already outperform, to our knowledge, all other heuristics on finite domain encodings of the latin square problems, and we believe that will also outperform other heuristics on problems with the same characteristics (low locality, large reduction of search space per choice).

References

1. C. Bessiére and J-C. Régin. Arc consistency for general constraint networks: preliminary results. In *Procs. of IJCAI'97*, pages 398–404, Nagoya, Japan, 1997.
2. Frédéric Boussemart, Fred Hemery, Christophe Lecoutre, and Lakhdar Sais. Boosting systematic search by weighting constraints. In *ECAI*, pages 146–150, Spain, 2004. IOS Press.
3. Peter Cheeseman, Bob Kanefsky, and William M. Taylor. Where the really hard problems are. In *Proceeding of IJCAI'91*, pages 331–340, 1991.
4. Marco Correia and Pedro Barahona. On the integration of singleton consistency and looh-ahead heuristics. In *Procs. of the annual ERCIM workshop on constraint solving and constraint logic programming*, Rocquencourt, France, June 2007.
5. Joseph Culberson. Graph coloring resources. on-line. http://web.cs.ualberta.ca/ joe/Coloring/Generators/generate.html.
6. Dotu, del Val, and Cebrian. Redundant modeling for the quasigroup completion problem. In *ICCP: Int. Conf. on Constraint Programming (CP), LNCS*, 2003.
7. Pieter Andreas Geelen. Dual viewpoint heuristics for binary constraint satisfaction problems. In *Procs. of ECAI '92*, pages 31–35, New York, NY, USA, 1992. John Wiley & Sons, Inc.
8. I. P. Gent, E. MacIntyre, P. Prosser, and T. Walsh. The constrainedness of search. In *Proceedings of AAAI'96*, volume 1, pages 246–252, 1996.
9. Ian P. Gent, E. MacIntyre, P. Prosser, Barbara M. Smith, and Toby Walsh. Random constraint satisfaction: Flaws and structure. *Constraints*, 6(4):345–372, 2001.
10. Kautz, Ruan, Achlioptas, Gomes, Selman, and Stickel. Balance and filtering in structured satisfiable problems. In *Procs. of IJCAI'01*, 2001.
11. Patrick Prosser, Kostas Stergiou, and Toby Walsh. Singleton consistencies. In Rina Dechter, editor, *Proceeding of CP'00*, volume 1894 of *Lecture Notes in Computer Science*, pages 353–368. Springer, 2000.
12. Philippe Refalo. Impact-based search strategies for constraint programming. In Mark Wallace, editor, *Procs. of CP'07*, volume 3258 of *Lecture Notes in Computer Science*, pages 557–571. Springer, 2004.