ToulBar2, an open source exact cost function network solver

D. Allouche, S. de Givry, T. Schiex
UR 875, INRA F-31320 Castanet Tolosan, France

Contributors: M. Sanchez (SP) , S. Bouveret (F), H. Fargier (F),
F. Heras (SP), P. Jégou (F), J. Larrosa (SP), K. L. Leung (CN), S. N'diaye (F),
E. Rollon (SP), C. Terrioux (F), G. Verfaillie (F), M. Zytnicki (F)

July 8, 2010

Abstract

This paper gives a quick description of the ToulBar2 solver following the UAI'2010 challenge on approximate inference in discrete stochastic graphical models where ToulBar2 finished respectively 2^{nd}, 1^{st} and 2^{nd} in the three categories representing optimization problems.

Cost function networks & graphical models

ToulBar2 is an exact combinatorial optimization tool targeted at Cost Function Networks (CFNs), also known as weighted Constraint Satisfaction problems [Meseguer et al., 2006]. This mathematical model has been derived from constraint satisfaction problems by replacing constraints with cost functions. In a CFN, we are given a set of variables with an associated finite domain and a set of local cost functions. Each cost function involves some variables and associates a non negative integer cost to each of the possible combinations of values they may take. The usual problem considered is to assign all variables in a way that minimizes the sum of all costs. Given an initial upper bound U, all costs above U can be considered as infinite, enabling the expression of usual constraints (strictly forbidden tuples). The first presentation of such "relaxed" constraint networks was given by [Rosenfeld et al., 1976] using max/min and max/* operations. It was later generalized to algebraic frameworks in [Schiex et al., 1995; Bistarelli et al., 1997]. Such frameworks allow for the concise description of global cost distributions, defined by the combination (usually the sum) of local cost functions.

Stochastic graphical models such as graph factors, Markov random field (MRF), chain graphs or bayesian nets represent (possibly non normalized) probability distributions by the product of local terms. Among the different inference problems that are usually considered, one is a discrete optimization problem called MPE (Maximum Probability Explanation), often denoted as MAP (Maximum A Posteriori) in MRF. To solve MPE, one must maximizes the joint probability. Maximizing the product of local functions is equivalent to minimizing the sum of the opposite of their logarithm. This transformation, followed by proper cost shifting, rescaling and discretization to get non negative integer costs, is used as the first step to translate stochastic graphical models into CFN.

1 The Toulbar2 solver

Toulbar2 is an exact solver for CFN with many bells and whistles. As a default, it uses a Depth First Branch and Bound algorithm to identify a minimum cost assignment and prove its optimality. The lower bound used for pruning during tree search is based on "soft local consistency enforcing" [Cooper and Schiex, 2004]. Soft local consistency allows to transform a CFN into an equivalent CFN with an associated non naive constant cost function that provides a strong incremental lower bound. The default level of enforcing used in ToulBar2 is known as "Existential Directed Arc Consistency" or EDAC [Larrosa et al., 2005]. This processing is only applied to cost function of arity 3 or less.

The tree explored is a binary tree where each node corresponds to either fixing a chosen variable to a chosen value or removing the value from its domain. For domains with a size above 10, a spliting strategy is used instead, where the domain of the chosen variable is split in 2 subddomains of similar size. Beyond this, ToulBar2 integrates on-the-fly variable elimination [Larrosa, 2000] which eliminates any variable with a low degree (default 3) during search. Each time a complete assignment is found, a new solution is printed, the upper bound is updated to the new solution cost and search proceeds until all possible combinations have been tried and optimality proven. In its default mode, ToulBar2 is therefore an exact anytime solver.

To choose a variable and a value at each node, dedicated heuristics are used. The variable chosen is selected using weighted degree [Boussemart et al., 2004] and last conflict heuristics [Lecoutre et al., 2009]. The value selected is a value with minimum cost in the associated unary cost function involving just this variable. This cost function can also be non naive, following local consistency enforcing, allowing to direct the search to promising area rapidly.

In order to better adapt to the stochastic graphical model area, different non default options have been activated for the challenge.

1. Bayesian nets often include large conditional probability tables involving many variables. The corresponding cost functions have a large arity and would have been ignored by local consistency enforcing until enough variables are assigned. This would lead to poor lower bounds and pruning. The "preproject" (`h`) option of ToulBar2 preprocesses the CFN by shifting cost from all cost functions of arity above 3 to binary and ternary cost functions.

2. to improve the anytime behavior of Toulbar2, an initial search phase combines restarts [Gagliolo and Schmidhuber, 2007] and Limited Discrepency Search [Harvey and Ginsberg, 1995] (`L` and `l` options). The number of nodes explored during this initial phase is bounded as well as the maximum number of discrepencies. This initial phase may already prove optimality. If not, it is followed by a complete tree search. For the UAI competition, the maximum number of discrepencies was set to the default value of 4. The maximum number of node for restarts was set to the default of $10,000$ in the 20" category and $100,000$ otherwise.

Beyond this version, two other variants of ToulBar2 have been submitted to the optimization categories of the challenge.

1. A first variant uses an additional stronger local consistency called "Virtual Arc Consistency" [Cooper et al., 2010] following EDAC enforcing (option `VA`) . This provides stronger lower bounds at the cost of extra time and may also help directing the value ordering. Globally, this variant was not better than the default but on the 20" category.

2. A second variant replaces the default Branch and Bound algorithm by a "graph structure-aware" Branch and Bound algorithm [Sanchez et al., 2009] exploiting an initially built cluster-tree decomposition. It is otherwise identical to the default above. This variant was no better.

2 Conclusion

ToulBar2 being an *exact* anytime solver, it is a nice surprise for us to see that it is competitive with the best *approximate* solvers that participated in the challenge. The table below gives the number of instances in each category and the number of problems solved to optimality by Toulbar2 strictly before the deadline has been reached.

Category	Nb. of instances	Nb. of opt. proofs
20"	160	130
20'	230	165
1 hour	195	135

Overall, ToulBar2 solved more than 73 % of all instances to optimality before the time deadline was reached. An optimality proof is not only a valuable information in itself. When it is produced quickly enough (in the majority of cases here), and cpu-resources are scarce, it allows to save time that can be used for solving other problems.

ToulBar2 includes other facilities which cannot be described here, including stronger preprocessing algorithms, and the ability to express global cost functions (involving a large number of variables, with a fixed semantics and dedicated efficient local consistency algorithms). Beyond stochastic graphical models, ToulBar2 and CFN have been used to solve real problems in resource allocation [Cabon et al., 1999], pedigree analysis [Sanchez et al., 2008] and bioinformatics [Zytnicki et al., 2008].

People interested in using or contributing to ToulBar2 can go the software forge hosting the project at https://mulcyber.toulouse.inra.fr/projects/toulbar2. More information on ToulBar2 and on the CostFunctionLib, a large collection of benchmarks (real, academic and random) of cost function networks can be found at http://carlit.toulouse.inra.fr/cgi-bin/awki.cgi/ToolBarIntro.

References

S. Bistarelli, U. Montanari, and F. Rossi. Semiring based constraint solving and optimization. *Journal of the ACM*, 44(2):201–236, 1997.

F. Boussemart, F. Hemery, C. Lecoutre, and L. Sais. Boosting systematic search by weighting constraints. In *ECAI 2004: 16th European Conference on Artificial Intelligence*, page 146, Valencia, Spain, 2004.

B. Cabon, S. de Givry, L. Lobjois, T. Schiex, and J.P. Warners. Radio link frequency assignment. *Constraints Journal*, 4:79–89, 1999.

M C. Cooper and T. Schiex. Arc consistency for soft constraints. *Artificial Intelligence*, 154(1-2):199–227, 2004.

M.C. Cooper, S de Givry, M Sanchez, T Schiex, M Zytnicki, and T Werner. Soft arc consistency revisited. *Artificial Intelligence*, 174(7-8):449–478, May 2010. ISSN 00043702. doi: 10.1016/j.artint.2010.02.001. URL http://linkinghub.elsevier.com/retrieve/pii/S0004370210000147.

M. Gagliolo and J. Schmidhuber. Learning restart strategies. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*, volume 1, pages 792–797, 2007.

W. D. Harvey and M. L. Ginsberg. Limited discrepency search. In *Proc. of the* 14^{th} *IJCAI*, volume 14, pages 607–615, Montréal, Canada, 1995.

J. Larrosa. Boosting search with variable elimination. In *Principles and Practice of Constraint Programming - CP 2000*, volume 1894 of *LNCS*, pages 291–305, Singapore, September 2000.

J. Larrosa, S. de Givry, F. Heras, and M. Zytnicki. Existential arc consistency: getting closer to full arc consistency in weighted CSPs. In *Proc. of the* 19^{th} *IJCAI*, pages 84–89, Edinburgh, Scotland, August 2005.

C. Lecoutre, L Saïs, S. Tabary, and V. Vidal. Reasoning from last conflict(s) in constraint programming. *Artificial Intelligence*, 173:1592,1614, 2009.

P. Meseguer, F. Rossi, and T. Schiex. Soft constraints processing. In F. Rossi, P. van Beek, and T. Walsh, editors, *Handbook of Constraint Programming*, chapter 9. Elsevier, 2006.

A. Rosenfeld, R. Hummel, and S. Zucker. Scene labeling by relaxation operations. *IEEE Trans. on Systems, Man, and Cybernetics*, 6(6):173–184, 1976.

M Sanchez, D Allouche, S de Givry, and T Schiex. Russian Doll Search with Tree Decomposition. In *Proc. of IJCAI'09*, Pasadena (CA), USA, 2009. URL http://www.inra.fr/mia/T/schiex/Export/IJCAI2009.pdf.

Marti Sanchez, Simon Givry, and Thomas Schiex. Mendelian Error Detection in Complex Pedigrees Using Weighted Constraint Satisfaction Techniques. *Constraints*, 13(1-2):130–154, January 2008. ISSN 1383-7133. doi: 10.1007/s10601-007-9029-5. URL http://www.springerlink.com/index/10.1007/s10601-007-9029-5.

T. Schiex, H. Fargier, and G. Verfaillie. Valued constraint satisfaction problems: hard and easy problems. In *Proc. of the* 14^{th} *IJCAI*, pages 631–637, Montréal, Canada, August 1995.

Matthias Zytnicki, Christine Gaspin, and Thomas Schiex. DARN! A Weighted Constraint Solver for RNA Motif Localization. *Constraints*, 13(1-2):91–109, February 2008. ISSN 1383-7133. doi: 10.1007/s10601-007-9033-9. URL http://www.springerlink.com/index/10.1007/s10601-007-9033-9.