

Multi-Language Evaluation of Exact Solvers in Graphical Model Discrete Optimization

Barry Hurley · Barry O’Sullivan ·
David Allouche · George Katsirelos ·
Thomas Schiex · Matthias Zytnicki ·
Simon de Givry

Received: date / Accepted: date

Abstract By representing the constraints and objective function in factorized form, graphical models can concisely define various NP-hard optimization problems. They are therefore extensively used in several areas of computer science and artificial intelligence. Graphical models can be deterministic or stochastic, optimize a sum or product of local functions, defining a joint cost or probability distribution. Simple transformations exist between these two types of models, but also with MaxSAT or linear programming.

In this paper, we report on a large comparison of exact solvers which are all state-of-the-art for their own target language. These solvers are all evaluated on deterministic and probabilistic graphical models coming from the Probabilistic Inference Challenge 2011, the Computer Vision and Pattern Recognition OpenGM2 benchmark, the Weighted Partial MaxSAT Evaluation 2013, the MaxCSP 2008 Competition, the MiniZinc Challenge 2012 & 2013, and the CFLib (a library of Cost Function Networks).

All 3026 instances are made publicly available in five different formats and seven formulations. To our knowledge, this is the first evaluation that encompasses such a large set of related NP-complete optimization frameworks, despite their tight connections. The results show that a small number of evaluated solvers are able to perform well on multiple areas. By exploiting the variability and complementarity of solver performances, we show that a simple portfolio approach can be very effective. This portfolio won the last UAI Evaluation 2014 (MAP task).

Keywords Graphical Model, Markov Random Field, Weighted Constraint Satisfaction Problem, Integer Linear Programming, MaxSAT

Barry Hurley · Barry O’Sullivan
Insight Centre for Data Analytics, University College Cork, Ireland
E-mail: firstname.lastname@insight-centre.org

David Allouche · George Katsirelos · Thomas Schiex · Matthias Zytnicki · Simon de Givry
MIAT, UR-875, INRA, F-31320 Castanet Tolosan, France
E-mail: firstname.lastname@toulouse.inra.fr

1 Introduction

Graphical Models can concisely represent highly dimensional multivariate functions using a factorization into local functions. We consider discrete variables.

Constraint Networks and weighted variants such as Cost Function Networks (CFNs), aka (Weighted) Constraint Satisfaction Problems ((W)CSPs), aim at finding an assignment of all variables that minimizes a joint cost function defined as the sum of local functions (constraints being represented as functions with values in $\{0, \infty\}$). With Boolean variables, and a language restricted to clausal form, the (partial weighted Max)-SAT problem has the same target. Constraint Programming (CP), an extension of Constraint Networks including non-deterministic programming language features, can also easily capture these optimization problems by introducing cost variables [43].

In AI and statistics, probabilistic graphical models [29] use the same idea to concisely represent probability distributions over random variables. These models include Bayesian Networks and Markov Random Fields (MRFs). The problem of identifying a variable assignment that has maximum probability is called the *Maximum Probability Explanation* in Bayesian networks, or *Maximum A-Posteriori* (MAP) in MRF. This NP-hard problem has an extremely large application scope, *e.g.*, in image processing or bioinformatics. By a simple ($-\log$) transformation, these problems can be reduced to CFNs.

Graphical Models can also be easily encoded as 0/1 Linear Programming (01LP) problems, a standard language for Operations Research (OR). We consider two encodings, including one based on the so-called *local polytope* [47, 30, 20], which has several interesting properties.

In this paper, we extract probabilistic and deterministic graphical models from various areas, each using a specific language. This covers competitions in MaxSAT, constraint programming, probabilistic inference and repositories in probabilistic image processing and cost function networks. We encode them in these underlying languages and close relatives, from AI (CFN, MaxSAT, MRF), CP, and OR (01LP). These benchmarks are traditionally used in competitions relying on a single language with dedicated solvers. We compare exact solvers which are all state-of-the-art for their own language on these encodings. We then define a novel portfolio hybrid solver exploiting them.

2 Combinatorial optimization languages

In this section we briefly describe the combinatorial optimization languages that will be used.

[CFN] Cost Function Networks, or Weighted Constraint Networks, extend Constraint Networks by using non-negative cost functions instead of constraints [38].

Definition 1 A Cost Function Network (CFN) is a triple (X, W, k) where $X = \{1, \dots, n\}$ is a set of n discrete variables, W is a set of non-negative

functions, and k , a (possibly infinite) maximum cost. Each variable $i \in X$ has a finite domain D_i of values that can be assigned to it. A function $w_S \in W$, with scope $S \subseteq X$, is a function $w_S : D_S \mapsto \{\alpha \in \mathbb{N} \cup \{k\} : \alpha \leq k\}$, where D_S denotes the Cartesian product of all D_i for $i \in S$.

In CFNs, the cost of a complete assignment is the sum of all cost functions. A solution has cost less than k . Therefore a cost of k denotes forbidden assignments, used in hard constraints. A solution of minimum cost is sought.

[MRF] Markov Random Fields define a joint probability distribution. The terminology of *Graphical Models* (GMs) originally designates *probabilistic graphical models* such as Markov Random Fields (MRFs) and Bayesian Networks (BNs) [29]. In this paper, we restrict ourselves to MRFs because they do not impose any restriction on the local functions that can be used in the decomposition of the joint probability distribution (BNs use local conditional probabilities with a normalization requirement).

Definition 2 A discrete Markov Random Field (MRF) is a pair (X, Φ) where $X = \{1, \dots, n\}$ is a set of n random variables, and Φ is a set of potential functions. Each variable $i \in X$ has a finite domain D_i of values that can be assigned to it. A potential function $\phi_S \in \Phi$, with scope $S \subseteq X$, is a function $\phi_S : D_S \mapsto \mathbb{R} \cup \{\infty\}$.

The probability of a tuple $t \in D_X$ is defined as:

$$P(t) \propto \prod_{\phi_S \in \Phi} \exp(-\phi_S(t[S])) = \exp\left(-\sum_{\phi_S \in \Phi} \phi_S(t[S])\right)$$

where $t[S]$ denotes the restriction of t to the set of variables S . The additive potentials ϕ_S are called energies, in relation with statistical physics. Alternatively, multiplicative $\exp(-\phi_S(t[S]))$ potentials can be used.

In this paper, we consider the MAP query that aims at finding a complete assignment of maximum probability (or equivalently, minimum energy).

[WPMS] Weighted Partial MaxSAT problems are CFNs restricted to Boolean domains and a language of weighted clauses [36].

Definition 3 A Weighted Partial MaxSAT (WPMS) instance is defined as a set of pairs $\langle C, w \rangle$ and an upper bound k . Each C is a clause and w is a number in $\mathbb{N} \cup \{k\}$, the *weight* of clause C . A clause is a disjunction of literals. A literal is a Boolean variable or its negation.

A clause with weight $\geq k$ is a *hard* clause, otherwise it is *soft*. The objective is to find an assignment to the variables appearing in the clauses that minimizes the sum of the weights of all falsified clauses, which should be of cost $< k$.

[01LP] A 0/1 Linear Program is defined by a linear objective function over a set of 0/1 variables to minimize under a conjunction of linear equalities and inequalities [25].

[CP] Constraint Programming problems are defined by a set of discrete variables and a set of constraints. The aim is to minimize the value of a given objective variable while satisfying all constraints [45].

3 Translations between formalisms

In this section we present encodings between graphical models represented in each of the AI/OR/CP languages we presented. We summarize in Table 1 for each input formalism the different translations used to produce every instance in the corresponding output formalism.

[MRF] Markov Random Field. With additive potentials, MRFs are essentially equivalent to CFNs except for the fact that they can use arbitrary real-valued potential functions instead of integer non-negative costs¹. Additive MRFs can therefore be reduced to CFNs using a fixed decimal point representation of energies which are then scaled to integers and shifted to enforce non-negativity. This preserves optimal solutions.

Multiplicative MRFs can be transformed to additive MRFs using a simple $(-\log)$ transform, and then to CFNs [18,17]. Conversely, CFNs can be transformed to multiplicative MRFs (as in the UAI *MARKOV* format) by exponentiating costs². Costs are all shifted by the same amount so that the largest multiplicative potentials are equal to 1. Hard costs ($\geq k$) are translated to a zero multiplicative potential (infinite energy) to preserve the ability to prune domain values based on constraint reasoning.

[WPMS] Weighted Partial MaxSAT. As weighted partial MaxSAT is a CFN with Boolean variables and a language of clauses, thus a WPMS instance is already a CFN. For a CFN, we consider two encodings to WPMS based on CSP to SAT encodings: the *direct* encoding [6], and the *tuple* encoding introduced by Bacchus [7]. WPMS costs are non-negative integers and the WCNF format allows to express an upper bound that will be used to represent k , preserving the ability to prune.

Direct encoding: for each variable i with domain size $|D_i| > 2$, we use one proposition $d_{i,r}$ for each value $r \in D_i$. This proposition is true iff variable i is assigned the value r . To ensure that exactly (At Least and At Most One) one value is used for each variable, we encode *At Most One* (AMO) with hard clauses $(\neg d_{i,r} \vee \neg d_{i,s})$ for all $i \in \{1, \dots, n\}$ and all $r < s$, $r, s \in D_i$, as well as *At Least One* (ALO) with one hard clause $(\bigvee_r d_{i,r})$ for each i . Boolean variables are directly encoded as propositions and do not require AMO/ALO clauses. Then, for each cost function $w_S \in W$ and each tuple $t \in D_S$ with $w_S(t) > 0$, we have a clause $(\bigvee_{i \in S} \neg d_{i,t[i]})$ with weight $w_S(t)$, where $d_{i,t[i]}$ denotes the proposition associated with assigning to variable i the value that it has in tuple t .

¹ Rational costs are also used in [11].

² Script available at genoweb.toulouse.inra.fr/~degivry/evalgm/scripts/wcsp2markov.py

Tuple encoding: it encodes domains as in the direct encoding. We have a proposition $d_{i,r}$ for each variable/value pair representing $i = r$, along with AMO/ALO clauses that enforce that each variable is assigned exactly one value (for non-Boolean variables). Nullary and unary cost functions are also represented as soft clauses exactly as in the direct encoding.

For each cost function $w_S, |S| > 1$, and each tuple $t \in D_S$ with $w_S(t) < k$ we have a proposition $p_{S,t}$. For non-zero cost $w_S(t) > 0$, we have the soft clause $(\neg p_{S,t})$ with weight $w_S(t)$. This represents the cost to pay if the tuple t is used. For every variable $i \in S$, we have a hard clause $(d_{i,t[i]} \vee \neg p_{S,t})$. These clauses enforce that if tuple t is used, its values $t[i]$ must be used. Then, for each variable $i \in S$ and each value $r \in D_i$, we have hard clauses $(\neg d_{i,r} \vee \bigvee_{t \in D_S, t[i]=r, w_S(t) < k} p_{S,t})$ that enforce that if a value $r \in D_i$ is used, one of the allowed tuples $t \in D_S$ such that $t[i] = r, w_S(t) < k$ must be used.

On CSP, it is known that Unit Propagation (UP) on the tuple encoding enforces arc consistency in the original CSP (the set of values that are deleted by enforcing AC have their corresponding literals set to false by UP) [7].

We express the asymptotic complexities of the two encodings in terms of the total number of tuples of cost 0 (t_0), k (t_k) or other (t_r) in the problem. For the direct encoding, this is $O(nd^2 + t_k + t_r)$, while for the tuple encoding this is $O(nd^2 + a(t_0 + t_r))$, where n is the number of variables, d is the maximum domain size, and a is the maximum cost function arity. The hidden big- O constants are larger for the tuple encoding, which has an additional linear factor a . In our experiments (see Table 2 in Sect. 4.1), we found that the tuple encoding is typically much larger, more than can be accounted for by the hidden constants. Hence it appears that our benchmark instances have many more tuples with zero cost than with infinite (k) cost ($t_0 \gg t_k$).

[01LP] 0/1 Linear Programming. The 01LP encodings of CFNs are similar to those for WPMS, using 0/1 variables. The additional expressivity of linear constraints enables further simplifications. These translations are used to generate 01LP in CPLEX “LP” format.

Direct encoding: AMO/ALO clauses are replaced by one linear constraint per non-Boolean variable $i \in X$: $\sum_{r \in D_i} d_{i,r} = 1$. For each cost function w_S , the soft clause encoding of a tuple t with non-zero soft cost $0 < w_S(t) < k$ is replaced by a linear constraint $\sum_{i \in S} (1 - d_{i,t[i]}) + p_{S,t} \geq 1$ that forces the value of $p_{S,t}$ to 1 if the tuple t is used. This $p_{S,t}$ variable appears in the objective function, with a coefficient $w_S(t)$. If t has cost k or above, a constraint $\sum_{i \in S} (1 - d_{i,t[i]}) \geq 1$ is used and no term appears in the objective function.

Tuple encoding: the same encoding as above is used for domains and for zero and unit-arity cost functions. For each cost function $w_S, |S| > 1$, for each variable $i \in S$, each value $r \in D_i$, a constraint $\sum_{t \in D_S, t[i]=r, w_S(t) < k} p_{S,t} = d_{i,r}$ enforces that a value (i, r) is used iff a tuple t s.t. $t[i] = r$ and $w_S(t) < k$ is used. The same 0/1 variable $p_{S,t}$ appears in the objective function with a $w_S(t)$ coefficient if $0 < w_S(t) < k$.

This encoding has been proposed by Koster in [30] to encode Partial Constraint Satisfaction Problems. Since all $d_{i,r}$ are 0/1 variables, the constraints enforce that the $p_{S,t}$ are also integral. We therefore relax the integrality constraint on $p_{S,t}$ variables.

Assuming there are no costs in $\{0, \infty\}$, for each cost function w_S , each variable i , and each value $r \in D_i$, by summing the linear constraints $\sum_{i \in S} (1 - d_{i,t[i]}) + p_{S,t} \geq 1$ from the direct encoding over all tuples $t \in D_S$ such that $t[i] = r$, we found:

$$\begin{aligned} M|S| - \sum_{j \in S \setminus \{i\}} \frac{M}{|D_j|} \left(\sum_{s \in D_j} d_{j,s} \right) - Md_{i,r} + \sum_{t \in D_S, t[i]=r} p_{S,t} &\geq M \\ M|S| - \sum_{j \in S \setminus \{i\}} \frac{M}{|D_j|} (1) - Md_{i,r} + \sum_{t \in D_S, t[i]=r} p_{S,t} &\geq M \\ M|S| - \frac{M(|S| - 1)}{d} - Md_{i,r} + \sum_{t \in D_S, t[i]=r} p_{S,t} &\geq M \\ M(|S| - 1) - \frac{M(|S| - 1)}{d} - Md_{i,r} + \sum_{t \in D_S, t[i]=r} p_{S,t} &\geq 0 \end{aligned}$$

$$\text{Thus, } \sum_{t \in D_S, t[i]=r} p_{S,t} \geq M(d_{i,r} - (|S| - 1)(1 - \frac{1}{d}))$$

with $d = \max_{j \in S \setminus \{i\}} |D_j|$ and $M = |D_{S \setminus \{i\}}|$, the Cartesian product of all domains of S except D_i . If $|S| = 2$, then $M = d$, and $M(d_{i,r} - (|S| - 1)(1 - \frac{1}{d}))$ is either negative ($d_{i,r} = 0$) or equal to 1 ($d_{i,r} = 1$). Therefore, the direct encoding can be seen as a relaxation of the tuple encoding.

The continuous relaxation of the tuple encoding is known in the MRF field as the *local polytope* [47, 50, 20]. This polytope is interesting for several reasons. First, the dual of the local polytope is exactly the Optimal Soft Arc Consistency (OSAC) LP for CFN described in [12, 11]. This polytope underlies also convergent message-passing bounds [50, 20] used for MRF optimization. Ignoring possible value pruning (by node consistency or substitutability [19]), OSAC and therefore the local polytope bound too, are known to be stronger than any other soft arc consistency [11]. Second, the dual variables of this polytope can be directly interpreted as the amount of cost that is shifted by arc consistency so-called Equivalence Preserving Transformations [13]. Therefore, existing soft arc consistencies that iteratively change blocks of costs can be analyzed as fast incremental approximate Block Coordinate Descent algorithms aiming at solving this dual LP [37]. This result establishes a strong link between 01LP solvers using the local polytope encoding and CFN/MRF solvers using soft arc consistencies or convergent message passing: in absence of pruning, the LP bound will always be at least as strong as the soft arc consistency bounds.

The significance of this connection is further strengthened by a recent result showing that the local polytope (or its dual) are “universal” in the sense that

any LP can be translated *in linear time* in a graphical model whose local polytope has the same optimum as the original LP [44]. Progress in solving this polytope (exactly or approximately by soft arc consistencies or message passing) and in solving a general LP are therefore tightly linked.

[CP] Constraint Programming. In [43], a translation of CFNs into crisp CSPs has been proposed. In this transformation, the decision variables of the CFN are preserved and every cost function is reified into a constraint whose scope is augmented by one extra variable, representing the assignment cost. This reification of costs into domain variables transforms a CFN in a crisp CSP with more variables and increased arities. Typically, unary and binary cost functions are converted into TABLE constraints of arity two and three respectively. Another extra cost variable encodes the objective function, connected by a SUM constraint to all other cost variables. All the cost variables are non-negative integers with the same initial upper bound k as provided in the WCSP format. The same approach applies to WPMSs, using reified Boolean expressions instead of TABLE constraints to encode hard and soft clauses. The resulting CSP models are expressed in the MINIZINC [39] CP language³.

The converse translation of CP models with a cost variable into a CFN (and then MRFs and WPMSs) that does not use cost variables is a complex task⁴. It requires identifying local⁵ cost functions, starting from the objective variable, while removing intermediate cost variables. We implemented a corresponding prototype in NUMBERJACK⁶ [22] reading the low-level FLATZINC format [39]. Global constraints are decomposed into ternary cost functions in extension (tables with costs in $\{0, \infty\}$, see [1]), requiring small input domain sizes.

Table 1: Summary of translations between formalisms and possible issues

In/Out	MRF (UAI)	CFN (WCSP)	WPMS (WCNF)	01LP (LP)	CP (MINIZINC)
MRF	-	$-\log(prob)$	Through CFN	Through CFN	Through CFN
CFN	$\exp(-cost)$	-	Direct/ tuple encod.	Direct/ tuple encod.	Extra cost ⁷ vars & table constraints
WPMS	Through CFN ⁸	Direct trans.	Direct encod. only	Through CFN ⁹	Extra cost ¹⁰ vars & reified logical <i>or</i>
CP	Through CFN	Decomposed objective & global constraints ¹¹	Through CFN	Through CFN ¹²	-

³ A 1-hour time limit was used to translate MINIZINC2 to FLATZINC, readable by CP solvers.

⁴ Directly lifting a CP model, with its cost variable, to a CFN would be of limited value since all AC in CFN are known to enforce AC on cost functions representing hard constraints.

⁵ We restrict the size of cost functions to be less than 10^6 tuples in our implementation.

⁶ <http://numberjack.ucc.ie/>

4 Graphical model evaluation

We have collected a set of benchmarks and performed experiments using state-of-the-art solvers coming from several research areas.

4.1 Collection of benchmarks

To gather an extensive set of benchmarks representing optimization problems from various areas, we collected problems from different sources including deterministic (CFN, MaxCSP, WPMS), probabilistic (MRF, BN), as well as CP collections. Each collection contains several categories of instances, each category corresponding to a specific class of problems.

[MRF]: the Probabilistic Inference Challenge (PIC) 2011 benchmark set¹³ and the (5-ary) genetic linkage analysis problem [18] from the Uncertainty in Artificial Intelligence (UAI) 2008 Evaluation¹⁴ were taken in UAI *MARKOV* format with multiplicative potentials. This PIC challenge on approximate inference in probabilistic graphical models is dedicated to a variety of queries and we only considered the MAP/MPE query. We used a subset of the instances available in PIC 2011, excluding Alchemy, CSP, Promedas, and ProteinProtein¹⁵. These problems have been translated to CFNs in WCSP format (then to WPMS, 01LP, CP) using a $(-\log)$ transform followed by fixed decimal point representation with 2-digit precision after the decimal point (the precision is constrained by CP solvers that typically accept only 32-bit integers)¹⁶.

[CVPR]: the Computer Vision and Pattern Recognition (CVPR) OpenGM2 benchmark¹⁷ [27] contains binary and ternary MRF instances in HDF5 format with additive potentials. We excluded Brain, Knott, and MatchingStereo/tedgm instances because of their size ($> 1GB$), and ModularityClustering because it came from outside the computer vision community. ColorSeg, Match-

⁷ Cannot represent large costs ($> 2^{31}$) using a single domain. CFLib benchmarks were manually translated, avoiding table constraints except for ProteinDesign and SPOT5.

⁸ Cannot represent large clauses (> 23 literals in our case) using complete tables.

⁹ No tuple encoding.

¹⁰ Cannot represent large costs ($> 2^{31}$) using a single domain.

¹¹ Cannot represent large domains in extension ($d > 1,000$) and non-decomposable objectives (requiring cost functions with $> 10^6$ tuples).

¹² This translation is far from being optimal, *e.g.*, linear constraints will be first decomposed in ternary cost functions.

¹³ <http://www.cs.huji.ac.il/project/PASCAL>

¹⁴ <http://graphmod.ics.uci.edu/uai08/Evaluation/Report/Benchmarks>

¹⁵ Alchemy and Promedas were solved by TOULBAR2 in less than 1 sec. each. CSP instances came from CFLib. ProteinProtein is already present in CVPR under the name of *Protein Prediction* ProteinInteraction.

¹⁶ The resulting WCSP instances were translated back to UAI instances (with *-digit2* extension) in order to optimize the same objective function.

¹⁷ <http://hci.iwr.uni-heidelberg.de/opengm2>

ingStereo, PhotoMontage have integer energies, directly defining non-negative costs. For the others, we used 8-digit precision after the decimal point.

[CFLib]: the CFLib¹⁸ is a collection of CFN and MaxSAT problems expressed in WCSP format. We extracted problems that are directly available in the WCSP format and further translated them into dedicated MINIZINC models manually. The extracted benchmarks include combinatorial auctions [35], CELAR/GRAPH radio-link frequency assignment problems [10], Mendelian error correction problems on complex pedigrees [46], computational protein design problems [3] (with 2-digit precision), SPOT5 satellite scheduling problems [8], and uncapacitated warehouse location problems [33,34].

[MaxCSP]: all binary CSP categories with table constraints and at least one inconsistent instance (BlackHole, Langford, Quasi-group Completion Problem, Graph Coloring, random Composed, random 3-SAT EHI, and random Geometric, excluding pure random categories) from the CSP 2008 Competition¹⁹ were translated from XCSP2.1/XML format to CFNs (as MaxCSPs) where allowed (resp. forbidden) tuples have zero (resp. unit) cost. We set $k = 1,000$.

[WPMS]: weighted partial MaxSAT instances coming from the MaxSAT 2013 Evaluation²⁰, including crafted MIPLib, DIMACS Max Clique, and industrial WPMS instances, have been directly encoded as CFNs, each clause being encoded as a cost function with just one non-zero cost tuple. Translation to MRF (resp. CP) was restricted to instances with small-arity clauses (resp. with 32-bit costs, excluding the WPMS/Upgradeability category).

[CP]: we extracted a selection of CFN-decomposable CP problems from the MiniZinc Challenges 2012 & 2013²¹. Only the smallest instances in FastFood, Golomb, and OnCallRostering categories could be decomposed in WCSP format using less than 1GB per instance (resp. 1, 3, and 3 instances per category).

Together, these benchmark resources contain problems offering a large variety in terms of size, maximum arity or domain size and cost range. WPMS and CVPR categories have the highest number of variables (close to 1 million variables for WPMS/TimeTabling, half a million for CVPR/PhotoMontage and ColorSeg). The WPMS benchmark also has the largest arities (a weighted clause on 580 variables appears in Haplotyping). For the other benchmarks, maximum arity varies from 2 to 5. Graph connectivities are usually very small for MRF&CVPR (often based on grid graphs where vertices represent pixels in images) and WPMS benchmarks. MRF/ObjectDetection, CFN/ProteinDesign, MaxCSP/Langford, and CVPR/Matching have complete graphs. MRF/ProteinFolding has the largest domain size (503 values). Most CVPR instances have very large cost ranges (8-digit precision), whereas MaxCSP instances contain only 0/1 costs. The emphasis between optimization and feasibility also varies a lot among the problems: almost all deterministic GM categories, except MaxCSPs and CFN/CELAR, contain forbidden (k) tuples in their cost

¹⁸ <http://costfunction.org/benchmark>

¹⁹ <http://www.cril.univ-artois.fr/{CPAI08|~lecoutre/benchmarks.html}>

²⁰ <http://maxsat.ia.udl.cat:81/13/benchmarks/>

²¹ <http://www.minizinc.org/challenge201{2|3}/results201{2|3}.html>.

functions. On the contrary, probabilistic GMs usually have no forbidden tuples (except for MRF/Linkage and DBN).

Table 2 reports the number of instances per benchmark resource and its gzipped size for the seven formulations. The UAI format appears to be the most compact to express local functions as tables. It relies on a *complete* ordered table of costs which does not require describing tuples whereas the other formats explicitly describe tuples associated to non-zero costs. The price to pay for this conciseness is the inability of the UAI format to represent large arity functions with a few non-zero costs (such as large weighted clauses). As seen before, the tuple encoding is usually larger than the direct one, except for MRF/CVPR LPs where the local polytope is a good choice since there are almost no zero costs. CP instances benefit from global constraints in the MINIZINC language, which are decomposed in large tables in the other formats.

Table 2: Number of instances and their total compressed (gzipped) size per format for each benchmark resource

Benchmark	Nb.	UAI	WCSP	LP(direct)	LP(tuple)	WCNF(direct)	WCNF(tuple)	MINIZINC
MRF	319	187MB	475MB	2.4G	2.0GB	518MB	2.9GB	473MB
CVPR	1461	430MB	557MB	9.8GB	11GB	3.0GB	15GB	N/A
CFN	281	43MB	122MB	300MB	3.5GB	389MB	5.7GB	69MB
MaxCSP	503	13MB	24MB	311MB	660MB	73MB	999MB	29MB
WPMS	427	N/A ²²	387MB	433MB	N/A	717MB	N/A	631MB
CP	35	7.5MB	597MB	499MB	1.2GB	378MB	1.9GB	21KB
Total	3026	0.68G	2.2G	14G	18G	5G	27G	1.2G

4.2 Experimental settings

We compared state-of-the-art MRF solvers²³ DAOPT²⁴ [42] (using its 1-hour settings), winner of PIC 2011, and TOULBAR2²⁵ [34, 18] (including Virtual Arc Consistency (VAC) as preprocessing [11], dominance rule pruning [19], and hybrid best-first search [2]), winner of MaxCSP 2008 and UAI 2010 & 2014 Evaluations, against WPMS MAXHS²⁶ solver [14, 15], winner of crafted WPMS MaxSAT 2013, the CP solver GECODE²⁷, winner of MiniZinc Challenges 2012, and IBM-ILOG CPLEX 12.6.0.0 (using parameters EPAGAP, EPGAP, and EPINT set to zero to avoid premature stop).

²² Only WPMS/MaxClique and WPMS/MIPLib (except mod008) could be translated in UAI format for a total of 8.8MB gzipped size.

²³ MPLP2 <http://cs.nyu.edu/~dsontag> version 2 (using 2.10^{-7} gap thres.) was tested but the results are not presented in Sec. 4.3 as it was dominated in most categories by TOULBAR2.

²⁴ <https://github.com/lotten/daoapt> open source version 1.1.2, not including the closed source and unavailable convergent message-passing bound tightening used in the PIC challenge.

²⁵ <http://www.inra.fr/mia/T/toulbar2> version 0.9.8, parameters `-A -V -dec -hbfs`.

²⁶ <http://www.maxhs.org> version 2.51, no parameter.

²⁷ <http://www.gecode.org/> version 4.4.0, using free search.

All computations were performed on a single core of AMD Opteron 6176 at 2.3 GHz and 8 GB of RAM with a 1-hour CPU time limit²⁸.

4.3 Experimental results

Table 3: Number of problems solved in less than 1 hour (N/A if RAM usage or 32-bit limit prevented encoding). In parentheses, mean CPU time in seconds on solved instances ('-' if none). Bold is best. The first column contains the category name followed by *s*: nb. of instances, *d*: max. dom. size, *a*: max. arity

Problem/ <i>s/d/a</i>	DIAGOPT	FOUDBAY2	CPLEX	CPLEX-mpplc	MAXHS	MAXHS-mpplc	CPXCDBE
MRF/319/503/5 (UAI)	151 (584.39)	226 (93.80)	156 (111.88)	210 (82.18)	118 (172.27)	72 (98.68)	1 (1509.93)
DBN/108/2/2	60 (626.79)	81 (192.42)	65 (124.66)	69 (155.12)	38 (366.15)	2 (1748.65)	0 (-)
Grid/21/2/2	5 (1223.67)	0 (-)	15 (120.90)	1 (3354.21)	8 (557.01)	0 (-)	0 (-)
ImageAlignment/10/93/2	10 (754.96)	10 (5.27)	0 (-)	9 (88.41)	0 (-)	0 (-)	0 (-)
Linkage/22/7/5	17 (576.94)	14 (364.73)	16 (365.09)	22 (21.99)	20 (52.62)	20 (124.04)	0 (-)
ObjectDetection/37/21/2	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)
ProteinFolding/21/503/2	0 (-)	21 (20.24)	10 (169.28)	9 (176.17)	2 (268.51)	0 (-)	0 (-)
Segmentation/100/21/2	59 (460.33)	100 (0.29)	50 (0.06)	100 (3.35)	50 (7.38)	50 (22.54)	1 (1509.93)
CVPR/1461/20/3 (HDF5)	1274 (481.02)	1340 (22.81)	382 (179.96)	1332 (8.70)	483 (355.71)	1038 (58.83)	N/A
ChineseChars/100/2/2	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	N/A
ColorSeg/21/12/2	0 (-)	15 (1340.56)	0 (-)	5 (190.33)	0 (-)	0 (-)	N/A
GeomSurf/600/7/3	555 (509.01)	600 (0.96)	382 (179.96)	600 (2.89)	387 (183.53)	321 (63.15)	N/A
InPainting/4/4/2	0 (-)	2 (325.72)	0 (-)	1 (339.90)	0 (-)	0 (-)	N/A
Matching/4/20/2	4 (319.24)	4 (3.20)	0 (-)	3 (765.25)	0 (-)	0 (-)	N/A
MatchingStereo/2/20/2	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	N/A
ObjectSeg/5/8/2	0 (-)	4 (2292.28)	0 (-)	5 (1057.88)	0 (-)	0 (-)	N/A
PhotoMontage/2/7/2	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	N/A
ProteinInteraction/8/2/3	0 (-)	0 (-)	0 (-)	3 (60.80)	0 (-)	2 (1019.63)	N/A
SceneDecomp/715/8/2	715 (460.20)	715 (0.07)	0 (-)	715 (1.11)	96 (1049.83)	715 (54.20)	N/A
CFN/281/300/3 (WCSP)	211 (768.93)	256 (109.84)	245 (33.85)	238 (34.00)	228 (14.91)	210 (121.20)	141 (224.77)
Auction/170/2/2	169 (663.04)	170 (93.10)	170 (0.03)	170 (0.14)	170 (0.03)	170 (121.16)	113 (231.55)
CELAR/16/44/2	4 (598.72)	14 (279.00)	0 (-)	3 (560.44)	0 (-)	0 (-)	0 (-)
Pedigree/10/28/3	4 (373.43)	10 (10.58)	5 (44.28)	9 (57.27)	10 (190.49)	6 (99.28)	0 (-)
ProteinDesign/10/198/2	4 (597.46)	9 (13.40)	0 (-)	7 (298.88)	0 (-)	4 (477.72)	0 (-)
SPOT5/20/4/3	6 (309.04)	4 (40.44)	16 (22.99)	12 (294.94)	6 (200.82)	5 (5.40)	0 (-)
Warehouse/55/300/2	24 (1752.42)	49 (163.23)	54 (142.57)	37 (6.46)	42 (6.78)	25 (92.83)	28 (197.39)
MaxCSP/503/50/2 (XCSP)	176 (603.56)	398 (386.08)	219 (152.73)	75 (876.84)	249 (76.21)	233 (538.93)	6 (115.39)

Continued on next page

²⁸ Using parameter `-pe parallel_smp 2` on a SUN Grid Engine to ensure half-load of the cores on the cluster.

Problem/s/d/a	DAKOPT	TOULBAR2	CPULEX	CPULEXtuple	MANUS	MANUStuple	CREODE
BlackHole/37/50/2	10 (222.19)	10 (0.08)	30 (141.91)	10 (2.22)	10 (0.30)	10 (2.78)	0 (-)
Coloring/22/6/2	17 (319.29)	17 (11.39)	17 (7.14)	16 (72.33)	14 (17.67)	14 (50.80)	4 (171.61)
Composed/80/10/2	26 (543.73)	80 (0.13)	80 (4.48)	37 (1667.07)	80 (79.81)	73 (1383.72)	0 (-)
EHI/200/7/2	0 (-)	179 (773.86)	0 (-)	0 (-)	1 (3078.96)	0 (-)	0 (-)
Geometric/100/20/2	92 (755.46)	95 (134.57)	65 (419.39)	0 (-)	89 (31.52)	84 (138.98)	0 (-)
Langford/4/29/2	2 (272.24)	2 (0.12)	2 (38.79)	1 (0.03)	2 (0.32)	2 (2.19)	2 (2.97)
QCP/60/9/2	29 (496.31)	15 (143.49)	25 (54.94)	11 (263.83)	53 (121.82)	50 (242.80)	0 (-)
WPMS/427/2/580 (WCNF)	11 (536.35)	197 (110.33)	269 (109.76)	N/A	321 (168.67)	N/A	28 (243.39)
Haplotyping/100/2/580	N/A	1 (784.32)	18 (679.90)	N/A	44 (674.01)	N/A	0 (-)
MIPLib/12/2/93	2 (365.31)	3 (102.39)	3 (49.85)	N/A	3 (9.47)	N/A	3 (28.61)
MaxClique/62/2/2	9 (574.36)	33 (209.07)	38 (229.33)	N/A	40 (362.26)	N/A	24 (280.38)
PackupWeighted/99/2/177	N/A	53 (167.82)	99 (0.72)	N/A	99 (7.14)	N/A	0 (-)
PlanningWithPref/29/2/372	N/A	7 (515.22)	11 (751.65)	N/A	28 (65.82)	N/A	1 (0.03)
TimeTabling/25/2/36	N/A	0 (-)	0 (-)	N/A	7 (1020.73)	N/A	0 (-)
Upgradeability/100/2/77	N/A	100 (12.43)	100 (0.84)	N/A	100 (2.73)	N/A	N/A
CP/35/163/4 (MINIZINC)	9 (387.13)	16 (354.57)	2 (0.99)	7 (584.10)	18 (145.94)	14 (400.03)	26 (138.55)
AMaze/6/17/4	0 (-)	3 (279.71)	0 (-)	4 (998.46)	6 (12.00)	5 (161.25)	4 (176.91)
FastFood/6/5/2	1 (200.32)	1 (0.00)	1 (0.00)	1 (0.00)	1 (0.00)	1 (0.00)	6 (14.22)
Golomb/6/163/3	0 (-)	3 (44.97)	0 (-)	0 (-)	3 (117.34)	1 (78.01)	6 (111.17)
OnCallRostering/5/89/4	1 (253.25)	2 (27.27)	1 (1.98)	2 (47.44)	3 (162.22)	3 (362.19)	2 (75.13)
ParityLearning/7/20/4	7 (432.94)	7 (663.51)	0 (-)	0 (-)	5 (343.24)	4 (907.40)	7 (248.10)
VRP/5/100/4	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	0 (-)	1 (255.45)
Total	1832 (534.34)	2433 (107.26)	1273 (123.70)	1862 (57.35)	1417 (191.45)	1567 (143.45)	202 (219.37)
Nb. of 1st position	0	16 [1]	7 [3]	3 [5]	9 [2]	0	4 [4]
Nb. of best solution	2209 [2]	2562 [1]	1355 [5]	1300 [6]	1626 [4]	1706 [3]	229[7]
Nb. of single best sol.	57 [4]	88 [2]	43 [5]	95 [1]	80 [3]	1 [7]	13 [6]
Zscore (time)	135.37 [6]	57.84 [1]	102.97 [3]	104.89 [4]	90.73 [2]	122.88 [5]	136.58 [7]
Zscore (cost)	63.00 [3]	26.25 [1]	59.24 [2]	69.92 [4]	80.55 [5]	108.76 [7]	100.55 [6]
Borda-score	89.40 [5]	182.50 [1]	129.60 [2]	102.78 [4]	114.37 [3]	59.54 [7]	60.64 [6]
Borda-score (norm)	2.08 [5]	4.24 [1]	3.01 [2]	2.86 [3]	2.66 [4]	1.65 [7]	1.84 [6]

The number of instances solved in less than 1 hour, excluding translation times between formats, is available in Table 3. Resource-based cactus plots are shown in Fig. 1²⁹. Beyond the number of problems solved and the mean CPU time on solved instances reported in this table, we refine our analysis in

²⁹ More detailed results are available at <http://genoweb.toulouse.inra.fr/~degivry/evalgm>.

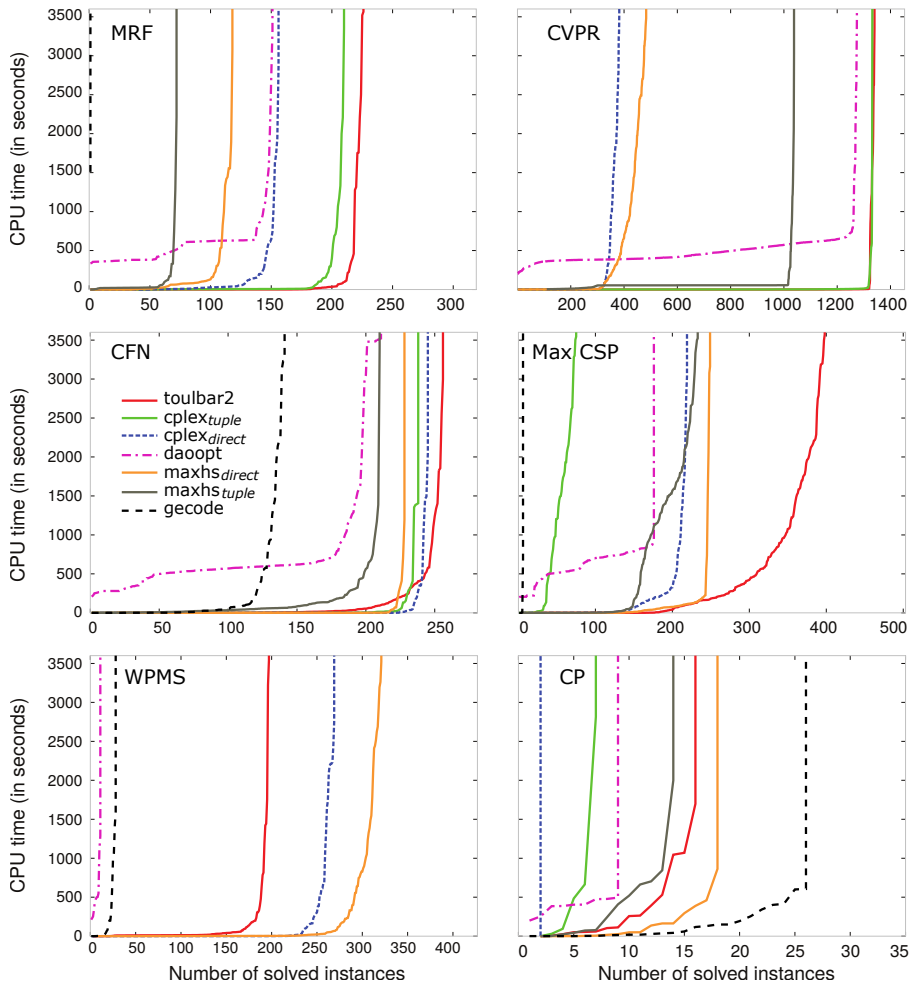


Fig. 1: Cactus plots for MRF, CVPR, CFN, MaxCSP, WPMS, and CP benchmark resources

two ways. First, we summarize the evolution of lower and upper bounds for each algorithm over all instances in Figure 2.

Specifically, for each instance I we normalize all costs as follows: the initial lower bound produced by TOULBAR2 (before VAC) is 0; the best – but potentially suboptimal – solution found by any algorithm is 1; the worst solution is 2. This normalization is invariant to translation and scaling. Additionally, we normalize time from 0 to 1 for each pair of algorithm A and instance I , so that each run finishes at time 1. This time normalization is different for different instances and for different algorithms on the same instance. A point $\langle x, y \rangle$ on the lower bound line for algorithm A in Figure 2 means that after




Fig. 2: Normalized lower and upper bounds on all instances (left) and a set of 1208 hardest instances (right)

normalized runtime x , algorithm A has proved on average over all instances a normalized lower bound of y and similarly for the upper bound. We show both the upper and lower bound curves for all algorithms evaluated here, except GECODE which produces no meaningful lower bound before it proves optimality. In order for the last point of each curve to be visible, we extend all curves horizontally after 1.0. Additionally, on the right of Figure 2, we show the same curves but excluding instances that took less than 5 seconds to solve with a simple version of TOULBAR2 that does not use either VAC preprocessing or hybrid best first search, for a final set of 1208 instances. We remove those easy instances because the runtime tends to be dominated by whatever preprocessing technique each solver uses. This means that the optimal solution is reported near the end of the search, although it is early in absolute terms.

Note that this plot highlights different aspects of the solvers' behavior than the cactus plots and should be interpreted in conjunction with those.

In the second part of our analysis, we compute global measures that try to compensate for the very different cardinalities of the categories. For each instance, we compute two Z-scores³⁰, one for the CPU time and another for the cost of the best solution found at the deadline. In the extreme case where a solver is the only one able to solve an instance (resp. is not able to solve it), we use a score of -4 (resp. 4). A mean Z-score is then computed for each category and the sum of all mean Z-scores is reported in Table 3.

To take into account the CPU time and cost in a common measure, we also computed Borda scores, following the MiniZinc Challenge's approach. For each instance, and each pair of solvers, a reward in $[0, 1]$ is granted to each solver as follows: if a solver reports a better cost than the other, it is granted a reward of 1 (and 0 for the other). For identical costs, if t_0 and t_1 are the CPU time for two solvers denoted 0 and 1, the solver i will receive a reward of $\frac{t_{|i-1|}}{t_0+t_1}$, favoring the fastest solver. A mean Borda score is computed for each category and the sum of mean scores reported.

³⁰ The Z-score of a value x in a set of values is $\frac{x-\mu}{\sigma}$ where μ is the mean of the set and σ its standard deviation.

The tuple encoding and CP approaches are not applicable to WPMS and CVPR, respectively. Quite fairly, these measures penalize these approaches for this limitation. To see if this penalty was enough to explain the scores of these approaches, we also report the Borda score normalized by the number of applicable categories (this optimistically assumes that these approaches would work as well on these inaccessible benchmarks as on the rest of the benchmarks). The only change is a swap of the order between CPLEX with the tuple encoding and MAXHS with the direct encoding.

As expected, for each source of benchmarks, the best solver (in terms of number of solved instances) is usually a solver that is dedicated to this type of problems (*i.e.*, TOULBAR2 for CFN, MAXHS for WPMS, GECODE for CP). However, some solvers, such as MAXHS, CPLEX, and TOULBAR2, performed well on several resources, respectively solving to optimality 2,043, 2,313, and 2,433 instances among a total of 3,026 (using the best encoding on each category for MAXHS and CPLEX). Using the number of solved instances per category, breaking ties by best mean CPU time on solved instances, these three solvers won the first position on 9, 10, and 16 categories respectively, among 43 categories. Looking at cactus plots in Fig. 1, TOULBAR2 and CPLEX dominate on MRF&CVPR, followed by DAOOPT. They also dominate on CFN, followed by MAXHS. TOULBAR2 performed well on MaxCSP. MAXHS and GECODE dominate on their own category (resp. WPMS and CP).

In terms of extreme size and solving difficulty, the CVPR/ColorSeg/colseg-cow4 instance defines the largest search space ($d^n = 2^{829,440}$) completely solved by TOULBAR2. MRF/ObjectDetection is the smallest totally unsolved category ($d^n \leq 2^{264}$). We now consider each benchmark resource, highlighting unexpected results.

[MRF]: on MRF/Linkage [28] (maximum number of variables $n = 1289$, maximum domain size $d = 7$), CPLEX_{tuple} , followed by MAXHS, got the best results, showing their suitability for non-binary (max. arity $a = 5$) cost functions with forbidden tuples. The tuple encoding is the only 01LP encoding usually considered in MRFs. Surprisingly, CPLEX with the direct encoding was the best on the Grid category ($n = 6400, d = 2$), benefiting from a large number of *zero-half cuts*. DAOOPT did not perform as well as for the PIC 2011 Evaluation. One explanation is the missing problem reformulation feature used in the PIC challenge [42], a piece of code which is not available in source or binary format. Another explanation is that DAOOPT spends more time finding good upper bounds (using local search in preprocessing) than on the optimality proof (as Fig. 2 seems to show). CP solvers performed poorly on MRFs due to the large costs, resulting in huge domains for cost variables in the CFN-to-CP translation.

[CVPR]: on CVPR/Scene Decomposition, using a superpixel model [27] with fewer variables ($n = 208, d = 8$), TOULBAR2 solved all 715 instances in 0.07 second each on average compared to 1.11 for CPLEX_t . The good performance of TOULBAR2 on CVPR instances is largely due to its virtual arc consistency initial problem reformulation [11]. On these problems, it offers a tight lower bound in much less time than LP on the tuple encoding. This encoding was

always better for CPLEX, consistent with the ubiquity of the local polytope formulation as a linear relaxation for MRFs. The tuple encoding also improved the performance of the MaxSAT solver (see MAXHS_t results in Table 3) on two categories (ProteinInteraction, SceneDecomp).

[CFN]: TOULBAR2 clearly dominates on CELAR ($n = 458, d = 44$), Pedigree ($n = 10017, d = 28$), and ProteinDesign ($n = 18, d = 198$), whereas CPLEX with direct encoding, followed by MAXHS, performed the best on Operations Research problems Auction ($n = 246, d = 2$) and Warehouse ($n = 1100, d = 300$). The 01LP tuple encoding still performed quite well when the problem size remains relatively small ($n \times d \leq 20,000$), otherwise memory errors sometimes occurred, as on the largest Warehouse instances (*capa-b-c-m*). GECODE performed relatively well on Auction and Warehouse, solving three large instances (*capmo-3-4-5* with $n = 200, d = 100$).

[MaxCSP]: MAXHS performed well on MaxCSP due to its ability to quickly solve all the satisfiable (zero cost optimum) instances that remained in the Geometric ($n = 50, d = 20$) and QCP ($n = 264, d = 9$) categories, thanks to its embedded MINISAT solver. The good results obtained by DAOPT can similarly be explained by its initial stochastic local search procedure [42], finding good initial upper bounds especially on satisfiable or random instances like EHI ($n = 315, d = 7$) and Geometric. TOULBAR2 won the first position on four MaxCSP categories, especially on EHI random category, thanks to its new hybrid best-first search strategy [2] which *simulates* restarts with memory. Surprisingly, the tuple encoding was always dominated by the direct encoding here.

[WPMS]: large clause arities make the tuple encoding or the use of exhaustive tables in UAI format space intractable. While MAXHS dominates the scene, it is interesting to notice the ability of CPLEX to outperform MAXHS on two categories (Upgradeability and PickupWeighted). In PickupWeighted ($n = 25554, d = 2$), CPLEX can be up to one order of magnitude faster than MAXHS. GECODE was the fastest solver to find and prove optimality on 11 MaxClique ($n = 3321, d = 2$) instances, whereas MAXHS won this category by solving 40 instances among 62.

[CP]: CP instances are difficult to translate into GMs with local functions and small domains: 10 instances among 35 MINIZINC instances could not be translated for space reasons. Moreover the translation is not appropriate for LP solvers (linear constraints are decomposed), explaining the poor performance of CPLEX. Here, GECODE performed the best in most of the cases. However, MAXHS, performed the best on two categories: Amaze and OnCallRostering. Similarly, DAOPT was faster than GECODE on the most difficult ParityLearning instance (52_26_6.3). DAOPT solved all the instances in preprocessing thanks to its complete bucket/variable elimination [16], with a memory space usage below 529MB (induced width less than 25), smaller than its limits (4GB and $i = 35$ -bound) [42].

With either the tuple or direct encoding, CPLEX was able to be the best in at least one category per benchmark resource (except for CP) showing very good robustness. For probabilistic models, the tuple encoding is the ideal

choice since the emphasis is on optimization (essentially no tuple with cost k). In this case, the tuple formulation offers a strong bound, an essential source of pruning. In several cases however, thanks to its incremental soft arc consistencies and strong virtual arc consistency preprocessing, TOULBAR2 outperformed CPLEX on such problems. These results can be analyzed in the light of the known relations between LP and soft arc consistencies [11]: thanks to pruning by node consistency and substitutability and to their efficiency and strong incrementality, soft arc consistencies seem capable of outperforming LP by finding a better trade off than LP in the compromise between tightness and computational cost on the local (universal) polytope.

In other benchmarks however, the direct encoding is always preferable. This could be explained by the better conciseness of the encoding on benchmarks with many 0 costs, as shown in Table 2, and to some extent by the lesser pruning of the optimization bound in the presence of hard constraints. This encoding seems essentially ignored in the MRF community.

Overall, this shows that significant speedups can be achieved by exploiting encodings to different optimization languages.

5 Exploitation: a portfolio approach

Solver portfolios [23,21,32] aim to exploit this diversity by replacing a single solver with a set of complimentary solvers and a mechanism for selecting a subset to use on a particular problem. By making decisions at an instance specific level, it is possible to make significant performance gains over any of the individual component solvers. Solver portfolios have been highly successful in constraint programming [41,24,4], satisfiability [51,26], MaxSAT [5], and many more fields. For an extensive survey of the wide-range of literature on the algorithm selection problem, we refer the reader to [32].

The majority of modern portfolio approaches employ some form of machine learning to take the role of the selection model. To enable this involves a training phase whereby for a reference set of instances, a domain-specific feature description, a candidate set of algorithms, and a performance metric are defined. Feature descriptions for each instance and performance data of each algorithm on each instance are recorded. The machine learning model is built such that the performance metric is maximized on this training data. Subsequently, to apply this trained model to a new test instance at runtime, first the feature description must be computed and passed to the model to make a solver selection. The chosen solver is then applied to the problem instance.

5.1 Graphical model instance features

To describe graphical model instances, we consider the following feature set: i) the input file size, ii) the CPU time to read the instance, iii) an initial upper bound on the solution, iv) the time to compute the initial upper bound, v) the

number of variables, vi) the number of cost functions. The ratio of vii) unary, viii) binary, and ix) ternary cost functions, i.e. the fraction of the total number of cost functions of each arity. x) The ratio of cost functions which have arity 4 or greater. Finally, a number of statistics such as the mean, standard deviation, coefficient of variation, minimum, and maximum for xi) domain size, and xii) cost function arity. By no means does this list constitute a comprehensive list of features for graphical models, nevertheless in initial evaluations these proved effective and have the benefit of being relatively cheap to compute.

Table 4 presents the Gini importances [9]³¹ of the above features according to a decision tree classifier aiming to predict the fastest solver. The most important features are the ratio of binary cost functions, the minimum domain size, and the value of the initial upper bound.

Table 4: Gini importances of features

Feature	Gini importance	Feature	Gini importance
Ratio of binary cost functions	0.14445	Arity coefficient of variation	0.07546
Minimum domain size	0.13928	Arity std. deviation	0.07149
Initial UB	0.10988	Arity mean	0.06555
Time to read	0.09211	Num. variables	0.04304
Time UB	0.08393	Num. cost functions	0.03875
File size	0.08305	Mean dom. size	0.01634

5.2 Machine learning offline evaluation results

Table 5 presents an offline evaluation of a simple portfolio approach based on 6 solvers from Sec. 4.3. We consider a subset of the benchmarks and the solvers such that all the instances could be translated to all the solvers, *i.e.*, we exclude the WPMS and CP benchmarks, and the GECODE solver.

The portfolio is built using LLAMA [31], with 10-fold stratified cross validation. This involves splitting the dataset into 10-equally sized folds with an equal distribution of the best solver across folds. For brevity, we present results only for the best performing regression, classification, and clustering methods, plus the Random Forest classifier. The *Virtual Best Solver* (VBS) corresponds to an oracle deciding the best solver for each instance. The table lists the mean (std. dev.) CPU time on the solved instances, the number of instances solved to optimality in less than 1 hour, the number of times each solver was the fastest. In addition, the misclassification penalty shows the contribution of each solver to the portfolio, *i.e.*, the number of instances that were not solved by any other solver, and, where another one solved the instance, the additional CPU time needed by the next best solver. From these statistics alone, it is clear that each of the component solvers (except MAXHS_{tuple}) play

³¹ The normalized total reduction brought by the feature.

Table 5: Summary of portfolio approaches sorted by decreasing number of problems solved over the 2,564 instances

Solver	Solved time (sec.)		Num. solved	Num. best	Misclass. pen.	
	Mean	Std. dev.			solved	total time
VBS(6)	93.0	385.1	2,321			
M5P regression	91.5	376.1	2,298			
J48 classification	84.7	368.1	2,294			
Random Forest	74.6	327.6	2,279			
<i>k</i> -means clustering	66.9	301.4	2,259			
TOULBAR2	105.2	408.3	2,220	1,863	224	28,000.1
CPLEX _{tuple}	55.4	316.6	1,852	27	3	10,345.3
DAOOPT	535.1	340.1	1,812	3	0	3,236.8
MAXHS _{tuple}	140.0	414.5	1,551	3	1	8.4
MAXHS	199.0	565.4	1,078	208	4	9,261.4
CPLEX	127.7	433.4	1,002	217	36	14,381.9

Table 6: Offline evaluation of the UAI 2014 portfolio on 2,564 instances

Solver	Solved time (sec.)		Num. solved	Num. best	Misclass. pen.	
	Mean	Std. dev.			solved	total time
VBS(5)	63.5	276.3	2,315			
UAI'14 portfolio	71.8	312.4	2,276			
INCOPT+TOULBAR2	87.6	361.2	2,227	352	23	82,616.2
TOULBAR2	105.2	408.3	2,220	1,449	13	56,339.3
CPLEX _{tuple}	55.4	316.6	1,852	27	6	9,584.7
MPLP2	66.2	424.6	1,537	198	0	1,183.3
CPLEX	127.7	433.4	1,002	289	46	13,276.0

a valuable contribution to the portfolio both in terms of being able to solve more instances, and reducing the overall CPU time needed. Additionally, each of the portfolio methods are able to outperform the single best solver and close most of the gap to the virtual best solver.

5.3 The UAI 2014 portfolio

A specific portfolio was developed and submitted to the UAI 2014 Inference Competition (MAP task). It was built from five constituent solvers: i) TOULBAR2, ii) a version of TOULBAR2 taking a starting solution from an initial run of the INCOPT [40] local search solver, iii) the Message Passing Linear Programming MPLP2 solver [48, 49], iv) CPLEX using the direct encoding, and v) CPLEX with the tuple encoding. These solvers were selected based on their complementary performances in previous empirical evaluations. Table 6 presents the results of an offline evaluation of this portfolio³².

³² <https://github.com/9thbit/uai-proteus> used a Random Forest classifier and an older version of TOULBAR2 version 0.9.7, with no parameter. Here we report the results using the

The effectiveness of this multi-language portfolio was independently verified in the UAI 2014 Inference Competition, achieving two first places in the MAP task under both the 20 and 60 minute timeouts³³. Three of the portfolio's component solvers were submitted to the same competition as independent entries. The two 01LP encodings performed extremely well on certain instances but extremely poorly on the remaining³⁴. Based on the competition's overall evaluation metric, the cumulative sum of a solver's rank on each instance, the 01LP encodings did not rank high overall but were the top-ranked solvers in a number of cases. Likewise, the INCOP+TOULBAR2 solver was the highest ranked in some cases but ranked in mid-field in many others³⁵. The UAI'14 portfolio solver was highly successful in deciding when to run these solvers or not, achieving first place overall. This independent empirical evaluation supports the findings demonstrated in this paper, that significant speedups can be achieved by exploiting various encodings to related languages.

6 Conclusions

Our empirical results demonstrate the effectiveness of a number of solvers on various graphical model formats, where no single solver consistently dominates the results. Rather, the best solver depends on each problem category, bringing to light the respective strengths, robustness and weaknesses of each solver family. They highlight the efficacy of encoding a problem to a related language and exploiting complementary solving technologies.

We demonstrate that it is possible to exploit these complementary strengths using a portfolio approach, built on this knowledge won the UAI 2014 Evaluation. We hope that our proposed collection of benchmarks, readily available in many formats, will enrich the various competitions in CP, AI, and OR, leading to more robust solvers and new solving strategies.

Acknowledgements We are grateful to the GenoToul (Toulouse) Bioinformatic platform for providing us computational support for this work. This work is supported by Science Foundation Ireland (SFI) Grant 10/IN.1/I3032. The Insight Centre for Data Analytics is supported by SFI Grant SFI/12/RC/2289.

References

1. Allouche, D., Bessiere, C., Boizumault, P., Givry, S., Gutierrez, P., Loudni, S., Métiévier, J., Schiex, T.: Decomposing global cost functions. In: Proc. of AAAI (2012)

same settings as in Sec. 4.3, INCOP+TOULBAR2 corresponds to TOULBAR2 using an extra parameter *-i* for the initial INCOP starting solution phase.

³³ See MAP/Proteus entry at <http://www.hlt.utdallas.edu/~vgogate/uai14-competition/leaders.html>.

³⁴ See MAP/MIP-UAI and MAP/MIP-T-UAI entries.

³⁵ See MAP/IncTb entry.

2. Allouche, D., de Givry, S., Katsirelos, G., Schiex, T., Zytnicki, M.: Anytime Hybrid Best-First Search with Tree Decomposition for Weighted CSP. In: Proc. of CP, pp. 12–28 (2015)
3. Allouche, D., Traoré, S., André, I., Givry, S., Katsirelos, G., Barbe, S., Schiex, T.: Computational protein design as a cost function network optimization problem. In: Proc. of CP, pp. 840–849 (2012)
4. Amadini, R., Gabbrielli, M., Mauro, J.: A Multicore Tool for Constraint Solving. In: Proc. of IJCAI, pp. 232–238 (2015)
5. Ansótegui, C., Malitsky, Y., Sellmann, M.: MaxSAT by Improved Instance-Specific Algorithm Configuration. In: Proc. of AAAI, pp. 2594–2600 (2014)
6. Argelich, J., Cabiscol, A., Lynce, I., Manyà, F.: Encoding Max-CSP into partial MaxSAT. In: Proc. of ISMVL, pp. 106–111 (2008)
7. Bacchus, F.: GAC via unit propagation. In: Proc. of CP, pp. 133–147 (2007)
8. Bensana, E., Lemaître, M., Verfaillie, G.: Earth observation satellite management. Constraints **4**(3), 293–299 (1999)
9. Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A.: Classification and regression trees. CRC press (1984)
10. Cabon, B., de Givry, S., Lobjois, L., Schiex, T., Warners, J.: Radio link frequency assignment. Constraints **4**, 79–89 (1999)
11. Cooper, M., de Givry, S., Sanchez, M., Schiex, T., Zytnicki, M., Werner, T.: Soft arc consistency revisited. Artificial Intelligence **174**, 449–478 (2010)
12. Cooper, M., de Givry, S., Schiex, T.: Optimal soft arc consistency. In: Proc. of IJCAI, pp. 68–73 (2007)
13. Cooper, M.C., Schiex, T.: Arc consistency for soft constraints. Artificial Intelligence **154**(1-2), 199–227 (2004)
14. Davies, J., Bacchus, F.: Solving MAXSAT by solving a sequence of simpler SAT instances. In: Proc. of CP, pp. 225–239 (2011)
15. Davies, J., Bacchus, F.: Exploiting the power of MIP solvers in MaxSAT. In: Proc. of SAT, pp. 166–181 (2013)
16. Dechter, R.: Bucket elimination: A unifying framework for reasoning. Artificial Intelligence **113**(1–2), 41–85 (1999)
17. Fargier, H., Lang, J., Martin-Clouaire, R., Schiex, T.: A constraint satisfaction framework for decision under uncertainty. In: Proc. of the 11th Int. Conf. on Uncertainty in Artificial Intelligence. Montréal, Canada (1995)
18. Favier, A., Givry, S., Legarra, A., Schiex, T.: Pairwise decomposition for combinatorial optim. in graphical models. In: Proc. of IJCAI, pp. 2126–2132 (2011)
19. de Givry, S., Prestwich, S., O’Sullivan, B.: Dead-end elimination for weighted CSP. In: Proc. of CP, pp. 263–272 (2013)
20. Globerson, A., Jaakkola, T.: Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations. In: Proc. of NIPS, pp. 553–560 (2007)
21. Gomes, C.P., Selman, B.: Algorithm Portfolios. Artificial Intelligence **126**(1-2), 43–62 (2001)
22. Hebrard, E., O’Mahony, E., O’Sullivan, B.: Constraint Programming and Combinatorial Optimisation in Numberjack. In: Proc. of CP-AI-OR, pp. 181–185 (2010)
23. Huberman, B.A., Lukose, R.M., Hogg, T.: An Economics Approach to Hard Computational Problems. Science **275**(5296), 51–54 (1997)
24. Hurley, B., Kotthoff, L., Malitsky, Y., O’Sullivan, B.: Proteus: A Hierarchical Portfolio of Solvers and Transformations. In: Proc. of CP-AI-OR, pp. 301–317 (2014)
25. Jünger, M., Lieblich, T., Naddef, D., Nemhauser, G., Pulleyblank, W., Reinelt, G., Rinaldi, G., Wolsey, L. (eds.): 50 Years of Integer Programming 1958-2008. Springer (2010)
26. Kadioglu, S., Malitsky, Y., Sellmann, M., Tierney, K.: ISAC – Instance-Specific Algorithm Configuration. In: Proc. of ECAI, pp. 751–756 (2010)
27. Kappes, J., Andres, B., Hamprecht, F., Schnörr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B., Kröger, T., Lellmann, J., Komodakis, N., Savchynskyy, B., Rother, C.: A comparative study of modern inference techniques for structured discrete energy minimization problems. International Journal of Computer Vision **115**(2), 155–184 (2015)

28. Kishimoto, A., Marinescu, R.: Recursive best-first and/or search with overestimation for genetic linkage analysis. In: Proc. of CP Workshop on Constraint Based Methods for Bioinformatics (2013)
29. Koller, D., Friedman, N.: Probabilistic graphical models: principles and techniques. The MIT Press (2009)
30. Koster, A.: Frequency assignment: Models and algorithms. Ph.D. thesis (1999)
31. Kotthoff, L.: LLAMA: leveraging learning to automatically manage algorithms. Tech. Rep. arXiv:1306.1031, arXiv (2013). URL <http://arxiv.org/abs/1306.1031>
32. Kotthoff, L.: Algorithm Selection for Combinatorial Search Problems: A Survey. AI Magazine **35**(3), 48–60 (2014)
33. Kratica, J., Tošić, D., Filipović, V., Ljubić, I.: Solving the simple plant location problem by genetic alg. RAIRO **35**(1), 127–142 (2001)
34. Larrosa, J., de Givry, S., Heras, F., Zytnicki, M.: Existential arc consistency: getting closer to full arc consistency in weighted CSPs. In: Proc. of IJCAI, pp. 84–89 (2005)
35. Larrosa, J., Heras, F., de Givry, S.: A logical approach to efficient max-sat solving. Artif. Intell. **172**(2-3), 204–233 (2008)
36. Li, C.M., Manyà, F.: Maxsat. In: Handbook of satisfiability, chap. 19. IOS Press (2009)
37. Meltzer, T., Globerson, A., Weiss, Y.: Convergent message passing algorithms: a unifying view. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, pp. 393–401. AUAI Press (2009)
38. Meseguer, P., Rossi, F., Schiex, T.: Soft constraints processing. In: F. Rossi, P. van Beek, T. Walsh (eds.) Handbook of Constraint Programming, chap. 9. Elsevier (2006)
39. Nethercote, N., Stuckey, P., Becket, R., Brand, S., Duck, G., Tack, G.: MiniZinc: Towards a standard CP modelling language. In: Proc. of CP, pp. 529–543 (2007)
40. Neveu, B., Trombettoni, G., Glover, F.: Id walk: A candidate list strategy with a simple diversification device. In: Proc. of CP, pp. 423–437 (2004)
41. O’Mahony, E., Hebrard, E., Holland, A., Nugent, C., O’Sullivan, B.: Using Case-based Reasoning in an Algorithm Portfolio for Constraint Solving. Irish Conference on Artificial Intelligence and Cognitive Science (2008)
42. Otten, L., Ihler, A., Kask, K., Dechter, R.: Winning the PASCAL 2011 MAP challenge with enhanced AND/OR branch-and-bound. In: NIPS DISCML Workshop (2012)
43. Petit, T., Régim, J., Bessière, C.: Meta constraints on violations for over constrained problems. In: Proc. of ICTAI, pp. 358–365 (2000)
44. Prusa, D., Werner, T.: Universality of the local marginal polytope. Pattern Analysis and Machine Intelligence, IEEE Transactions on **37**(4), 898–904 (2015)
45. Rossi, F., van Beek, P., Walsh, T. (eds.): Handbook of Constraint Programming. Elsevier (2006)
46. Sánchez, M., de Givry, S., Schiex, T.: Mendelian error detection in complex pedigrees using weighted constraint satisfaction techniques. Constraints **13**(1-2), 130–154 (2008)
47. Schlesinger, M.: Syntactic analysis of two-dimensional visual signals in noisy conditions. Kibernetika **4**, 113–130 (1976)
48. Sontag, D., Choe, D., Li, Y.: Efficiently searching for frustrated cycles in MAP inference. In: Proc. of UAI, pp. 795–804 (2012)
49. Sontag, D., Meltzer, T., Globerson, A., Weiss, Y., Jaakkola, T.: Tightening LP relaxations for MAP using message-passing. In: Proc. of UAI, pp. 503–510 (2008)
50. Werner, T.: A linear programming approach to max-sum problem. Pattern Analysis and Machine Intelligence **29**(7), 1165–1179 (2007)
51. Xu, L., Hutter, F., Hoos, H.H., Leyton-Brown, K.: SATzilla: Portfolio-based Algorithm Selection for SAT. In: Journal of Artificial Intelligence Research, pp. 565–606 (2008)