

Lifted Variable Elimination for Probabilistic Logic Programming

ELENA BELLODI¹, EVELINA LAMMA¹, FABRIZIO RIGUZZI²

VITOR SANTOS COSTA³, RICCARDO ZESE¹

¹ *Dipartimento di Ingegneria – Università di Ferrara
Via Saragat 1, 44122, Ferrara, Italy*

² *Dipartimento di Matematica e Informatica – Università di Ferrara
Via Saragat 1, 44122, Ferrara, Italy*

³ *CRACS and DCC-FCUP – Universidade do Porto
Rua do Campo Alegre, 1021/1055, 4169-007 Porto, Portugal
(e-mail: name.surname@unife.it, vsc@dcc.fc.up.pt)*

Abstract

Lifted inference has been proposed for various probabilistic logical frameworks in order to compute the probability of queries in a time that depends on the size of the domains of the random variables rather than the number of instances. Even if various authors have underlined its importance for probabilistic logic programming (PLP), lifted inference has been applied up to now only to relational languages outside of logic programming. In this paper we adapt Generalized Counting First Order Variable Elimination (GC-FOVE) to the problem of computing the probability of queries to probabilistic logic programs under the distribution semantics. In particular, we extend the Prolog Factor Language (PFL) to include two new types of factors that are needed for representing ProbLog programs. These factors take into account the existing causal independence relationships among random variables and are managed by the extension to variable elimination proposed by Zhang and Poole for dealing with convergent variables and heterogeneous factors. Two new operators are added to GC-FOVE for treating heterogeneous factors. The resulting algorithm, called LP² for Lifted Probabilistic Logic Programming, has been implemented by modifying the PFL implementation of GC-FOVE and tested on three benchmarks for lifted inference. A comparison with PITA and ProbLog2 shows the potential of the approach.

KEYWORDS: Probabilistic Logic Programming, Lifted Inference, Variable Elimination, Distribution Semantics, ProbLog, Statistical Relational Artificial Intelligence

1 Introduction

Over the last years, there has been increasing interest in models that combine first-order logic and probability, both for domain modeling under uncertainty, and for efficiently performing inference and learning (Getoor and Taskar 2007; De Raedt et al. 2008). Probabilistic Logic Programming (PLP) has recently received an increasing attention for its ability to incorporate probability in logic programming. Among the various proposals, the one based on the distribution semantics (Sato 1995) has gained popularity as the basis of languages such as Probabilistic Horn Abduction (Poole 1993), PRISM (Sato

1995), Independent Choice Logic (Poole 1997), Logic Programs with Annotated Disjunctions (Vennekens et al. 2004), and ProbLog (De Raedt et al. 2007).

Nonetheless, research in Probabilistic Logic Languages has made it very clear that it is crucial to design models that can support efficient inference, while preserving intensional, and declarative modeling. Lifted inference (Poole 2003; de Salvo Braz et al. 2005; Milch et al. 2008; Van den Broeck et al. 2011) is one of the major advances in this respect. The idea is to take advantage of the regularities in structured models to decrease the number of operations. Originally, the idea was proposed as an extension of variable elimination (VE for short in the following). Work on lifting VE started with (Poole 2003). Lifted VE exploits the symmetries present in first-order probabilistic models, so that it can apply the same principles behind VE to solve a probabilistic query without grounding the model.

Most work in probabilistic inference compute statistics from a sum of products representation, where each element is named a *factor*. Lifted inference generates templates, named *parametric factors* or *parfactors*, which stand for a set of similar factors found in the inference process, thus delaying as much as possible the use of fully instantiated factors.

In (Gomes and Costa 2012), the *Prolog Factor Language* (PFL, for short) was proposed as a Prolog extension to support probabilistic reasoning with parfactors. PFL exploits the state-of-art algorithm GC-FOVE (Taghipour et al. 2013), which redefines the operations described in (Milch et al. 2008) to be correct for whatever constraint representation is being used. This decoupling of the lifted inference algorithm from the constraint representation mechanism allows any constraint language that is closed under these operators to be plugged into the algorithm to obtain an inference system. In fact, the lifted VE algorithm of (Gomes and Costa 2012) represents the adaptation of GC-FOVE to the PFL constraints.

In this work, we move further towards exploiting efficient inference via lifted VE for PLP Languages under the distribution semantics. To support reasoning compliant with the distribution semantics, we introduce two novel operators (named *heterogeneous* lifted multiplication and sum) in the PFL, and modify the GC-FOVE algorithm for computing them. We name LP² (for Lifted Probabilistic Logic Programming) the resulting system. An experimental comparison between LP² and ProbLog2 (Fierens et al. 2014) and PITA (Riguzzi and Swift 2011) shows that inference time increases linearly with the number of individuals of the program domain for LP², rather than exponentially as with ProbLog2 and PITA.

This is an exciting development towards the goal of preserving the declarativeness and conciseness of Probabilistic Logic Languages, while extremely gaining in performances.

The paper is organized as follows. Section 2 introduces preliminaries regarding ProbLog, PFL, Causal Independence Variable Elimination, and GC-FOVE. Section 3 discusses the translation of ProbLog into the extended PFL. Section 4 presents the new operators introduced in GC-FOVE. Section 5 reports the experiments performed and Section 6 concludes the paper.

2 Preliminaries

2.1 ProbLog

ProbLog (De Raedt et al. 2007) is a Probabilistic Logic Programming (PLP) language. A ProbLog program consists of a set of *ground probabilistic facts* plus a definite logic program, i.e. a set of rules. A ground probabilistic fact, written $p :: f$, is a ground fact f annotated with a number p such that $0 \leq p \leq 1$. An atom that unifies with a ground probabilistic fact is called a *probabilistic atom*, while an atom that unifies with the head of some rule in the logic program is called a *derived atom*.

If a set of probabilistic facts has the same probability p , it can be defined intensionally through the syntax $p :: f(X_1, X_2, \dots, X_n) : -B$, where $f(X_1, X_2, \dots, X_n)$ is the signature of the set, and B is a conjunction of non-probabilistic goals, as shown in Example 1. Such rules are range-restricted: all variables in the head of a rule should also appear in a positive literal in the body.

Example 1 (Running example)

Here we present an example inspired by the *workshop attributes* problem of (Milch et al. 2008). The ProbLog program models the scenario in which a workshop is being organized and a number of people have been invited. **series** indicates whether the workshop is successful enough to start a series of related meetings while **attends**(P) indicates whether person P will attend the workshop.

```
series :- s.
series :- attends(P).
attends(P) :- at(P,A).
0.1::s.
0.3::at(P,A) :- person(P), attribute(A).
```

The first two rules define when the workshop becomes a series: either because of its own merits or because people attend. The third rule states that whether a person attends the workshop depends on its attributes (location, date, fame of the organizers, etc).

The probabilistic fact **s** represents the merit of the workshop. The probabilistic fact **at**(P,A) represents whether person P attends because of attribute A. Notice that the last statement corresponds to a set of ground probabilistic facts, one for each person P and attribute A. For brevity we do not show the (non-probabilistic) facts describing **person**/1 and **attribute**/1 predicates.

A ProbLog program specifies a probability distribution over normal logic programs. In this work, we consider the semantics in the case of no function symbols to be restricted to finite programs, and assume all worlds have a *two-valued* well-founded model.

For each ground probabilistic fact $p_i :: f_i$, an *atomic choice* specifies whether to include f in a world (with probability p_i) or not (with probability $1 - p_i$). A *total choice* C is a set of atomic choices, one for each ground probabilistic fact. These choices are assumed to be independent, hence the probability of a total choice is the product of the probabilities of the individual atomic choices, $P(C) = \prod_i p_i$. A total choice C also identifies a normal logic program $W = F \cup R$ called a *world*, where F is the set of facts to be included according to C and R denotes the rules in the ProbLog program. Let $\mathcal{W}$ be the set of all possible worlds. The probability of a world is equal to the probability of its total choice. The conditional probability of a query (a ground atom) Q given a world W $P(Q|W)$ is 1 if the Q is true in the well-founded model of W and 0 otherwise. The probability of a query can therefore be obtained as $P(Q) = \sum_{W \in \mathcal{W}} P(Q, W) = \sum_{W \in \mathcal{W}} P(Q|W)P(W) = \sum_{W \in \mathcal{W}: W \models Q} P(W)$.

2.2 The Prolog Factor Language

Most graphical models provide a concise representation of a joint distribution by encoding it as a set of factors. The probability of a set of variables $\mathbf{X}$ taking the value $\mathbf{x}$ can be expressed as product of n factors if:

$$P(\mathbf{X} = \mathbf{x}) = \frac{\prod_{i=1,\dots,n} \phi_i(\mathbf{x}_i)}{Z}$$

where $\mathbf{x}_i$ is a sub-vector of $\mathbf{x}$ that depends on the i -th factor and Z is a normalization constant (i.e. $Z = \sum_{\mathbf{x}} \prod_{i=1,\dots,n} \phi_i(\mathbf{x}_i)$). Bayesian networks are an example where there is a factor for each variable that is a function of the variable X_i and its parents $X_j \dots X_k$, such that $\phi(X_i, X_j \dots X_k) = P(X_i | X_j \dots X_k)$ and $Z = 1$. As progress has been made on managing large networks, it has become clear that often the same factor appears repeatedly in the network, thus suggesting the use of templates generalizing individual factors, or *parametric factors* (Kisynski and Poole 2009a).

The Prolog Factor Language (PFL) (Gomes and Costa 2012) extends Prolog to support probabilistic reasoning with parametric factors or *parfactors*. The PFL syntax for a factor is *Type F ; ϕ ; C*. *Type* refers to the type of the network over which the parfactor is defined (*bayes* for directed networks or *markov* for undirected ones); *F* is a sequence of Prolog terms that define sets of random variables under the constraints in *C*. The set of all logical variables in *F* is named *L*. *C* is a list of Prolog goals that impose bindings on the logical variables in *L* (the successful substitutions for the goals in *C* are the valid values for the variables in *L*). ϕ is the table defining the factor in the form of a list of real values. By default all random variables are boolean but a different domain may be defined. An example of a factor is `series,attends(P);[0.51,0.49,0.49,0.51];[person(P)]`: it has the Boolean random variables `series` and `attends(P)` as arguments, `[0.51,0.49,0.49,0.51]` as table and `[person(P)]` as constraints. The semantics of a PFL program is given by the set of factors obtained by grounding parfactors: each parfactor stands for the set of its grounding obtained by replacing variables of *L* with the values allowed by the constraints in *C*. The set of ground factors define a factorization of the joint probability distribution over all random variables.

Example 2 (PFL Program)

A version of the *workshop attributes* problem presented in Example 1 can be modeled by a PFL program such as

```
bayes attends(P), at(P,A) ; [0.7, 0.3, 0.3, 0.7] ; [person(P),attribute(A)].
bayes series, attends(P) ; [0.51, 0.49, 0.49, 0.51] ; [person(P)].
```

2.3 Variable Elimination and Causal Independence

Quite often we want to find out the probability distribution of a set of random variables $\mathbf{X}$ given that we know the values, or have evidence $\mathbf{y}$, on a set of variables $\mathbf{Y}$, where $\mathbf{X}$ is often a single variable X . Variable Elimination (VE) (Zhang and Poole 1996) is an algorithm for computing this *posterior* probability in factorized joint probability distributions. The key idea is to eliminate the random variables from a set of factors one by one until only the query variable X remains. To do so VE eliminates a variable V by first multiplying all

the factors that include V into a single factor; V can then be discarded through summing it out from the newly constructed factor.

More formally, suppose $\phi_1(X_1 \dots X_i, Y_1 \dots Y_j)$ and $\phi_2(Y_1 \dots Y_j, Z_1 \dots Z_k)$ are factors. The product $(\phi_1 \times \phi_2)(x_1 \dots x_i, y_1 \dots y_j, z_1 \dots z_k)$ is simply $\phi_1(x_1 \dots x_i, y_1 \dots y_j) \times \phi_2(y_1 \dots y_j, z_1 \dots z_k)$ for every value of $x_1 \dots x_i, y_1 \dots y_j, z_1 \dots z_k$. To eliminate a variable X_1 from the factors $\phi(x_1 \dots x_i)$ one observes that the cases for X_1 are mutually exclusive, thus $\phi'(x_2 \dots x_i)$ is simply $(\sum_{x_1} \phi)(x_2 \dots x_i) = \phi(\alpha_1, x_2 \dots x_i) + \dots + \phi(\alpha_m, x_2 \dots x_i)$, where $\alpha_1, \dots, \alpha_m$ are the possible values of X_1 .

The full VE algorithm takes as input a set of factors $\mathcal{F}$, an elimination order ρ , a set of query variables $\mathbf{X}$ and a list $\mathbf{y}$ of observed values. First, it sets the observed variables in all factors to their corresponding observed values. Then it repeatedly selects the first variable Z from the elimination order ρ and it calls SUM-OUT on $\mathcal{F}$ and Z , until ρ becomes empty. In the final step, it multiplies together the factors of $\mathcal{F}$ obtaining a new factor γ that is normalized as $\gamma(x)/\sum_{x'} \gamma(x')$ to give the posterior probability.

Noisy OR-Gates Bayesian networks take advantage of conditional independence between variables to reduce the size of the representation. *Causal independence* (Zhang and Poole 1996) goes one step further and looks at independence conditioned on *values* of the random variables. One important example is the *noisy OR-gate*, where we have a Boolean variable X with parents $\mathbf{Y}$, and ideally X should be true if any of the Y_i is true. In practice, each parent Y_i has a noisy inhibitor that independently blocks or activates Y_i , so X is true if either **any** of the causes Y_i holds true *and* is not inhibited.

A noisy OR can be expressed as a factor ϕ . In fact, it can be also expressed as a combination of factors by introducing intermediate variables that represent the effect of each cause *given the inhibitor*. For example, if X has two causes Y_1 and Y_2 , we can introduce a variable X' to account for the effect of Y_1 and X'' for Y_2 , and the factor $\phi(Y_1, Y_2, X)$ can be expressed as

$$\phi(y_1, y_2, x) = \sum_{x' \vee x'' = x} \psi(y_1, x') \gamma(y_2, x'') \quad (1)$$

where the summation is over all values x' and x'' of X' and X'' whose disjunction is equal to x . The X variable is called *convergent* as it is where independent contributions from different sources are collected and combined. Non-convergent variables will be called *regular variables*.

The noisy OR thus allows for a $O(n)$ representation of a conditional probability table with n parents. Unfortunately, straightforward use of VE for inference would lead to construct $O(2^n)$ tables. A modified algorithm, called VE1 (Zhang and Poole 1996), combines factors through a new operator $\otimes$, that generalizes formula (1) as follows. Let ϕ and ψ be two factors that share convergent variables $E_1 \dots E_k$, let $\mathbf{A}$ be the list of regular variables that appear in both ϕ and ψ , let $\mathbf{B}_1$ ($\mathbf{B}_2$) be the list of variables appearing only in ϕ (ψ). The combination $\phi \otimes \psi$ is given by

$$\begin{aligned} \phi \otimes \psi(E_1 = \alpha_1, \dots, E_k = \alpha_k, \mathbf{A}, \mathbf{B}_1, \mathbf{B}_2) = \\ \sum_{\alpha_{11} \vee \alpha_{12} = \alpha_1} \dots \sum_{\alpha_{k1} \vee \alpha_{k2} = \alpha_k} \phi(E_1 = \alpha_{11}, \dots, E_k = \alpha_{k1}, \mathbf{A}, \mathbf{B}_1) \psi(E_1 = \alpha_{12}, \dots, E_k = \alpha_{k2}, \mathbf{A}, \mathbf{B}_2) \end{aligned} \quad (2)$$

Factors containing convergent variables are called *heterogeneous* while the remaining factors are called *homogeneous*. Heterogeneous factors sharing convergent variables must be combined with $\otimes$ that we call *heterogeneous multiplication*.

Algorithm VE1 exploits causal independence by keeping two lists of factors instead of one: a list of homogeneous factors $\mathcal{F}_1$ and a list of heterogeneous factors $\mathcal{F}_2$. Procedure SUM-OUT is replaced by SUM-OUT1 that takes as input $\mathcal{F}_1$ and $\mathcal{F}_2$ and a variable Z to be eliminated. First, all the factors containing Z are removed from $\mathcal{F}_1$ and combined with multiplication to obtain factor ϕ . Then all the factors containing Z are removed from $\mathcal{F}_2$ and combined with heterogeneous multiplication obtaining ψ . If there are no such factors set $\psi = \text{nil}$. In the latter case, SUM-OUT1 adds the new (homogeneous) factor $\sum_z \phi$ to $\mathcal{F}_1$ otherwise it adds the new (heterogeneous) factor $\sum_z \phi\psi$ to $\mathcal{F}_2$. Procedure VE1 is the same as VE with SUM-OUT replaced by SUM-OUT1 and with the difference that two sets of factors are maintained instead of one.

The $\otimes$ operator assumes that the convergent variables are independent given the regular variables. This can be ensured by *deputising* the convergent variables: every such variable E is replaced by a new convergent variable E' (called a *deputy variable*), E' that replaces E in the heterogeneous factors containing E , E becomes a regular variable, and a new factor $\iota(E, E')$ is introduced, called *deputy factor*, that represents the identity function between E and E' , i.e., it is defined by

$\iota(E, E')$	ff	ft	tf	tt
	1.0	0.0	0.0	1.0

Deputising ensures that we do not have descendants of a convergent variable in an heterogeneous factor as long as the elimination order for VE1 is such that $\rho(E') < \rho(E)$.

2.4 GC-FOVE

Work on lifting VE started with (Poole 2003), and the current state of the art is the algorithm GC-FOVE (Taghipour et al. 2013), which redefines the operations of C-FOVE (Milch et al. 2008). The lifted VE algorithm of (Gomes and Costa 2012) represents the adaptation of GC-FOVE to the PFL language.

First-order Variable Elimination (FOVE) (Poole 2003; de Salvo Braz et al. 2005) computes the marginal probability distribution for query random variables (randvars) by repeatedly applying operators that are lifted counterparts of VE's operators. Models are in the form of a set of parfactors that are essentially the same as in PFL. A parametrized random variable (PRV) $\mathcal{V}$ is of the form $A|C$, where $A = F(X_1, \dots, X_n)$ is a non-ground atom and C is a constraint on logical variables (logvars) $\mathbf{X} = \{X_1, \dots, X_n\}$. Each PRV represents the set of randvars $\{F(\mathbf{x}) | \mathbf{x} \in C\}$, where $\mathbf{x}$ is the tuple of constants $(x_1, \dots, x_n)$. Given a PRV $\mathcal{V}$, we use $RV(\mathcal{V})$ to denote the set of randvars it represents. Each ground atom is associated with one randvar, which can take any value in $range(F)$.

GC-FOVE tries to eliminate all (non-query) PRVs in a particular order. To do so, GC-FOVE supports several operators. It first tries *Lifted Sum-Out*, that excludes a PRV from a parfactor ϕ if the PRV only occurs in ϕ . Next, *Lifted Multiplication*, that multiplies two aligned parfactors. Matching variables must be properly aligned and the new coefficients must be computed taking into account the number of groundings in C . Third, *Lifted Absorption* eliminates n PRVs that have the same observed value. If the two operations cannot be applied, a chosen parfactor must be *split* so that some of its PRVs match another parfactor. In the worst case, when none of the lifted operators can be applied, GC-FOVE resorts to propositionalization: it completely grounds the parametrized randvars and parfactors and performs inference on the ground level.

GC-FOVE further extends PRVs with counting formulas, introduced in C-FOVE (Milch et al. 2008). A counting formula takes advantage of symmetry existing in factors that are products of independent variables. It represents a factor of the form $\phi(F(x_1), F(x_2), \dots, F(x_n))$, where all variables have the same domain, as $\phi(\#_X[F(X)])$. The factor implements a multinomial distribution, such that its values depend on the number of variables n and domain size. The lifted counted variable is named a PCR. PCRVs may result from summing-out, when we obtain factors with a single PRV, or through *Counting Conversion* that searches for factors of the form $\phi(\prod_i (S(X_j)F(x_j, y_i)))$ and counts on the occurrences of Y . Definitions for counting formulas are reported in Appendix B.

GC-FOVE employs a constraint-tree to represent arbitrary constraints C , whereas the PFL simply uses sets of tuples. Arbitrary constraints can capture more symmetries in the data, which potentially offers the ability to perform more operations at a lifted level.

3 Translating ProbLog into PFL

In order to translate ProbLog into PFL, let us start from the conversion of a ProbLog program into a Bayesian network with noisy OR nodes. Here we adapt the conversion for Logic Programs with Annotated Disjunctions presented in (Vennekens et al. 2004; Meert et al. 2008) to the case of ProbLog. The first step is to generate the grounding of the ProbLog program. For each atom A in the Herbrand base of the program, the Bayesian network contains a Boolean random variable with the same name. Each probabilistic fact $p :: A$ is represented by a parentless node with the conditional probability table (CPT):

A	f	t
	1-p	p

For each ground rule $R_i = H \leftarrow B_1, \dots, B_n, \text{not}(C_1), \dots, \text{not}(C_m)$ we add to the network a random variable called H_i that has as parents $B_1, \dots, B_n, C_1, \dots, C_m$ and the following CPT:

H_i	$B_1 = t, \dots, B_n = t, C_1 = f, \dots, C_m = f$	all other columns
f	0.0	1.0
t	1.0	0.0

In practice H_i is the result of the conjunction of random variables representing the atoms in the body. Then for each ground atom H in the Herbrand base not appearing in a probabilistic fact, we add H to the network with parents all H_i of ground rules with H in the head and with the CPT

H	at least one $H_i = t$	all other columns
f	0.0	1.0
t	1.0	0.0

representing the result of the disjunction of random variables H_i .

Translating ProbLog into PFL allows us to stay in the lifted (non-ground) program.

Example 3 (Translation of a ProbLog program into PFL)

The translation of the ProbLog program of Example 1 into PFL is

```

bayes series1, s; identity ; [].
bayes series2, attends(P); identity; [person(P)].
bayes series, series1, series2; disjunction; [].
bayes attends1(P), at(P,A); identity; [person(P),attribute(A)].
bayes attends(P), attends1(P); identity; [person(P)].
bayes s; [0.9, 0.1]; [].
bayes at(P,A); [0.7, 0.3] ; [person(P),attribute(A)].

identity([1,0,0,1]).

```

```
disjunction([1,0,0,0,
            0,1,1,1]).
```

Notice that **series2** and **attends1(P)** can be seen as or-nodes. Thus, after grounding, factors derived from the second and the fourth parfactor should not be multiplied together but should be combined with heterogeneous multiplication, as variables **series2** and **attends1(P)** are in fact convergent variables. To do so, we need to identify heterogeneous factors and add deputy variables and factors. We thus introduce two new types of factors to PFL, **het** and **deputy**. The first factor is such that its ground instantiations are heterogeneous factors. The convergent variables are assumed to be represented by the first atom in the factor's list of atoms. Lifting identity is straightforward, it corresponds to two atoms and imposes an identity factor between their ground instantiations. Since the factor is fixed, it is not indicated.

Example 4 (Extended PFL program)

The PFL program of Example 3, extended with the two new factors **het** and **deputy**, becomes:

```
het series1p, s; identity ; [].
het series2p, attends(P); identity; [person(P)].
deputy series2, series2p; [].
deputy series1, series1p; [].
bayes series, series1, series2; disjunction ; [].
het attends1p(P), at(P,A); identity; [person(P),attribute(A)].
deputy attends1(P), attends1p(P); [person(P)].
bayes attends(P), attends1(P); identity; [person(P)].
bayes s; [0.9, 0.1]; [].
bayes at(P,A); [0.7, 0.3] ; [person(P),attribute(A)].
```

where **series1p**, **series2p** and **attends1p(P)** are the convergent deputy random variables, and **series1**, **series2** and **attends1(P)** are their corresponding new regular variables. The fifth Bayesian factor represents the combination of the contribution to **series** of the two rules for it. Causal independence could be applied here as well since the combination is really an OR, but for simplicity we decided to concentrate only on exploiting causal independence for the convergent variables represented by the head of rules which is the hard part.

4 Heterogeneous Lifted Multiplication and Summation

GC-FOVE must be modified in order to take into account heterogeneous factors and convergent variables. The VE algorithm must be replaced by VE1, i.e., two lists of factors must be maintained, one with homogeneous and the other with heterogeneous factors. When summing out a variable, first the homogeneous factors must be combined together with homogeneous lifted multiplication. Then the heterogeneous factors must be combined together with heterogeneous lifted multiplication and, finally, the two results must be combined to produce a final factor from which the random variable is eliminated.

Lifted heterogeneous multiplication is defined as Operator 1, considering the case in which the two factors share convergent random variables. We assume familiarity with set and relational algebra (e.g., join $\bowtie$) while some useful definitions are reported in Appendix B. PRVs must be *count-normalized*, that is, the corresponding parameters must be scaled to take into account domain size and number of occurrences in the parfactor. PRVs are then aligned and the joint domain is computed as the natural join between the

Operator HET-MULTIPLY**Inputs:**

- (1) $g_1 = \phi_1(\mathcal{A}_1)|C_1$: a parfactor in model G with convergent variables $\mathcal{A}_1 = \{A_{11}, \dots, A_{1k}\}$
- (2) $g_2 = \phi_2(\mathcal{A}_2)|C_2$: a parfactor in model G with convergent variables $\mathcal{A}_2 = \{A_{21}, \dots, A_{2k}\}$
- (3) $\theta = \{\mathbf{X}_1 \rightarrow \mathbf{X}_2\}$: an alignment between g_1 and g_2

Preconditions:

- (1) for $i = 1, 2$: $\mathbf{Y}_i = \text{logvar}(\mathcal{A}_i) \setminus \mathbf{X}_i$ is count-normalized w.r.t. $\mathbf{X}_i$ in C_i

Output: $\phi(\mathcal{A})|C$, such that

- (1) $C = C_1 \theta \bowtie C_2$
- (2) $\mathcal{A} = \mathcal{A}_1 \theta \cup \mathcal{A}_2$
- (3) Let $\mathcal{A}$ be $(A_1, \dots, A_k, \mathcal{B})$ with $A_j = A_{1j} \theta = A_{2j}$ for $j = 1, \dots, k$, $\mathcal{B}$ the set of regular variables
- (4) for each assignment $\mathbf{a} = (a_1, \dots, a_k, \mathbf{b})$ to $\mathcal{A}$ with $\mathbf{b}_1 = \pi_{\mathcal{A}_1 \theta}(\mathbf{b})$, $\mathbf{b}_2 = \pi_{\mathcal{A}_2}(\mathbf{b})$

$$\phi(a_1, \dots, a_k, \mathbf{b}) = \sum_{a_{11} \vee a_{21} = a_1} \dots \sum_{a_{1k} \vee a_{2k} = a_k} \phi_1(a_{11}, \dots, a_{1k}, \mathbf{b}_1)^{1/r_2} \phi_2(a_{21}, \dots, a_{2k}, \mathbf{b}_2)^{1/r_1}$$

with $r_i = \text{COUNT}_{\mathbf{Y}_i|\mathbf{X}_i}(C_i)$

Postcondition: $G \sim G \setminus \{g_1, g_2\} \cup \{\text{HET-MULTIPLY}(g_1, g_2, \theta)\}$

Operator 1. Operator HET-MULTIPLY.

set of constraints. Following standard lifted multiplication, we assume the same PRV will have a different instance in each grounded factor. We thus proceed very much as in the grounded case, and for each case $(a_{11}, \dots, a_{1k}, \mathbf{b}_1, \mathbf{b}_2)$ we sum the potentials obtained by multiplying the $\phi_1(a_{11}, \dots, a_{1k}, \mathbf{b}_1)$ and $\phi_2(a_{21}, \dots, a_{2k}, \mathbf{b}_2)$. Note that although potentials need not be normalised to sum to 1 until the end, the relative counts of ϕ_1 and ϕ_2 must be weighed by considering the number of instances ϕ_2 for each ϕ_1 .

Example 5

Consider the heterogeneous parfactors $g_1 = \phi_1(p(X_1))|C_1$ and $g_2 = \phi_2(p(X_2), q(X_2, Y_2))|C_2$ and suppose that we want to multiply g_1 and g_2 ; $p(X)$ is convergent in g_1 and g_2 ; $\{X_1 \rightarrow X_2\}$ is an alignment between g_1 and g_2 ; Y_2 is count-normalized w.r.t. X_2 in C_2 ; $r_1 = \text{COUNT}_{Y_1|X_1}(C_1) = 2$ and $\text{COUNT}_{Y_2|X_2}(C_2) = 3$. Then $\text{HET-MULTIPLY}(g_1, g_2, \{X_1 \rightarrow X_2\}) = \phi(p(X_2), q(X_2, Y_2))|C$ with ϕ given by:

	$\phi(p(X_2), q(X_2, Y_2))$
ff	$\phi_1(f)^{1/3} \phi_2(f, f)^{1/2}$
ft	$\phi_1(f)^{1/3} \phi_2(f, t)^{1/2}$
tf	$\phi_1(f)^{1/3} \phi_2(t, f)^{1/2} + \phi_1(t)^{1/3} \phi_2(f, f)^{1/2} + \phi_1(t)^{1/3} \phi_2(t, f)^{1/2}$
tt	$\phi_1(f)^{1/3} \phi_2(t, t)^{1/2} + \phi_1(t)^{1/3} \phi_2(f, t)^{1/2} + \phi_1(t)^{1/3} \phi_2(t, t)^{1/2}$

Example 6

Consider the heterogeneous parfactors $g_1 = \phi_1(p(X_1), q(X_1, Y_1))|C_1$ and $g_2 = \phi_2(p(X_2), q(X_2, Y_2))|C_2$ and suppose we want to multiply g_1 and g_2 ; all randvars are convergent; $\{X_1 \rightarrow X_2, Y_1 \rightarrow Y_2\}$ is an alignment between g_1 and g_2 (so $r_1 = r_2 = 1$). Then $\text{HET-MULTIPLY}(g_1, g_2, \{X_1 \rightarrow X_2, Y_1 \rightarrow Y_2\}) = \phi(p(X_2), q(X_2, Y_2))|C$, with ϕ given by

	$\phi(p(X_2), q(X_2, Y_2))$
ff	$\phi_1(f, f) \phi_2(f, f)$
ft	$\phi_1(f, f) \phi_2(f, t) + \phi_1(f, t) \phi_2(f, f) + \phi_1(f, t) \phi_2(f, t)$
tf	$\phi_1(t, f) \phi_2(f, f) + \phi_1(f, f) \phi_2(t, f) + \phi_1(t, f) \phi_2(t, f)$
tt	$\phi_1(f, f) \phi_2(t, t) + \phi_1(f, t) \phi_2(t, f) + \phi_1(f, t) \phi_2(t, t) +$ $\phi_1(t, f) \phi_2(f, t) + \phi_1(t, t) \phi_2(f, f) + \phi_1(t, t) \phi_2(f, t) +$ $\phi_1(t, f) \phi_2(t, t) + \phi_1(t, t) \phi_2(t, f) + \phi_1(t, t) \phi_2(t, t)$

The SUM-OUT operator must be modified as well. In fact, consider the case in which a random variable must be summed out from a heterogeneous factor (i.e. a factor that

Operator HET-SUM-OUT**Inputs:**

- (1) $g = \phi(\mathcal{A})|C$: a parfactor in model G
- (2) let $\mathcal{A} = (A_1, \dots, A_k, A_{k+1}, \mathcal{B})$ where $A_1, \dots, A_k$ are convergent atoms
- (3) A_{k+1} is the atom to be summed out

Preconditions:

- (1) For all PRVs $\mathcal{V}$, other than $A_{k+1}|C$, in G : $RV(\mathcal{V}) \cap RV(A_{k+1}|C) = \emptyset$
- (2) A_{k+1} contains all the logvars $X \in \text{logvar}(\mathcal{A})$ for which $\pi_X(C)$ is not singleton
- (3) $\mathbf{X}^{excl} = \text{logvar}(A_{k+1}) \setminus \text{logvar}(\mathcal{A} \setminus A_{k+1})$ is count-normalized w.r.t.
 $\mathbf{X}^{com} = \text{logvar}(A_{k+1}) \cap \text{logvar}(\mathcal{A} \setminus A_{k+1})$ in C

Output: $\phi'(\mathcal{A}')|C'$, such that

- (1) $\mathcal{A}' = \mathcal{A} \setminus A_{k+1}$
- (2) $C' = \pi_{\mathbf{X}}(C)$
- (3) for each assignment $(\mathbf{a}', \mathbf{b}) = (a'_1, \dots, a'_k, \mathbf{b})$, to $\mathcal{A}'$

$$\begin{aligned} \phi'(\mathbf{a}', \mathbf{b}) = & \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \sum_{a_{k+1} \in \text{range}(A_{k+1})} \text{MUL}(A_{k+1}, a_{k+1}) \phi(a_1, \dots, a_k, a_{k+1}, \mathbf{b}) \right)^r - \\ & - \sum_{\mathbf{a} < \mathbf{a}'} \phi'(a_1, \dots, a_k, \mathbf{b}) \text{ with} \\ r = & \text{COUNT}_{\mathbf{X}^{excl}|\mathbf{X}^{com}}(C) \end{aligned}$$

Postcondition: $\mathcal{P}_{G \setminus \{g\}} \cup \{\text{HET-SUM-OUT}(g, (A_1, \dots, A_k), A_{k+1})\} = \sum_{RV(A_{k+1})} \mathcal{P}_G$

Operator 2. Operator HET-SUM-OUT. The order $\leq$ between truth values is the obvious one and between tuples of truth values is the product order induced by $\leq$ between values, i.e., $(a_1, \dots, a_k) \leq (a'_1, \dots, a'_k)$ iff $a_i \leq a'_i$ for $i = 1, \dots, k$ and $\mathbf{a} < \mathbf{a}'$ iff $\mathbf{a} \leq \mathbf{a}'$ and $\mathbf{a}' \not\leq \mathbf{a}$. Function MUL is defined in Appendix B.

contains a convergent variable). Consider for example the factor $\phi(p(X), q(X, Y))|C$ with $C = \{x_1, x_2\} \times \{y_1, y_2\}$ and suppose we want to eliminate the PRV $q(X, Y)$. This factor stands for four ground factors of the form $\phi(p(x_i), q(x_i, y_j))$ for $i, j = 1, 2$ where $p(x_i)$ is convergent. Given an individual x_i , the two factors $\phi(p(x_i), q(x_i, y_1))$ and $\phi(p(x_i), q(x_i, y_2))$ share a convergent variable and cannot be multiplied together with regular multiplication. In order to sum out $q(X, Y)$ however we must first combine the two factors with heterogeneous multiplication. To avoid generating first the ground factors, we have added to GC-FOVE HET-SUM-OUT (operator 2) that performs the combination and the elimination of a random variable at the same time. We provide a correctness proof for this operator in Appendix C.

Example 7

Consider the heterogeneous parfactor $g = \phi(r, p(X), q(X, Y))|C$ and suppose that we want to sum out $q(X, Y)$, that r and $p(X)$ are convergent, that Y is count-normalized w.r.t. X and that $\text{COUNT}_{Y|X}(C) = 2$. Then $\text{HET-SUM-OUT}(g, (r, p(X)), q(X, Y)) = \phi'(r, p(X))|C'$ with ϕ' given by

	$\phi'(r, p(X))$
ff	$(\phi(f, f, f) + \phi(f, f, t))^2$
ft	$(\phi(f, t, f) + \phi(f, t, t) + \phi(f, f, f) + \phi(f, f, t))^2 - \phi'(f, f)$
tf	$(\phi(t, f, f) + \phi(t, f, t) + \phi(f, f, f) + \phi(f, f, t))^2 - \phi'(f, f)$
tt	$(\phi(t, t, f) + \phi(t, t, t) + \phi(t, f, f) + \phi(t, f, t) + \phi(f, t, f) + \phi(f, t, t) + \phi(f, f, f) + \phi(f, f, t))^2 -$ $-\phi'(t, f) - \phi'(f, t) - \phi'(f, f)$

5 Experiments

In order to evaluate the performance of LP² algorithm, we compare it with PITA and ProbLog2 in two problems: *workshops attributes* (Milch et al. 2008) and Example 7 in (Poole 2008) that we call *plates*. Code of all the problems can be found in Appendix A.




Fig. 1. Runtime of LP², PITA and ProbLog2 on the *workshops attributes* problems. The Y-axis (runtime) is drawn in log scale. Note that, with the number of people fixed to 50, the problem is intractable by ProbLog2 which can manage at most two attributes.

Moreover, we did a scalability test on a third problem: *competing workshops* (Milch et al. 2008). All the tests were done on a machine with an Intel Dual Core E6550 2.33GHz processor and 4GB of main memory. The *workshops attributes* problem differs from Example 1 because the first clause for **series** is missing and the second clause contains a probabilistic atom in its body. The *competing workshops* problem differs from *workshops attributes* because it considers, instead of workshop attributes, a set of competing workshops W each one associated with a binary random variable $hot(W)$, which indicates whether it is focusing on popular research areas. The *plates* problem is an artificial example which contains two sets of individuals, X and Y . The distribution is defined by 7 probabilistic facts and 9 rules.

Figure 1(a) shows the runtime of LP², PITA and ProbLog2 on the *workshops attributes* problem for the query **series** where we fixed the number of people to 50 and we increased the number of attributes m . As expected, LP² is able to solve a much larger set of problems than PITA and ProbLog2. Figure 1(b) shows the time spent by LP² with up to 10^5 attributes. Figure 2(a) shows the runtime of LP², PITA and ProbLog2 on the *plates* problem, while figure 2(b) shows that of LP² with up to 12×10^4 Y individuals. For this test we executed the query **f** and we fixed the number of different X individuals to 5 and we increased the number of Y individuals.

Finally, we used the *competing workshops* problem for testing the scalability of LP². The trend was calculated performing the query **series** with 10 competing workshops and an increasing number n of people problems. Figure 3 shows LP² computation time. The trend is almost linear in the number of people contained in the problem.

As the results show, LP² can manage domains that are of several orders of magnitude larger than the ones managed by PITA and ProbLog2 in a shorter time.

6 Conclusions

We have shown that the Lifted Variable Elimination approach is very effective at resolving queries w.r.t. probabilistic logic programs containing large amount of facts. We have proved that with the introduction of the two new heterogeneous operators we can




Fig. 2. Performance on *plates* with an increasing number of Y individuals. The Y-axis (runtime) is drawn in log scale.




Fig. 3. LP² runtime in the *competing workshops* problem. The Y-axis (runtime) is drawn in log scale.

compute the probability of queries following the distribution semantics in a very efficient way. Experimental evidence shows that LP² can achieve several orders of magnitude improvements, both in runtime and in the number of facts that can be managed effectively. In the future, we plan to compare our approach with that of (Kisynski and Poole 2009b; Takikawa and D'Ambrosio 1999; Díez and Galán 2003) for dealing with noisy or factors, and to compare with weighted first order model counting (Van den Broeck et al. 2014).

Acknowledgments: VSC was partially nanced by the North Portugal Regional Operational Programme (ON.2 O Novo Norte), under the NSRF, through the ERDF and the Fundação para a Ciência e a Tecnologia within project ADE/PTDC/EIA-EIA/121686/2010.

References

- DE RAEDT, L., FRASCONI, P., KERSTING, K., AND MUGGLETON, S., Eds. 2008. *Probabilistic Inductive Logic Programming - Theory and Applications*. LNCS, vol. 4911. Springer.
- DE RAEDT, L., KIMMIG, A., AND TOIVONEN, H. 2007. ProbLog: A Probabilistic Prolog and its

- application in link discovery. In *20th International Joint Conference on Artificial Intelligence (IJCAI-2007)*. AAAI Press, 2462–2467.
- DE SALVO BRAZ, R., AMIR, E., AND ROTH, D. 2005. Lifted first-order probabilistic inference. In *19th International Joint Conference on Artificial Intelligence*, L. P. Kaelbling and A. Saffiotti, Eds. Professional Book Center, 1319–1325.
- DÍEZ, F. J. AND GALÁN, S. F. 2003. Efficient computation for the noisy max. *International Journal of Intelligent Systems*, 165–177.
- FIERENS, D., VAN DEN BROECK, G., RENKENS, J., SHTERIONOV, D., GUTMANN, B., THON, I., JANSSENS, G., AND DE RAEDT, L. 2014. Inference and learning in probabilistic logic programs using weighted boolean formulas. *Theory and Practice of Logic Programming First View Articles*.
- GETOOR, L. AND TASKAR, B., Eds. 2007. *Introduction to Statistical Relational Learning*. MIT Press.
- GOMES, T. AND COSTA, V. S. 2012. Evaluating inference algorithms for the prolog factor language. In *22nd International Conference on Inductive Logic Programming*, F. Riguzzi and F. Zelezný, Eds. LNCS, vol. 7842. Springer, 74–85.
- KISYNSKI, J. AND POOLE, D. 2009a. Constraint processing in lifted probabilistic inference. In *25th Conference on Uncertainty in Artificial Intelligence*, J. Bilmes and A. Y. Ng, Eds. AUAI Press, 293–302.
- KISYNSKI, J. AND POOLE, D. 2009b. Lifted aggregation in directed first-order probabilistic models. In *24th International Joint Conference on Artificial Intelligence*, C. Boutilier, Ed. 1922–1929.
- MEERT, W., STRUYF, J., AND BLOCKEEL, H. 2008. Learning ground CP-Logic theories by leveraging Bayesian network learning techniques. *Fundamenta Informaticae* 89, 131–160.
- MILCH, B., ZETTMLOYER, L. S., KERSTING, K., HAIMEs, M., AND KAEHLING, L. P. 2008. Lifted probabilistic inference with counting formulas. In *23rd AAAI Conference on Artificial Intelligence*, D. Fox and C. P. Gomes, Eds. AAAI Press, 1062–1068.
- POOLE, D. 1993. Probabilistic horn abduction and Bayesian networks. *Artificial Intelligence* 64, 1, 81–129.
- POOLE, D. 1997. The Independent Choice Logic for modelling multiple agents under uncertainty. *Artificial Intelligence* 94, 7–56.
- POOLE, D. 2003. First-order probabilistic inference. In *18th International Joint Conference on Artificial Intelligence*, G. Gottlob and T. Walsh, Eds. Morgan Kaufmann Publishers Inc., 985–991.
- POOLE, D. 2008. The independent choice logic and beyond. In *Probabilistic Inductive Logic Programming*, L. De Raedt, P. Frasconi, K. Kersting, and S. Muggleton, Eds. LNCS, vol. 4911. Springer, 222–243.
- RIGUZZI, F. AND SWIFT, T. 2011. The PITA system: Tabling and answer subsumption for reasoning under uncertainty. *Theory and Practice of Logic Programming, International Conference on Logic Programming (ICLP) Special Issue 11*, 433–449.
- SATO, T. 1995. A statistical learning method for logic programs with distribution semantics. In *12th International Conference on Logic Programming*, L. Sterling, Ed. MIT Press, 715–729.
- TAGHIPOUR, N., FIERENS, D., DAVIS, J., AND BLOCKEEL, H. 2013. Lifted variable elimination: Decoupling the operators from the constraint language. *Journal of Artificial Intelligence Research* 47, 393–439.
- TAKIKAWA, M. AND D’AMBROSIO, B. 1999. Multiplicative factorization of noisy-max. In *15th Conference on Uncertainty in Artificial Intelligence*. 622–630.
- VAN DEN BROECK, G., MEERT, W., AND DARWICHE, A. 2014. Skolemization for weighted first-order model counting. *ArXiv e-prints* 1312.5378v2. To appear in the 14th International Conference on Principles of Knowledge Representation and Reasoning.

- VAN DEN BROECK, G., TAGHIPOUR, N., MEERT, W., DAVIS, J., AND RAEDT, L. D. 2011. Lifted probabilistic inference by first-order knowledge compilation. In *21st International Joint Conference on Artificial Intelligence*, T. Walsh, Ed. IJCAI/AAAI, 2178–2185.
- VENNEKENS, J., VERBAETEN, S., AND BRUYNOOGHE, M. 2004. Logic Programs With Annotated Disjunctions. In *20th International Conference on Logic Programming*, B. Demoen and V. Lifschitz, Eds. Springer, LNCS 3131, 195–209.
- ZHANG, N. L. AND POOLE, D. L. 1996. Exploiting causal independence in bayesian network inference. *Journal of Artificial Intelligence Research* 5, 301–328.

Appendix A Problems code

In this section we present the Problog and PFL code of the testing problems.

A.1 Workshops Attributes

To all programs of this section we added 50 workshops and an increasing number of attributes.

Problog program

```
series:- person(P),attends(P),sa(P).

0.501::sa(P):-person(P).

attends(P):- person(P),attr(A),at(P,A).

0.3::at(P,A):-person(P),attr(A).
```

PFL program

```
het series1,ch1(P);[1.0, 0.0, 0.0, 1.0];[person(P)].

deputy series,series1;[].

bayes ch1(P),attends(P),sa(P);[1.0,1.0,1.0,0.0,
                                0.0,0.0,0.0,1.0];[person(P)].

bayes sa(P);[0.499,0.501];[person(P)].

het attends1(P),at(P,A);[1.0, 0.0, 0.0, 1.0];[person(P),attr(A)].

deputy attends(P),attends1(P);[person(P)].

bayes at(P,A);[0.7,0.3];[person(P),attr(A)].
```

A.2 Competing Workshops

For the *competing workshops* problem we report only the PFL version. For testing purpose we added to this code 10 workshops and an increasing number of people.

PFL program

```
bayes ch1(P),attends(P),sa(P);[1.0,1.0,1.0,0.0,
                                0.0,0.0,0.0,1.0];[person(P)].

het series1,ch1(P);[1.0, 0.0, 0.0, 1.0];[person(P)].
```

```

deputy series,series1; [].

bayes sa(P);[0.499,0.501];[person(P)].

het attends1(P),ch2(P,W);[1.0, 0.0, 0.0, 1.0];[person(P),workshop(W)].

deputy attends(P),attends1(P);[person(P)].

bayes ch2(P,W),hot(W),ah(P,W);[1.0,1.0,1.0,0.0,
                                0.0,0.0,0.0,1.0];[person(P),workshop(W)].

bayes ah(P,W);[0.2,0.8];[person(P),workshop(W)].

```

A.3 Plates

For the *plates* problem we added 5 individuals for X and an increasing number of individuals for Y .

Problog program

```

f:- e(Y).

e(Y) :- d(Y),n1(Y).
e(Y) :- y(Y),\+ d(Y),n2(Y).

d(Y):- c(X,Y).

c(X,Y):-b(X),n3(X,Y).
c(X,Y):- x(X),\+ b(X),n4(X,Y).

b(X):- a, n5(X).
b(X):- \+ a,n6(X).

a:- n7.

0.1::n1(Y) :-y(Y).
0.2::n2(Y) :-y(Y).
0.3::n3(X,Y) :- x(X),y(Y).
0.4::n4(X,Y) :- x(X),y(Y).
0.5::n5(X) :-x(X).
0.6::n6(X) :-x(X).
0.7::n7.

```

PFL program

```

het f1,e(Y);[1.0, 0.0, 0.0, 1.0];[y(Y)].

deputy f,f1; [].

bayes e1(Y),d(Y),n1(Y);[1.0, 1.0, 1.0, 0.0,
                        0.0, 0.0, 0.0, 1.0];[y(Y)].

bayes e2(Y),d(Y),n2(Y);[1.0, 0.0, 1.0, 1.0,
                        0.0, 1.0, 0.0, 0.0];[y(Y)].

bayes e(Y),e1(Y),e2(Y);[1.0, 0.0, 0.0, 0.0,
                        0.0, 1.0, 1.0, 1.0];[y(Y)].

het d1(Y),c(X,Y);[1.0, 0.0, 0.0, 1.0];[x(X),y(Y)].

deputy d(Y),d1(Y);[y(Y)].

bayes c1(X,Y),b(X),n3(X,Y);[1.0, 1.0, 1.0, 0.0,

```

```

0.0, 0.0, 0.0, 1.0];[x(X),y(Y)].

bayes c2(X,Y),b(X),n4(X,Y);[1.0, 0.0, 1.0, 1.0,
0.0, 1.0, 0.0, 0.0];[x(X),y(Y)].

bayes c(X,Y),c1(X,Y),c2(X,Y);[1.0, 0.0, 0.0, 0.0,
0.0, 1.0, 1.0, 1.0];[x(X),y(Y)].

bayes b1(X),a,n5(X);[1.0, 1.0, 1.0, 0.0,
0.0, 0.0, 0.0, 1.0];[x(X)].

bayes b2(X),a,n6(X);[1.0, 0.0, 1.0, 1.0,
0.0, 1.0, 0.0, 0.0];[x(X)].

bayes b(X),b1(X),b2(X);[1.0, 0.0, 0.0, 0.0,
0.0, 1.0, 1.0, 1.0];[x(X)].

bayes a,n7;[1.0, 0.0, 0.0, 1.0];[].

bayes n1(Y);[0.9, 0.1];[y(Y)].
bayes n2(Y);[0.8, 0.2];[y(Y)].
bayes n3(X,Y);[0.7, 0.3];[x(X),y(Y)].
bayes n4(X,Y);[0.6, 0.4];[x(X),y(Y)].
bayes n5(X);[0.5, 0.5];[x(X)].
bayes n6(X);[0.4, 0.6];[x(X)].
bayes n7;[0.3, 0.7];[].

```

Appendix B Definitions

Definition 1 (counting formula)

A counting formula is a syntactic construct of the form $\#_{X_i \in C}[F(\mathbf{X})]$, where $X_i \in \mathbf{X}$ is called the counted logvar.

A *grounded* counting formula is a counting formula in which all arguments of the atom $F(\mathbf{X})$, except for the counted logvar, are constants. It defines a counting randvar (CRV) as follows.

Definition 2 (counting randvar)

A parametrized counting randvar (PCRIV) is a pair $(\#_{X_i}[F(\mathbf{X})], C)$. For each instantiation of $\mathbf{X} \setminus X_i$, it creates a separate counting randvar (CRV). The value of this CRV is a histogram, and it depends deterministically on the values of $F(\mathbf{X})$. Given a valuation for $F(\mathbf{X})$, it counts how many different values of X_i occur for each $r \in \text{range}(F)$. The result is a *histogram* of the form $\{(r_1, n_1), \dots, (r_k, n_k)\}$, with $r_i \in \text{range}(F)$ and n_i the corresponding count.

Definition 3 (multiplicity)

The multiplicity of a histogram $h = \{(r_1, n_1), \dots, (r_k, n_k)\}$ is a multinomial coefficient, defined as

$$\text{MUL}(h) = \frac{n!}{\prod_{i=1}^k n_i!}.$$

As multiplicities should only be taken into account for (P)CRVs, never for regular PRVs, we define for each PRV A and for each value $v \in \text{range}(A)$: $\text{MUL}(A, v) = 1$ if A is a regular PRV, and $\text{MUL}(A, v) = \text{MUL}(v)$ if A is a PCRIV. This MUL function is identical to (Milch et al. 2008)'s NUM-ASSIGN.

Definition 4 (Count function)

Given a constraint $C_{\mathbf{X}}$, for any $\mathbf{Y} \subseteq \mathbf{X}$ and $\mathbf{Z} \subseteq \mathbf{X} - \mathbf{Y}$, the function $\text{COUNT}_{\mathbf{Y}|\mathbf{Z}} : C_{\mathbf{X}} \rightarrow \mathbb{N}$ is defined as follows:

$$\text{COUNT}_{\mathbf{Y}|\mathbf{Z}}(t) = |\pi_{\mathbf{Y}}(\sigma_{\mathbf{Z}=\pi_{\mathbf{Z}}(t)}(C_{\mathbf{X}}))|$$

That is, for any tuple t , this function tells us how many values for $\mathbf{Y}$ co-occur with t 's value for $\mathbf{Z}$ in the constraint. We define $\text{COUNT}_{\mathbf{Y}|\mathbf{Z}}(t) = 1$ when $\mathbf{Y} = \emptyset$.

Definition 5 (Count-normalized constraint)

For any constraint $C_{\mathbf{X}}$, $\mathbf{Y} \subseteq \mathbf{X}$ and $\mathbf{Z} \subseteq \mathbf{X} - \mathbf{Y}$, $\mathbf{Y}$ is count-normalized w.r.t. $\mathbf{Z}$ in $C_{\mathbf{X}}$ if and only if

$$\exists n \in \mathbb{N} : \forall t \in C_{\mathbf{X}} : \text{COUNT}_{\mathbf{Y}|\mathbf{Z}}(t) = n.$$

When such an n exists, we call it the conditional count of $\mathbf{Y}$ given $\mathbf{Z}$ in $C_{\mathbf{X}}$, and denote it $\text{COUNT}_{\mathbf{Y}|\mathbf{Z}}(C_{\mathbf{X}})$.

Definition 6 (substitution)

A substitution $\theta = \{X_1 \rightarrow t_1, \dots, X_n \rightarrow t_n\} = \{\mathbf{X} \rightarrow \mathbf{t}\}$ maps each logvar X_i to a term t_i , which can be a constant or a logvar. When all t_i are constants, θ is called a grounding substitution, and when all are different logvars, a renaming substitution. Applying a substitution θ to an expression α means replacing each occurrence of X_i in α with t_i ; the result is denoted $\alpha\theta$.

Definition 7 (alignment)

An alignment θ between two parfactors $g = \phi(\mathcal{A})|C$ and $g' = \phi'(\mathcal{A}')|C'$ is a one-to-one substitution $\{\mathbf{X} \rightarrow \mathbf{X}'\}$, with $\mathbf{X} \subseteq \text{logvar}(\mathcal{A})$ and $\mathbf{X}' \subseteq \text{logvar}(\mathcal{A}')$, such that $\rho(\pi_{\mathbf{X}'}(C)) = \pi_{\mathbf{X}'}(C')$ (with ρ the attribute renaming operator).

An alignment tells the multiplication operator that two atoms in two different parfactors represent the same PRV, so it suffices to include it in the resulting parfactor only once.

Appendix C Correctness proof for heterogeneous multiplication

Theorem 1

Given a model $(\mathcal{F}_1, \mathcal{F}_2)$, two heterogeneous parfactors $g_1, g_2 \in \mathcal{F}_2$ and an alignment θ between g_1 and g_2 , if the preconditions of the HET-MULTIPLY operator are fulfilled then the postcondition

$$G \setminus \{g_1, g_2\} \cup \{\text{HET-MULTIPLY}(g_1, g_2, \theta)\}$$

holds.

Proof

Immediate from the definition of heterogeneous multiplication. $\square$

Theorem 2

Given a model $(\mathcal{F}_1, \mathcal{F}_2)$, a heterogeneous parfactor $g \in \mathcal{F}_2$ and an atom A_{k+1} to be summed out, if the preconditions of the HET-SUM-OUT operator are fulfilled then the postcondition

$$\mathcal{P}_{G \setminus \{g\} \cup \{\text{HET-SUM-OUT}(g, (A_1, \dots, A_k), A_{k+1})\}} = \sum_{RV(A_{k+1})} \mathcal{P}_g$$

holds.

Proof

We prove the formula giving ϕ' in Operator 2 by double induction over $r = \text{COUNT}_{\mathbf{X}_{excl}|\mathbf{X}_{com}}(C)$ and the number n of values at t in the tuple $(a'_1, \dots, a'_k)$. For simplicity we assume that the variable to be summed out, A_{k+1} , is not a counting variable, but the same reasoning can be applied for a counting variable. For $n = 0, r = 1$

$$\phi'(f, \dots, f, \mathbf{b}) = \phi(f, \dots, f, f, \mathbf{b}) + \phi(f, \dots, f, t, \mathbf{b})$$

so the thesis is proved. For $n = 0, r > 1$, let us call $\phi'_r(a'_1, \dots, a'_k, \mathbf{b})$ the value of ϕ' for r . Let us assume that the formula holds for $r - 1$. For $r > 1$, there is an extra valuation $\mathbf{x}_{excl}$ for $\mathbf{X}_{excl}$ given $\mathbf{X}_{com}$ so there is an extra factor $g''(\mathbf{x}_{excl})$. Eliminating A_{k+1} from g'' gives

$$\phi'_1(f, \dots, f, \mathbf{b}) = \phi(f, \dots, f, f, \mathbf{b}) + \phi(f, \dots, f, t, \mathbf{b})$$

This must be multiplied by ϕ'_{r-1} with heterogeneous multiplication as $A_1, \dots, A_k$ are shared obtaining

$$\begin{aligned} \phi'_r(f, \dots, f, \mathbf{b}) &= \sum_{\hat{a}_1 \vee \tilde{a}_1 = f} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = f} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \times \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, f, \mathbf{b}) = \\ &= \phi'_{r-1}(f, \dots, f, \mathbf{b}) \times \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, f, \mathbf{b}) = \\ &= (\phi(f, \dots, f, f, \mathbf{b}) + \phi(f, \dots, f, t, \mathbf{b}))^{r-1} \times (\phi(f, \dots, f, f, \mathbf{b}) + \phi(f, \dots, f, t, \mathbf{b})) = \\ &= (\phi(f, \dots, f, f, \mathbf{b}) + \phi(f, \dots, f, t, \mathbf{b}))^r \end{aligned}$$

so the thesis is proved.

For $n > 0$ of values at t in the tuple $\mathbf{a}' = (a'_1, \dots, a'_k)$, we assume that the formula holds for $(n - 1, r)$ and $(n, r - 1)$. For the definition of heterogeneous multiplication

$$\phi'_r(a'_1, \dots, a'_k, \mathbf{b}) = \sum_{\hat{a}_1 \vee \tilde{a}_1 = a'_1} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = a'_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b})$$

By adding and removing

$$\sum_{\mathbf{a} < \mathbf{a}'} \sum_{\hat{a}_1 \vee \tilde{a}_1 = a_1} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = a_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b})$$

we get

$$\begin{aligned} \phi'_r(a'_1, \dots, a'_k, \mathbf{b}) &= \\ &= \sum_{\hat{a}_1 \vee \tilde{a}_1 = a'_1} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = a'_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) + \\ &+ \sum_{\mathbf{a} < \mathbf{a}'} \sum_{\hat{a}_1 \vee \tilde{a}_1 = a_1} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = a_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) - \\ &- \sum_{\mathbf{a} < \mathbf{a}'} \sum_{\hat{a}_1 \vee \tilde{a}_1 = a_1} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = a_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) = \\ &= \sum_{\hat{a}_1 \leq a'_1} \sum_{\tilde{a}_1 \leq a'_1} \dots \sum_{\hat{a}_k \leq a'_k} \sum_{\tilde{a}_k \leq a'_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) - \\ &- \sum_{\mathbf{a} < \mathbf{a}'} \sum_{\hat{a}_1 \vee \tilde{a}_1 = a_1} \dots \sum_{\hat{a}_k \vee \tilde{a}_k = a_k} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) = \end{aligned}$$

For the definition of heterogeneous multiplication we obtain $\phi'_r(a'_1, \dots, a'_k, \mathbf{b}) =$

$$= \left(\sum_{\hat{\mathbf{a}} \leq \mathbf{a}'} \phi'_{r-1}(\hat{a}_1, \dots, \hat{a}_k, \mathbf{b}) \right) \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b})$$

By applying the inductive hypothesis for $r - 1$ we get $\phi'_r(a'_1, \dots, a'_k, \mathbf{b}) =$

$$\begin{aligned}
&= \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right) \right)^{r-1} - \sum_{\mathbf{a} < \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \Bigg) \times \\
&\quad \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) = \\
&= \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right) \right)^{r-1} \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \sum_{\mathbf{a} < \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) = \\
&= \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^{r-1} \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) + \\
&\quad + \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right) \right)^{r-1} \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \sum_{\mathbf{a} < \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) =
\end{aligned}$$

By applying the formula for $r = 1$

$$\phi'_r(a'_1, \dots, a'_k, \mathbf{b}) =$$

$$\begin{aligned}
&= \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^{r-1} \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right) + \\
&\quad + \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right) \right)^{r-1} \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \sum_{\mathbf{a} < \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) = \\
&= \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^r - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) + \\
&\quad + \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right) \right)^{r-1} \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) - \\
&\quad - \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \sum_{\mathbf{a} < \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right)
\end{aligned}$$

$$\begin{aligned}
& \text{By collecting the factor } \sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \phi'_r(a'_1, \dots, a'_k, \mathbf{b}) = \\
& = \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^r - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) + \\
& + \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^{r-1} - \sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \\
& \times \left(\sum_{\tilde{\mathbf{a}} \leq \mathbf{a}'} \phi'_1(\tilde{a}_1, \dots, \tilde{a}_k, \mathbf{b}) \right) = \\
& = \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^r - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) + \\
& + \left(\sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \left(\sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^{r-1} - \sum_{\tilde{\mathbf{a}} < \mathbf{a}'} \sum_{\mathbf{a} \leq \tilde{\mathbf{a}}} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) - \right. \\
& \left. - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi'_1(a_1, \dots, a_k, \mathbf{b}) \right) = \\
& = \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^r - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b}) + \\
& + \left(\sum_{\mathbf{a} < \mathbf{a}'} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_{r-1}(a_1, \dots, a_k, \mathbf{b}) \right) \times \sum_{\mathbf{a} \leq \mathbf{a}'} \phi'_1(a_1, \dots, a_k, \mathbf{b}) = \\
& = \left(\sum_{\mathbf{a} \leq \mathbf{a}'} \phi(a_1, \dots, a_k, f, \mathbf{b}) + \phi(a_1, \dots, a_k, t, \mathbf{b}) \right)^r - \sum_{\mathbf{a} < \mathbf{a}'} \phi'_r(a_1, \dots, a_k, \mathbf{b})
\end{aligned}$$

□