

Automatic Reordering for Dataflow Safety of DATALOG

Mistral Contrastin
University of Cambridge
Mistral.Contrastin@cl.cam.ac.uk

Dominic Orchard
University of Kent
d.a.orchard@kent.ac.uk

Andrew Rice
University of Cambridge
Andrew.Rice@cl.cam.ac.uk

ABSTRACT

Clauses and subgoals in a DATALOG program can be given in any order without affecting program meaning. However, practical applications of the language require the use of built-in or external predicates with particular dataflow requirements. These can be expressed as input or output “modes” on arguments. We describe a static analysis of moding for DATALOG which calculates how to transform an ill-moded program into a well-moded program by reordering clause subgoals to satisfy dataflow requirements. We describe an incremental algorithm which efficiently finds a reordering if it exists. This frees the programmer to focus on the declarative specification of a program rather than implementation details of external predicates. We prove that our computed reorderings yield well-moded programs (soundness) and if a program can be made well-moded, then we compute a reordering to do so (completeness).

CCS CONCEPTS

• **Theory of computation** → **Database query processing and optimization (theory)**; **Program analysis**; Logic and databases;
• **Software and its engineering** → **Constraint and logic languages**; • **Computing methodologies** → *Logic programming and answer set programming*;

ACM Reference Format:

Mistral Contrastin, Dominic Orchard, and Andrew Rice. 2018. Automatic Reordering for Dataflow Safety of DATALOG. In *The 20th International Symposium on Principles and Practice of Declarative Programming (PPDP '18)*, September 3–5, 2018, Frankfurt am Main, Germany. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3236950.3236954>

1 INTRODUCTION

Declarative languages aim to free the programmer from implementation details, allowing them to focus on the essence of a problem. However, in practice, implementation details often creep back in. In logic programming, one such implementation detail is *subgoal ordering* which rears its head once we start introducing external functions into pure logic programs or when performance becomes a concern. The aim of this paper is to move the implementation concern of subgoal ordering from the programmer back to the language. We consider DATALOG, a syntactic subset of PROLOG which has recently been (re)growing in popularity. Among other things,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PPDP '18, September 3–5, 2018, Frankfurt am Main, Germany

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6441-6/18/09...\$15.00

<https://doi.org/10.1145/3236950.3236954>

it is used for managing enterprise data [1] and as a language for concisely and efficiently expressing program analyses [10, 14].

Sentences in pure DATALOG are Horn clauses, hence subgoals can be specified in *any order* since any ordering is trivially safe to invoke. However, this is not true in practice. Many systems require external functionality such as arithmetic, comparison, and input/output functions. For example, a function for printing to the console might choose not to print unbound values or an external hashing function may require its first parameter to be an input and its second to be an output. In large systems, ordering can also have performance implications, e.g., a database query might make more efficient use of an index if an argument is ground.

The concept of *moding* [12] allows us to specify the dataflow requirements of predicates. For example, a programmer can specify via moding that a particular argument of a clause must be bound before the clause is executed (as supported in MERCURY [16]). An invocation error occurs if this moding requirement is not met. *Well-moded* programs do not give invocation errors in the same way that well-typed programs do not go wrong.

We describe a few examples of ill-moded programs which can be fixed using the information from our analysis. Consider the following DATALOG clause with moding annotations as superscripts:

```
1 auth(U) :- hash+(P,H), password(U,P), valid(U,H).
```

The superscript $+$ specifies that an invocation of hash is safe when the first argument is bound and the second is free or bound. Therefore, this example is not well-moded since P is not bound in the context of the first subgoal (it would need to be bound by the clause head). However, in the absence of side-effects, it is sound to reorder the subgoals to meet the moding constraints by swapping the first and second subgoals so that password is invoked first, binding P under the usual left-to-right semantics.

As an alternative, one might write our example as two clauses:

```
1 auth(U) :- check(U,P), password(U,P).  
2 check(U,P) :- hash+(P,H), valid(U,H).
```

The hash subgoal in check is ill-moded because P is unbound in the invocation of check on line 1 in auth. The reordering to fix this is therefore non-local: one must reorder the body of auth to make the invocation of hash in check safe. Searching for valid orderings amongst all permutations of subgoals in all clauses of a program is infeasible. We instead propagate moding constraints to the callers.

We also provide information that can be used in conjunction with clause cloning. Imagine an interactive system with client- and server-side facilities for checking password strength:

```
1 client_check(P) :- weak(P,H).  
2 server_check(H) :- weak(P,H).  
3 weak(P,H) :- hash+(P,H), rainbow+(H,P).
```

The client side does not have access to the hash and the server side does not have access to the plaintext password. However both parties want to check if the password is compromised by looking

up the hash in a rainbow table¹ and confirming it is indeed the password corresponding to the hash. There is no reordering of `weak`'s subgoals that satisfy moding requirements of both `hash` and `rainbow`. If we generate two versions of `weak` with different subgoal orderings and use the appropriate one according to the binding pattern at the call site, then queries involving both client- and server-side checks can be well-moded:

```

1 client_check(P) :- weak1(P,H).
2 server_check(H) :- weak2(P,H).
3 weak1(P,H) :- hash+(P,H), rainbow+(H,P).
4 weak2(P,H) :- rainbow+(H,P), hash+(P,H).

```

Cloning arises as a by-product of a program transformation called *adornment* [17] (Section 3.2) which generates versions of each clause annotated with a variable binding pattern. Ordering information calculated by our analysis guides the adornment procedure and ensures that the subgoals used in the generated clauses are ordered such that they are safe to invoke.

In this paper, we show how to statically verify that DATALOG programs are well-moded. Furthermore, we give an algorithm that computes orderings of subgoals that satisfy the moding constraints of programs, if such an ordering exists. This is valuable because it frees the programmer to focus on the specification of the higher-level goals rather than their syntactic order. Although orderings are computed statically, we do not prescribe a time of use: they can be used to transform the DATALOG program statically (suitable for bottom-up or top-down evaluation) or to reorder subgoals dynamically as they are evaluated (suitable for top-down evaluation).

The key contributions of this paper are:

- a natural formalisation of well-modedness for DATALOG in terms of the adornment program transformation;
- a sound and complete incremental analysis algorithm for finding suitable subgoal orderings of program clauses and verifying the well-modedness of programs.

The rest of the paper is structured as follows. We first state our assumptions and notation (Section 2) and then formalise well-modedness in terms of the adornment program transformation for DATALOG (Section 3). This leads us to our mode analysis algorithm establishing well-modedness and producing subgoal orderings for binding patterns. This is presented in intra-clausal (Section 4) and inter-clausal (Section 5) stages along with properties of the algorithm. We modify the intra-clausal analysis to accommodate some DATALOG extensions (Section 6). Finally, we discuss previous approaches to mode analysis (Section 7) and conclude (Section 8). Proofs omitted from the main text are given in Appendix A.

2 NOTATION AND ASSUMPTIONS

Definition 2.1 (DATALOG program structure). A DATALOG program is a set of *clauses* of the form $p :- s_1, \dots, s_n$, where:

- p is called the clause *head*;
- s_i collectively forms the *body* of the clause, where each s_i is individually called a *subgoal*;
- the clause head and subgoals are *atomic formulae*;
- an *atomic formula* is a predicate symbol applied to a tuple of terms, e.g., $p(X, Y)$.

¹A pre-computed reverse lookup table from hashes to plaintext.

We assume subgoals are executed left-to-right and variables are bound to their values whenever possible. Binding a variable is the same as grounding it as there are no function symbols in DATALOG.

We consider programs that come with a *query*. A query is of the form $?- s_1, \dots, s_n$. Internally, this expands into a normal clause with a head predicate that does not appear elsewhere, and whose head parameters are all the variables that appear in the query.

We use a function *vars* to map the syntax of a logical formula or clause to a set of its variables. Functions *head* and *body* map a clause to its head and the set of its subgoals respectively.

Traditionally a predicate symbol and its arity uniquely identify a predicate however, for presentation purposes, we assume a predicate symbol alone uniquely identifies a predicate. A predicate may appear many times in a clause body, e.g., $p(X) :- q(X), q(X)$. has two distinct subgoals invoking the same predicate. The function *pred* maps a subgoal to its predicate and *arity* maps a predicate to its arity. If p is a predicate and Pr is a program, then Pr_p is the set of clauses that have p as their head.

Throughout, the “min” operator refers to the minimal elements of a set of subsets under the partial order defined by subset relation. For example, $\min\{\{1, 2\}, \{2, 3\}, \{1\}, \{1, 2, 3\}\}$ is $\{\{1\}, \{2, 3\}\}$.

The family of unary operators, π_i , projects the i^{th} component of appropriately-sized tuples.

We assume that the mode requirements of predicates are available and do not give a syntax for their declarations. This information may be hard-coded in the case of built-in predicates or supplied through mode declarations akin to type declarations.

3 ADORNMENT AND WELL-MODEDNESS

Informally, a well-moded program does not produce runtime errors arising from insufficient variable binding. In this section, we introduce the definitions needed to formalise this notion for DATALOG.

3.1 Mode annotations and constraints

Each variable only ever needs one of two modes in DATALOG. The mode $+$ indicates that the argument should be bound at the time of invocation and $?$ indicates it can either be bound or free.

As in the introduction, we use *mode vectors* written as superscripts to indicate the mode requirements of a predicate. Though these superscripts are placed on subgoals in our examples, they should be considered as global specifications for the predicate in the whole program. If a subgoal has no superscripts, then the underlying predicate has no mode requirements (equivalent to having $?$ for all argument positions). If a predicate is annotated with a set of mode vectors, e.g., $p^{(+?,?+)}(X, Y)$, then any one of them can be used to satisfy the dataflow requirements of the predicate. This set may arise from multiple implementations backing the predicate or, as we explore below, due to different orderings of subgoals leading to different moding requirements for user-defined predicates.

Invocation safety is based upon bound arguments, so instead of working with mode vectors directly, we use *constraints*.

Definition 3.1 (Constraint). For a predicate, p , any subset of its argument positions, $1 \leq i \leq \text{arity}(p)$, forms an *atomic constraint*. A set of atomic constraints which is minimal under the subset relation is a *constraint*. Thus, the domain of constraints for a predicate p is

$D_p = \{\min S \mid S \subseteq \mathcal{P}(\{i \mid 1 \leq i \leq \text{arity}(p)\})\}$. Throughout this text C is used to range over constraints and AC over atomic constraints.

Definition 3.2 (Mode requirement semantics). A mode vector is translated into an atomic constraint by taking the set of indices for which the mode is \pm . A set of mode vectors is converted to a constraint by translating each mode vector and removing all super sets. This translation is done by $\llbracket - \rrbracket$. For example, $\llbracket \{++?, ?++\} \rrbracket$ is $\{\{1, 2\}, \{2, 3\}\}$. Given a *mode function*, mv , from predicates to a set of mode vectors, $\llbracket - \rrbracket_F$ produces a *constraint function* from predicates to constraints and is defined as $\llbracket - \rrbracket \circ mv$. Throughout the text, mv and f range over mode and constraint functions respectively.

Definition 3.3 (Ill constraint). A constraint denotes alternative moding requirements and if this set is empty, then there are no alternatives that can be used for safe predicate invocation. Hence, \emptyset for any constraint domain is the unique *ill constraint*.

Definition 3.4 (Trivial constraint). For any constraint domain, the *trivial constraint*, $\{\emptyset\}$, contains an atomic constraint that does not require any variables to be bound.

3.2 Adornments and ordering

Program adornment [17] annotates the atomic formulae in clauses with binding patterns. This nicely formalises well-modedness. We use a generalised form of adornment which relies on a subgoal reordering function. This generalisation allows different binding patterns to be produced for subgoals depending on the ordering and allows us to derive a versatile well-modedness definition. We give a high-level definition of adornment transformation.

Definition 3.5 (Adornment). An *adornment* associates to an argument/parameter of an atomic formula either f or b indicating the binding status of the argument/parameter: *free* or *bound*. An *adornment vector* (or a *binding pattern*) for an atomic formula is a vector of adornments whose size matches the predicate's arity. Throughout, \mathbf{a}, \mathbf{b} range over adornment vectors which we index with natural numbers, i.e., a_i is the i^{th} adornment in the vector.

For some subgoal, sub , of an adorned clause, we denote the adornment vector of sub as $\text{adornment}(sub)$.

Definition 3.6 (Ordering). An *ordering* is a bijection between two lists of subgoals. When applied to a list of subgoals it permutes them. We use σ to range over orderings.

Definition 3.7 (Generalised clause adornment). Let adorn be a function that takes a clause cl , a binding pattern \mathbf{a} , an ordering σ for the clause, and returns the adorned and reordered (according to σ) version of the clause.

We calculate $\text{adorn}(cl, \mathbf{a}, \sigma)$ as follows: first, the clause head is assigned the binding pattern \mathbf{a} . Next, the body of the clause cl is reordered using σ . Finally, the list of subgoals is traversed left-to-right and adorned. For each argument that is a literal or a variable that is known to be bound, the argument receives the adornment b , otherwise it is given the adornment f . For each processed subgoal, we add its variables to the list of known bound variables. Initially, only the variables bound in the head pattern are known to be bound.

Example 3.8. Let cl be the clause:

```
1 p(X, Y) :- q(Y, Z), r(X, Y).
```

Given an adornment $\mathbf{a} = \text{bf}$ and a local ordering $\sigma(q, r) = r, q$, then $\text{adorn}(cl, \mathbf{a}, \sigma)$ produces:

```
1 p(X, Y)bf :- r(X, Y)bf, q(Y, Z)bf.
```

where Y is bound in the invocation of q by the preceding invocation of r , which at runtime computes and binds a value to Y .

Definition 3.9 (Reordering functions). A *local reordering function* for a particular clause is a function mapping binding patterns (for the clause head) to orderings for that clause. A *global reordering function* maps clauses to local reordering functions.

We use r to range over local orderings and gr for the global ones.

The purpose of this is to allow different binding patterns of the clause head to imply different reorderings for the clause. In Section 4 and Section 5, we compute reordering functions so that, given a head binding pattern, adorning the reordered clause left-to-right with respect to this pattern yields subgoal binding patterns that are consistent with mode requirements.

Definition 3.10 (Generalised program adornment). Let the function adornProgram take a program Pr , a query clause cl_q , and a global reordering function, gr . The adorned version of the program is generated by invoking adorn on cl_q with a binding pattern (adornment vector) comprising f (free) for each parameter of the clause head, and with the reordering function $gr(cl_q)$.

For each subgoal in the adorned query clause, we generate an adorned version of the subgoal's predicate by applying adorn to the predicate's clauses using the binding pattern given to the subgoal along with the corresponding local reordering from gr . This process is repeated for newly generated adorned clauses until no more clauses can be generated.

An adorned program is equivalent to the original program in the answers it computes [2]. When all the local reordering functions are the identity function (preserving source ordering), then adornProgram is the traditional adornment transformation [17].

Example 3.11. Consider the following program Pr where q is the query, hash is a built-in hash function, and password is an external database predicate to look up a user's password.

```
1 q(H) :- hashByUser("Rebecca", H).
2 hashByUser(U, H) :- password(U, P), hash(P, H).
```

The adorned program is given by $\text{adornProgram}(Pr, cl_q, \text{identity})$, where identity is the trivial global reordering function, mapping every clause to itself for every binding pattern. The output is:

```
1 qf(H) :- hashByUserbf("Rebecca", H).
2 hashByUserbf(U, H) :- passwordbf(U, P), hashbf(P, H).
```

The head adornment of the second line matches the adornment of hashByUser in q 's body as it generates this clause.

3.3 Well-modedness

We define *well-modedness* of clauses and programs in terms of adornment and a notion of *consistency* between mode constraints and adornment vectors.

Definition 3.12 (Mode & adornment consistency). An adornment, \mathbf{a} , is consistent with an atomic constraint, AC , when AC is a set of indices into a indicating bound adornments alone.

$$\mathbf{a} \blacktriangleleft AC \triangleq \forall i \in AC. a_i = b$$

The function $findAC$ selects all atomic constraints in a constraint that are consistent with a given binding pattern:

$$findAC(a, C) \triangleq \{AC \in C \mid a \blacktriangleleft AC\}$$

A binding pattern, a , is consistent with a constraint, C , when there are some atomic constraints in C consistent with a .

$$a \triangleleft C \triangleq findAC(a, C) \neq \emptyset$$

Definition 3.13 (Clause well-modedness). A clause cl is *well-moded* with respect to a constraint function f (mapping predicates to constraints), a binding pattern a , and a local reordering σ , if adornment procedure assigns to each subgoal of cl a binding pattern which is consistent with the moding constraints given by f , i.e.:

$$\begin{aligned} wellModed(cl, a, f, \sigma) &\triangleq \\ \forall sub \in body(adorn(cl, a, \sigma)). adornment(sub) \triangleleft f(pred(sub)) \end{aligned}$$

Definition 3.14 (Program well-modedness). A program Pr with a goal clause cl_q is *well-moded* with respect to a constraint function f and a global reordering function gr , if every subgoal in the adorned program has a binding pattern consistent with the constraints in f :

$$\begin{aligned} wellModedProgram(Pr, cl_q, f, gr) &\triangleq \\ \forall cl \in adornProgram(Pr, cl_q, gr), \\ \forall sub \in body(cl). adornment(sub) \triangleleft f(pred(sub)) \end{aligned}$$

This definition of well-modedness permits ordering-based transformation of clauses as well as retaining multiple versions of the same clause (with different subgoal orderings).

Somogyi [15] noted that modes generalise adornments. This is indeed the case for PROLOG which was the subject of their work. For DATALOG, however, adornment precisely formalises well-modedness because DATALOG does not deal with function symbols as PROLOG does. Hence, a variable can only be instantiated to a value at the time of subgoal invocation or not at all, whereas in PROLOG context, it is possible to partially instantiate variables, e.g. a list of variables, which calls for finer-grained moding instead of binary adornments to fully express the dataflow behaviour of predicates.

3.4 Properties of consistency

We briefly cover several results on the definition of mode consistency which will be of use in later results.

LEMMA 3.15 (ILL AND TRIVIAL CONSTRAINTS). *The trivial constraint, $\{\emptyset\}$, is consistent with all binding patterns whilst the ill constraint, \emptyset , is consistent with none.*

We define a partial order on constraints and functions that output constraints (constraint functions). Later, we establish monotonicity of various functions and operators to prove termination and incrementality of the analysis in Theorem 5.11 and Theorem 5.15.

Definition 3.16. Let C_1 and C_2 be two constraints for the same predicate. We define a relation \leq and say C_1 is less restrictive than C_2 , if every binding pattern that is consistent with C_2 is also consistent with C_1 . The relation \leq is the pointwise extension of \leq for constraint functions, thus:

$$C_1 \leq C_2 \triangleq \forall a. a \triangleleft C_2 \implies a \triangleleft C_1 \quad f \leq g \triangleq \forall p. f(p) \leq g(p)$$

e.g., $\{\{1\}, \{2\}\} \leq \{\{1, 2\}\}$ since the left constraint indicates either the first or second argument needs to be bound, but the right constraint indicates both must be bound (more restrictive).

LEMMA 3.17. *For a fixed predicate, \leq is a bounded partial order with \emptyset as the top element (most restrictive) and $\{\emptyset\}$ as the bottom element (least restrictive). For a fixed domain, \leq is also a bounded partial order with constant functions returning \emptyset and $\{\emptyset\}$ as top and bottom elements respectively.*

The partial order defined by binding pattern consistency implies a subset relation between atomic constraints of two constraints.

LEMMA 3.18. *If a constraint C_2 is more restrictive than C_1 , this means C_1 has a less restrictive atomic constraint for each atomic constraint in C_2 . That is:*

$$\forall C_1, C_2. C_1 \leq C_2 \implies \forall AC \in C_2 \exists AC' \in C_1. AC' \subseteq AC$$

4 INTRA-CLAUSAL ANALYSIS

We start mode analysis by considering the individual clauses of a predicate in isolation, deriving a moding constraint (Def. 3.1) for a clause in terms of its body and the constraints of its subgoals alone.

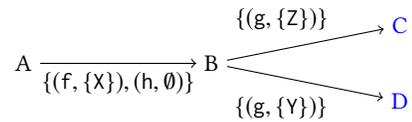
One aim of this analysis is to perform better than brute-force search. We do this by following the path of least resistance, that is: as soon as we find some subgoal that can be scheduled without any constraints, we commit to it. This may discard some valid orderings of clauses but always calculates at least one valid ordering that makes the program well-moded, if such an ordering exists. This analysis is performed by a graph construction.

Before explaining this construction, we first characterise a more general graph structure and explain how orderings are stored in it.

4.1 Scheduling graphs

A scheduling graph encodes orderings of a clause's subgoals. We work towards its formal definition and then explain how orderings can be retrieved from it. We introduce scheduling graphs prior to the specific construction used in the analysis for two reasons. First, it assists in our proof of completeness (Lemma 4.31). Second, it provides a framework that abstracts most details of the construction and allows us to focus on subgoal choices alone.

Fig. 1 shows an example scheduling graph for a clause r . Paths of a scheduling graph capture possible orderings of subgoals in a clause. For example, the path ABC represents two subgoal orderings f, h, g and h, f, g . Nodes comprise the subgoals yet to be scheduled with their dataflow requirements and a set of variables to be resolved at the call site, e.g., node B contains the unscheduled subgoal g and X to be resolved by the caller of the clause. We refer back to this example in the following definitions, explaining its details.



$$A = (\{(f, \{\{X\}\}), (g, \{\{X, Z\}, \{Y\}\}), (h, \{\})\}, \emptyset)$$

$$B = (\{(g, \{\{Z\}, \{Y\}\}), \{X\}\} \quad C = (\emptyset, \{X, Z\}) \quad D = (\emptyset, \{X, Y\})$$

$$\text{Clause: } r(X, Y, Z) :- f^+(X), g^{\{+?+, ?+\}}(X, Y, Z), h(W).$$

Figure 1: Example scheduling graph for a clause of r

Part of the construction requires a way to translate between constraints, given in terms of argument positions of predicates, and a clause context which contains variables.

Definition 4.1 (Obligation). An *obligation* (as opposed to a constraint) is a set of variables that a subgoal is constrained upon. The empty set is the *trivial obligation*.

In Fig. 1, the subgoal $g(X, Y, Z)$ is constrained in either its first and third arguments or just its second argument therefore $\{\{1, 3\}, \{2\}\}$ is its constraint and $\{X, Z\}$ and $\{Y\}$ are the corresponding *obligations*. We use two functions to translate between sets of obligations and constraints: for an atomic formula p , $osToC_p$ maps a set of obligations to a constraint and $cToOs_p$ maps in the opposite direction.

Definition 4.2 (Scheduling graph). A *scheduling graph* g for a clause cl and a constraint function f (over predicates in cl) is a *directed acyclic graph* with a set of vertices $V(g)$ and edges $E(g)$.

An edge is a triple of a source vertex, label, and destination vertex. The label is a non-empty set of pairs of a subgoal from the body of cl and a discharged obligation. In Fig. 1, the edge between B and C has the subgoal g in its label along with the obligation $\{Z\}$.

We assume helper functions src , $label$, and dst to access components of edges. The $paths$ function give paths of the graph.

A vertex is a tuple of the form (Alt, Acc) . The set Alt stores *alternatives*: tuples of the form (s, Obg) where s is a subgoal of the clause and Obg is a set of obligations. The alternatives represent what can be scheduled after a given point in the graph. The set of obligations Obg , associated with the subgoal, represent the variables that can be bound to satisfy the moding constraints. To put it another way, they are the cost of scheduling the subgoal as the caller of the clause must bind these variables. In the running example, at vertex A all subgoals can be scheduled and f and h get scheduled first.

The second component of the vertex, Acc , is an accumulated obligation, keeping track of the variables that have to be bound at the head of the clause. In other words, it keeps track of the total cost for the path. In Fig. 1, this cost is $\{X\}$ at vertex B because the preceding edge is labelled with f , which constrains X . As the head of the clause needs to bind these variables, the accumulator has to be a subset of the head variables, *i.e.*, $Acc \subseteq vars(head(cl))$.

Every edge in a scheduling graph is of the form:

$$(Alt, Acc) \xrightarrow{l} (Alt', Acc')$$

and the vertices and the edge satisfy the following properties which are jointly referred to as the *valid scheduling* property:

- (1) The accumulator of the target vertex extends that of the source with the obligations in the label:

$$Acc' = nextAcc(Acc, l)$$

$$where \ nextAcc(Acc, l) \triangleq Acc \cup \bigcup obgs(l)$$

with $obgs(l) \triangleq \{o \mid (s, o) \in l\}$. Thus scheduling adds a binding requirement to be resolved in the head of the clause.

- (2) Alternatives in the target are computed from source alternatives:

$$Alt' = nextAlt(Alt, l)$$

where $nextAlt(Alt, l) \triangleq \{(s, Obg) \in Alt \mid s \notin subs(l)\} \ominus \bigcup_{s \in subs(l)} vars(s)$

with $subs(l) \triangleq \{s \mid (s, o) \in l\}$. The set of alternatives in the target has the labelling subgoals removed and variables in the subgoals that are marked as discharged by the label are *released*, by the operator \ominus :

$$A \ominus Vars \triangleq \{(s, \min \{o \setminus Vars \mid o \in Obg\}) \mid (s, Obg) \in A\}$$

The release operator removes the variables $Vars$ from the obligations of the alternatives in A , and minimises the set of obligations to remove any redundancies.

Consider the edge between A and B in Fig. 1. The only alternative subgoal in B is g as f and h appear in the edge label. Further, at A , the set of obligations for g includes $\{X, Z\}$, but at B the corresponding set of obligations contains only $\{Z\}$ because X is an argument of f which was scheduled before B .

- (3) Edge labels are picked from the preceding vertex's alternatives:

$$\forall (s, o) \in l, \exists (s', Obg) \in Alt. s = s' \wedge o \in Obg$$

Finally, a scheduling graph has a root vertex, $Root_{cl, f}$. It is a transcription of the predicate constraint into the clause context coupled with an empty accumulator.

$$Root_{cl, f} \triangleq (\{(s, \min cToOs_s(f(pred(s)))) \mid s \in body(cl)\}, \emptyset)$$

Although the constraints being translated are minimal by definition (Def. 3.1), we minimise the sets of obligations in the alternatives after translation because two argument positions may point to the same variable, hence creating a subset relation that did not exist between atomic constraints. For example, a constraint $\{\{1, 2\}, \{2, 3\}\}$ for a subgoal $p(X, X, Y)$ produces the alternative obligation $\min\{\{X\}, \{X, Y\}\}$, thus $\{\{X\}\}$.

If a path in a scheduling graph has labels covering every subgoal of a clause, then the path represents a clause ordering. Such paths are characterised by having a terminal vertex:

Definition 4.3 (Terminal vertex). A terminal vertex is of the form (\emptyset, Acc) for some Acc . We can check if a path has such a vertex:

$$terminal(p) \triangleq \exists Acc. (\emptyset, Acc) \in V(p)$$

A terminal vertex must be the last vertex on a path since the alternative set size is strictly decreasing. Consequently, all subgoals are scheduled in a path with a terminal vertex.

Paths from the root vertex to a terminal vertex represent subgoal orderings. Since edge labels may have multiple subgoals, paths in the graph form *compact orderings*. Each edge expands to all permutations of its members. Concatenating the resulting ordering fragments leads to full clause orderings. For example, a sequence of subgoals in labels $\{p\}, \{q, r\}, \{s, t\}$ leads to orderings mapping the syntactic order of the subgoals to the following:

$$p, q, r, s, t \quad p, r, q, s, t \quad p, q, r, t, s \quad p, r, q, t, s$$

Using these definitions, we define *well-modedness* of paths.

Definition 4.4 (Well-moded path). A path p in a scheduling graph constructed for clause cl is *well-moded* with respect to a binding pattern a and a constraint function f when all orderings in the path make the clause well-moded. The path must be terminal as the adornment procedure needs to know where to place every subgoal.

$wellModedPath(p, cl, a, f) \triangleq$

$$terminal(p) \wedge \forall \sigma \in orderings(p). wellModed(cl, a, f, \sigma)$$

In Fig. 1, the path ABC is well-moded with respect to the binding pattern fbf as all orderings stored in this path require the first and the third arguments of the head to be bound. However, ABD is not well-moded with respect to the same pattern as the path requires the second argument to be bound but the pattern marks it as free.

4.2 Minimal obligation graphs

A valid scheduling graph is determined solely by the subgoal scheduling choices at each edge. Here, we describe the construction of a particular scheduling graph (for some clause) referred to as a *minimal obligation graph* (MOG). Such a graph is *greedy* (Def. 4.10) by construction, meaning it schedules subgoals with trivial obligations as soon as possible. It is minimal because it stores just enough information to derive from each path a minimal set of obligations with respect to subset inclusion. These minimal sets are the dataflow-requirement summaries arising from subgoal orderings.

Definition 4.5 (Minimal obligation graph). For a clause, cl , and a constraint function, f , over predicates in cl , let a *minimal obligation graph*, $\text{mog}_{cl,f}$, be a scheduling graph.

We construct $\text{mog}_{cl,f}$ in a breadth-first manner. We start with a complete example of a MOG construction and use this to expand on the algorithm.

Example 4.6. Consider the following mode-annotated clause cl_r :
 $r(Y, Z) :- f^+(X), g^{++?, +?+}(X, Y, Z), h^+(Z), i(X), j(X, W)$.
 The moding annotations induce a constraint function f mapping predicate names to constraints defined:

$$f(f) = \{\{1\}\} \quad f(g) = \{\{1, 2\}, \{1, 3\}\} \quad f(h) = \{\{1\}\} \quad f(i) = f(j) = \{\emptyset\}$$

Fig. 2 then shows the MOG constructed by $\text{mog}_{cl_r, f}$.

MOGs are constructed iteratively starting from a root node. We write $\text{mog}_{cl,f}^i$ for the i^{th} step of this process.

Base case. All MOGs stem from a graph of a root vertex (as in Def. 4.2) and no edges, defined:

$$V(\text{mog}_{cl,f}^0) = \text{Root}_{cl,f} \quad E(\text{mog}_{cl,f}^0) = \emptyset$$

Thus $\text{mog}_{cl,f}^0$ initiates the construction of the graph by adding no edges and one root. In the running example, A is the root vertex. The alternatives set of the root is a translation from predicate constraints to the clause context (Def. 4.2). For example, g has the constraint $\{\{1, 2\}, \{1, 3\}\}$ and thus $g(X, Y, Z)$ is represented by $\{\{X, Y\}, \{X, Z\}\}$. The accumulator component of the vertex is empty.

Inductive step. We expand the MOG one unit distance at a time. At distance $n + 1$, vertices come from the destination of the edges at the same distance. Edges are added for each vertex at distance n .

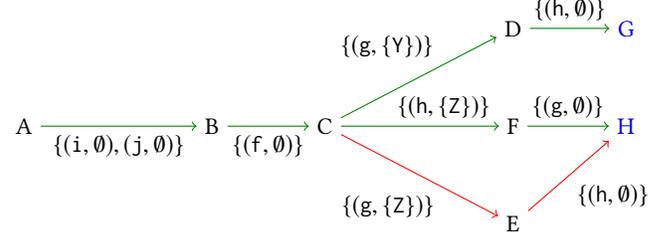
$$V(\text{mog}_{cl,f}^{n+1}) = \{dst \mid (src, l, dst) \in E(\text{mog}_{cl,f}^{n+1})\}$$

$$E(\text{mog}_{cl,f}^{n+1}) = \bigcup_{v \in V(\text{mog}_{cl,f}^n)} \{mkEdge(v, l) \mid l \in \text{pickLabel}_{cl}(\pi_1 v)\}$$

where $mkVertex$ is defined using scheduling graph primitives (Def. 4.2):

$$mkEdge(src, l) \triangleq (src, l, mkVertex(src, l))$$

$$mkVertex((Alt, Acc), l) \triangleq (nextAlt(Alt, l), nextAcc(Acc, l))$$



$$A = (\{(f, \{X\}), (g, \{X, Y, \{X, Z\}), (h, \{Z\}), (i, \{\emptyset\}), (j, \{\emptyset\}), \emptyset\})$$

$$B = (\{(f, \{\emptyset\}), (g, \{\{Y\}, \{Z\}), (h, \{\{Z\}\}), \emptyset\})$$

$$C = (\{(g, \{\{Y\}, \{Z\}), (h, \{\{Z\}\}), \emptyset\})$$

$$D = (\{(h, \{\emptyset\}), \{Y\}\}) \quad E = (\{(h, \{\emptyset\}), \{Z\}\}) \quad F = (\{(g, \{\emptyset\}), \{Z\}\})$$

$$G = (\emptyset, \{Y\}) \quad H = (\emptyset, \{Z\})$$

Clause:

$$r(Y, Z) :- f^+(X), g^{++?, +?+}(X, Y, Z), h^+(Z), i(X), j(X, W)$$

Figure 2: Minimal obligation graph for a clause of r

From a given vertex, we decide which subgoals to schedule (which edge label to generate) using pickLabel . We give preference to scheduling of “natural” subgoals within the alternative set over the others. We explain what each case of these label choices means.

$$\text{pickLabel}_{cl}(Alt) \triangleq \begin{cases} \{nats(Alt)\} & nats(Alt) \neq \emptyset \\ nonNats_{cl}(Alt) & \text{otherwise} \end{cases}$$

Case 1: Extending the graph with natural edges. The *natural subgoals* at a vertex are those subgoals in the alternatives set that are paired with a singleton set of the trivial obligation. This means that the binding requirements of this subgoal are satisfied at this point.

$$nats(Alt) \triangleq \{(sub, \emptyset) \mid (sub, \{\emptyset\}) \in Alt\}$$

In the example, applying $nats$ to $\pi_1 A$ yields $\{(i, \emptyset), (j, \emptyset)\}$. This set is collectively used by $mkEdge$ as a label. Thus, the subgoals i and j are scheduled between the nodes A and B .

One consequence of natural subgoals all having the trivial obligation in the edge label is that $nextAcc$ does not add new obligations to the accumulator in the generated destination vertex. In the example, this is why the accumulator of B remains as the empty set.

Subgoal naturality is contextual. In the example, f is not natural at vertex A , but it becomes natural at B due to edge AB releasing the variables in i and j . Thus, BC schedules f as a natural subgoal.

Case 2: Extending the graph with non-natural edges. If there are no natural subgoals to schedule, the pickLabel function tries alternatives with non-trivial obligations. Unlike the natural case, labels picked in this manner comprise a single subgoal-obligation pair, formally expressed as follows:

$$nonNats_{cl}(Alt) \triangleq \{(s, o) \mid (s, os) \in Alt, o \in os, o \subseteq \text{vars}(\text{head}(cl))\}$$

To select a label this way, we check if the obligation is a subset of the head variables. This makes binding the variables in the head possible and satisfies the valid scheduling property (Def. 4.2).

In the example, all edges from C are created in this manner as all the alternatives in C have non-empty obligations.

Unlike the natural case, *nextAcc* may augment the accumulator in the destination vertex since labels of this kind have non-empty obligations; the caller must satisfy these obligations. In the example, vertices D , F , and E are generated using this function and since at C the accumulator is empty, the accumulator in these three vertices is the single obligation specified by the label of the incoming edge.

Building the whole graph. The full $\text{mog}_{cl,f}$ graph is given as a union of its components:

$$\mathbf{V}(\text{mog}_{cl,f}) = \bigcup_{n \geq 0} \mathbf{V}(\text{mog}_{cl,f}^n) \quad \mathbf{E}(\text{mog}_{cl,f}) = \bigcup_{n \geq 0} \mathbf{E}(\text{mog}_{cl,f}^n)$$

LEMMA 4.7. *MOG construction leads to a scheduling graph.*

LEMMA 4.8 (MOG TERMINATION). *For any clause and constraint function, the MOG construction terminates.*

The only scenario in which a clause cannot lift its constraints to the caller is if its subgoals have variables that do not appear in the head. MOG construction gets stuck in this case, leading to a graph with no terminal vertices.

Example 4.9 (Stuck construction). Consider the following clause with the given moding requirements:

$$\text{r}(X) :- \text{g}^+(Y), \text{f}^+(X).$$

These moding requirements cannot be satisfied by any ordering of subgoals as the variable Y cannot be bound by just binding X . This is reflected by the MOG for this clause, which fails to contain any paths that schedule all the subgoals of the clause.

The root vertex contains the alternatives: $\{(f, \{\{X\}\}), (g, \{\{Y\}\})\}$. Since none of these are natural subgoals, *pickLabel* selects a non-natural edge which requires augmenting the accumulator. Only one edge is possible since only one obligation has all of its variables in the head of the clause, namely f . This yields the MOG:

$$((f, \{\{X\}\}), (g, \{\{Y\}\}), \emptyset) \xrightarrow{\{(f, \{X\})\}} ((g, \{\{Y\}\}), \{X\})$$

Further edges cannot be generated (i.e., $\mathbf{V}(\text{mog}_{cl,f}^2) = \mathbf{E}(\text{mog}_{cl,f}^2) = \emptyset$) and the MOG has no terminal vertices. Hence, there are no orderings making this clause well-moded with respect to its constraints.

MOGs are “greedy” scheduling graphs. A path is considered greedy if it schedules natural subgoals as soon as possible.

Definition 4.10 (Greedy paths). A path p is said to be *greedy* if for all edges p_i with source vertex v , whenever v has a natural subgoal *sub* (denoted $(\text{sub}, \{\emptyset\})$ in its alternatives) then (sub, \emptyset) is in the label of p_i , or equivalently (due to *valid scheduling*) the subgoal does not appear in the alternatives of the destination vertex:

$$\text{greedy}(p) \triangleq \forall i, s. (s, \{\emptyset\}) \in \pi_1 \text{src}(p_i) \implies \forall \text{Alt}. (s, \text{Alt}) \notin \pi_1 \text{dst}(p_i)$$

PROPOSITION 4.11. *All paths in a MOG are greedy.*

4.2.1 Optimising MOG construction. MOG construction is described here in a breadth-first manner. A depth-first reconstruction is possible which, when accompanied with some global state, may prune the graph. First, we switch to depth first-construction by branching only when we reach a terminal vertex or there is no

possible edge to add to a path. Second, we maintain a global list of the accumulator components of terminal vertices encountered so far. Each time we extend a path (via natural subgoals or otherwise), we check if the accumulator component of the newly created vertex is a superset of any member of the global list. If so, then we do not create that vertex at all as we already have a path that is at least as good as the one we are about to explore.

In Example 4.6, the red edges show the edges that would not have been constructed via this approach. If the terminal vertex G is discovered first, the global list contains $\{Y\}$. Then, the terminal vertex H is discovered via F , hence we add $\{Z\}$ to the global list. When we attempt to explore the third edge from C to reach E , however, we realise the accumulator at E would be $\{Z\}$ which is a superset of one of the elements in the global list. Hence, we do not add E and the two edges from C to H via E to the graph.

The upside of this change is that we have fewer vertices to explore. The downside is, in the end, we will have fewer orderings to choose from for a given binding pattern and which orderings are retained depends on which terminal vertex is encountered first.

4.3 Extracting a clause constraint out of a MOG

The second component of a MOG vertex is an accumulated obligation. Using this, we can derive a new constraint for the given clause. Only terminal vertices contribute to constraint derivation.

Obligations in terminal vertices of a clause are viable options for a caller to resolve. Since variable names in obligations are meaningful only in a clause context, we convert obligations back into singleton constraints represented in terms of parameter positions. The \otimes operator then combines these while eliminating redundancies:

Definition 4.12. If A and B are in the constraint domain D_p , we define $A \otimes B$ as $\min(A \cup B)$.

The union retains alternative moding patterns in each constraint and \min eliminates the redundantly restrictive constraints.

Example 4.13. Consider a mode-annotated clause cl_s :

$$\text{s}(X, Y) :- \text{f}^{++}(X, Y), \text{g}^{+?}(X, Y).$$

The mode annotations induce a constraint function f , then $\text{mog}_{cl_s, f}$:

$$\begin{array}{ccc} \{(f, \{X, Y\})\} & \xrightarrow{\{((g, \{\emptyset\}), \{X, Y\})\}} & \{(g, \emptyset)\} \otimes \{(f, \emptyset)\} \\ \downarrow & & \downarrow \\ ((f, \{\{X, Y\}\}), (g, \{\{X\}\}), \emptyset) & & \{(f, \emptyset)\} \\ \downarrow & & \downarrow \\ \{(g, \{X\})\} & \xrightarrow{\{((f, \{\emptyset\}), \{X\})\}} & \{(f, \emptyset)\} \otimes \{(g, \{X\})\} \end{array}$$

Different subgoal orderings of s have different constraints. If f is first (as in the source), X and Y together are the obligation, hence we constrain both argument positions leading to $\{\{1, 2\}\}$. If g is first, only X is in the obligation as the use of g binds Y and satisfies the moding requirement of f . Thus, this ordering implies the constraint $\{\{1\}\}$. The latter ordering is more favourable as it leads to a less restrictive constraint (Def. 3.16). Hence, \otimes eliminates the former:

$$\{\{1, 2\}\} \otimes \{\{1\}\} = \min(\{\{1, 2\}\} \cup \{\{1\}\}) = \{\{1\}\}$$

Definition 4.14 (Constraint generated from a MOG). We derive an overall constraint for the clause of a MOG by translating the

accumulators (obligations) of all terminal vertices into constraints, combining these via \otimes . This operation is called *extract*:

$$\text{extract}_{cl}(\text{mog}) \triangleq \bigotimes_{(\emptyset, o) \in \mathcal{V}(\text{mog})} \text{osToC}_{\text{head}(cl)}(\{o\})$$

Example 4.15. In the running example from Figure 2, the terminal vertices are $G = (\emptyset, \{Y\})$ and $H = (\emptyset, \{Z\})$ therefore:

$$\text{extract}_{cl_r}(\text{mog}_{cl_r, f}) = \{\{1\}\} \otimes \{\{2\}\} = \{\{1\}, \{2\}\}$$

i.e., either the first or the second parameter of r should be bound.

LEMMA 4.16. *Composition of the extract function with the MOG construction is monotonically increasing.*

$$\forall cl, f, g. f \leq g \implies \text{extract}_{cl}(\text{mog}_{cl, f}) \leq \text{extract}_{cl}(\text{mog}_{cl, g})$$

4.4 Extracting subgoal orders out of a MOG

Now that we have a constraint for the clause, we build a partial function mapping from binding patterns to subgoal orderings of the clause. This function is partial as, for some patterns, no ordering of subgoals leads to a well-moded clause. Thus, we have orderings only for binding patterns that are consistent with the clause constraint.

We define a function $\text{collect}_{\text{mog}}$, parameterised by a MOG, which maps atomic constraints to sets of subgoal orderings. Orderings are computed in two steps. First, $\text{collect}_{\text{mog}}(AC)$ finds all paths from the root to a terminal vertex whose accumulated obligation translates to the given atomic constraint AC . Second, the compact orderings represented by these paths are expanded into subgoal orderings.

In general, a MOG may provide multiple valid orderings for a binding pattern as there may be multiple distinct atomic constraints consistent with the binding pattern. We assume a function, choose , that selects from a set of orderings. This function can be based on a metric defined on orderings. For example, it may use a metric that measures the distance from the original source ordering, choosing the one that is closest. This is useful for preserving some optimisations based on estimated relation sizes.

Equipped with these definitions, we build the desired ordering (partial) function for a given clause cl and constraint function f :

$$\text{reordering}_{cl, f}(\mathbf{a}) \triangleq \text{choose}\left(\bigcup_{AC \in C} \text{collect}_{\text{mog}_{cl, f}}(AC)\right)$$

where $C = \text{findAC}(\mathbf{a}, \text{extract}_{cl}(\text{mog}_{cl, f}))$

4.5 Results

We present useful properties of intra-clausal analysis: propagation of ill constraints, soundness, and completeness. All of these are generalised to inter-clausal analysis in Section 5.1.

LEMMA 4.17. *If any predicates in the body of a clause cl has the ill constraint according to the constraint function f , the constraint for the clause as extracted from its MOG is also ill.*

$$\forall sub \in \text{body}(cl). f(\text{pred}(sub)) = \emptyset \implies \text{extract}_{cl}(\text{mog}_{cl, f}) = \emptyset$$

We now formalise soundness, *i.e.*, when a clause constraint is consistent with the binding pattern of the head formula, the reordered clause subgoals are dataflow safe.

LEMMA 4.18 (INTRA-CLAUSAL SOUNDNESS). *For a given clause, cl , and a constraint function, f , if a binding pattern is consistent with the constraint from the intra-clausal analysis, then the binding patterns of all the subgoals of the adorned clause (with the reordering function from the intra-clausal analysis) are also consistent with their respective constraints in f .*

$$\forall \mathbf{a}. \mathbf{a} \triangleleft \text{extract}_{cl}(\text{mog}_{cl, f}) \implies \text{wellModed}(cl, \mathbf{a}, f, \text{reordering}_{cl, f}(\mathbf{a}))$$

PROOF. Pick an arbitrary adornment \mathbf{a} that is consistent with the extracted constraint from the MOG, *i.e.*, $\mathbf{a} \triangleleft \text{extract}_{cl}(\text{mog}_{cl, f})$ (see Def. 3.12 for \triangleleft). This implies $\text{reordering}_{cl, f}$ is defined at \mathbf{a} , which can only happen if there is a path p in $\text{mog}_{cl, f}$ leading to this ordering.

We proceed by establishing a contradiction if p represents an unsound path, *i.e.*, there is at least one subgoal on the path with a binding pattern inconsistent with the constraint function f .

Let $s(\mathbf{X})$ to be the earliest subgoal (where \mathbf{X} is a list of variables) that is inconsistent with its constraint after adorning the clause with \mathbf{a} and let p_i be the edge that contains this subgoal in its label. Let \mathbf{b} be that inconsistent binding pattern of s and C be its constraint from f , then we have $\forall AC \in C. \mathbf{b} \not\triangleleft AC$ by definition of consistency. This can only happen if for each atomic constraint there is at least one index j such that \mathbf{X}_j is adorned free because all indices in AC s are bound by definition of \triangleleft . Let F be the set of all such offending variables. Since all of F has to be free at the point s is scheduled, none of F can appear in the head of the clause as bound or as an argument to any subgoals before p_i .

There are two ways s could have been scheduled. Either in an edge with (possibly) other natural subgoals (extension by natural subgoals) or by extending the accumulated obligation (extension by non-natural subgoals).

Consider extension by natural subgoals, requiring an alternative in the source of p_i of the form $(s, \{\emptyset\})$. At the root, the alternatives include $(s, \min c\text{ToOs}_s(C))$. Each obligation in this alternative must have at least one element from F as they are generated from C . This cannot be the case as none of F appeared as an argument before p_i and hence these variables cannot be released with \ominus , thus we cannot obtain the alternative $(s, \{\emptyset\})$ or schedule s at p_i .

Consider extension by non-natural subgoals. As before we know that the alternatives at the source of p_i each contain at least one member of F . This means regardless of which alternative is used, the accumulated obligation at the destination of p_i must contain a member of F . Hence, the terminal vertex's accumulator must contain at least one member of F . All variables in this accumulator must be *bound* in the head due to the fact that the reordering is generated using a terminal vertex with an accumulated obligation corresponding to an atomic constraint of the head. This contradicts our assumption that the offending variable is *free* in the head.

Since having an inconsistent subgoal in the path contradicts the formation of the path, all subgoals must be consistent with their constraints when the head is adorned with \mathbf{a} . \square

Intra-clausal completeness is more involved. It states that:

$$\forall \mathbf{a}. (\exists r. \text{wellModed}(cl, \mathbf{a}, f, r(\mathbf{a}))) \implies \mathbf{a} \triangleleft \text{extract}_{cl}(\text{mog}_{cl, f})$$

(full statement in Lemma 4.31). That is, if there exists a local ordering r that makes a clause well-moded with respect to head binding pattern \mathbf{a} , then \mathbf{a} is consistent with the constraint computed by

our analysis. We prove completeness by showing than an arbitrary ordering r is captured (in some way) by our MOG construction, *i.e.*, that $\text{mog}_{cl,f}$ is complete. This involves first converting arbitrary orderings to scheduling graphs, and showing that paths in such a graph can be transformed into effectively equivalent paths in our MOG. Key to this is Proposition 4.11: that MOG paths are *greedy*.

Definition 4.19 (Ordering to scheduling graph). Given a clause, cl , with n subgoals, a constraint function, f , defined for all predicates invoked in cl , a binding pattern \mathbf{a} , and an ordering σ of cl giving an adorned and reordered clause $cl' = \text{adorn}(cl, \mathbf{a}, \sigma)$, then there is a scheduling graph $g = \text{schedule}(cl, \mathbf{a}, f, \sigma)$, where:

$$\begin{aligned} V(g) &= \bigcup_{0 \leq i < n} V_i & E(g) &= \bigcup_{0 \leq i < n} E_i & V_0 &= \text{Root}_{cl,f} & E_0 &= \emptyset \\ V_i &= \text{dst}(E_i) \\ E_i &= \{ \text{mkEdge}(v, \{(s, o)\}) \mid v \in V_{i-1}, (s, \text{obgs}) \in \pi_1 v, s = cl'_i, \\ &\quad o \in \text{obgs}, \text{adornment}(s) \triangleleft \text{osToC}_s(\{o\}) \} \end{aligned}$$

LEMMA 4.20 (WELL-MODED ORDERING TO TERMINAL PATH). *For a constraint function, f , clause, cl , binding pattern, \mathbf{a} , and ordering, σ , where $\text{wellModed}(cl, \mathbf{a}, f, \sigma)$, a scheduling graph $\text{schedule}(cl, \mathbf{a}, f, \sigma)$ has a terminal path p .*

Definition 4.21 (Conversion). We next convert scheduling graph paths, where each edge has a singleton set label, into greedy paths.

Find the earliest vertex, v , that has a subgoal with the trivial obligation in its alternatives. For all such subgoals, use the *swap* operation (below) to place these subgoals in adjacent edges starting from v in any order. Use *merge* repeatedly to merge all such edges. Repeat this process until the path cannot be changed anymore.

Definition 4.22 (Swap). The *swap* operation on a path p in a scheduling graph takes an index i and assuming p_i and p_{i+1} exist, produces a new path where the subgoals in the edge p_i comes before the subgoals in the edge p_{i+1} . The operation is applied when all the subgoals in the edge p_{i+1} have trivial obligations in $\text{src}(p_i)$ and consequently are natural subgoals scheduled by p_{i+1} .

Let L be the edge label at p_i and R be the edge label at p_{i+1} . The new path produced by *swap* is called q , defined as follows:

$$\begin{aligned} \forall j < i. q_j &= p_j \\ q_i &= \text{mkEdge}(\text{src}(p_i), R) \\ q_{i+1} &= \text{mkEdge}(\text{dst}(q_i), \{(s, o \setminus \bigcup_{(s', o) \in R} \text{vars}(s')) \mid (s, o) \in L\}) \\ \forall j > i + 1. q_j &= \text{mkEdge}(\text{dst}(q_{j-1}), \text{label}(p_j)) \end{aligned}$$

LEMMA 4.23 (TRIVIAL OBLIGATION CONSISTENCY). *If a scheduling graph vertex v has a subgoal s with obligation $\{\emptyset\}$ in its alternative set, any ordering with a partial ordering derived from the root to v makes the predicate s consistent with respect to any binding pattern.*

LEMMA 4.24 (SWAP PRESERVATION). *If a scheduling graph path is well-moded, performing a swap on this path preserves well-modedness and the path still belongs to some scheduling graph.*

Definition 4.25 (Merge). Let p be a path in a scheduling graph, a merge of edges p_i and p_{i+1} removes them both and replaces them

with a single edge with a label that is the union of all the labels. The operation is applied if every subgoal in both of these edges has the trivial obligation at the source of p_i . Let q be the resulting path with the following specification:

$$\begin{aligned} \forall j < i. q_j &= p_j \\ q_i &= (\text{src}(p_i), \text{label}(p_i) \cup \text{label}(p_{i+1}), \text{dst}(p_{i+1})) \\ \forall j > i. q_j &= p_{j+1} \end{aligned}$$

LEMMA 4.26 (MERGE PRESERVATION). *If a scheduling graph path is well-moded, performing a merge on this path preserves well-modedness and the path still belongs to some scheduling graph.*

LEMMA 4.27 (CONVERSION PRESERVATION). *If a scheduling graph path is well-moded, its conversion is also a well-moded path of some scheduling graph.*

LEMMA 4.28 (CONVERSION GREED). *If a path is in a scheduling graph, then its conversion produces a greedy path.*

Thus, we have shown that conversion creates greedy scheduling paths, and preserves well-modedness of paths. We then show such paths are in the MOG of the clause (Lemma 4.29) and are consistent with adornment (Lemma 4.30), finally leading to completeness.

LEMMA 4.29 (GREEDY PATH COMPLETENESS). *For a clause cl , every greedy scheduling path for cl ending in a terminal vertex and conforming to a constraint function f is present in the MOG determined by cl and f . That is:*

$$\forall p, \mathbf{a}. \text{greedy}(p) \wedge \text{wellModedPath}(p, \mathbf{a}, cl, f) \implies p \in \text{paths}(\text{mog}_{cl,f})$$

LEMMA 4.30 (PATH EXTRACT CONNECTION). *For a fixed binding pattern \mathbf{a} , existence of a well-moded path in a MOG implies consistency of \mathbf{a} with the constraint extracted from the MOG.*

$$\begin{aligned} \forall \mathbf{a}. \text{wellModedPath}(p, \mathbf{a}, cl, f) \wedge p \in \text{paths}(\text{mog}_{cl,f}) \\ \implies \mathbf{a} \triangleleft \text{extract}_{cl,f}(\text{mog}_{cl,f}) \end{aligned}$$

LEMMA 4.31 (INTRA-CLAUSAL COMPLETENESS). *For a given clause, cl , a constraint function, f , and an adornment \mathbf{a} for the head of cl , if there is a local reordering that makes the adornment of the subgoals consistent with their constraints, the head adornment is consistent with the constraint extracted from the MOG. That is:*

$$\forall \mathbf{a}. (\exists r. \text{wellModed}(cl, \mathbf{a}, f, r(\mathbf{a}))) \implies \mathbf{a} \triangleleft \text{extract}_{cl}(\text{mog}_{cl,f})$$

PROOF. Fix an arbitrary adornment \mathbf{a} and assume the antecedent. We need to show that $\text{mog}_{cl,f}$ contains a path that ends in a terminal vertex leading to an atomic constraint consistent with \mathbf{a} .

By Def. 4.19, we convert the ordering for \mathbf{a} into a scheduling graph, *i.e.* $g = \text{schedule}(cl, \mathbf{a}, f, r(\mathbf{a}))$. Since, $\text{wellModed}(cl, \mathbf{a}, f, r(\mathbf{a}))$, there exists at least one path p which is terminal in g (Lemma 4.20). From Def. 4.21 (with Lemma 4.28) we convert this path p into a greedy path p' , which is terminal and well-moded (Lemma 4.27, Lemma 4.28). Since the path p' is well moded and terminal we have that, $\text{wellModedPath}(p, \mathbf{a}, cl, f)$ (Def. 4.4).

Using Lemma 4.29 (greedy path completeness), it then follows that $p' \in \text{paths}(\text{mog}_{cl,f})$, *i.e.*, that p' is constructed by our MOG-based analysis. Combined with Lemma 4.30 (well-moded MOG path implies extract consistency) then $\mathbf{a} \triangleleft \text{extract}_{cl,f}(\text{mog}_{cl,f})$. \square

5 INTER-CLAUSAL ANALYSIS

Having defined how to determine a moding constraint for a given clause, we are ready to find constraints for a program.

Any clause of a predicate can be used to evaluate a subgoal involving that predicate. Consequently, for a subgoal invocation to be safe, the bodies of all such clauses must be safe to evaluate. Therefore, the constraint of a predicate must jointly reflect the constraints of its clauses, which we capture with the \oplus operator:

Definition 5.1. If A and B are in the constraint domain D_p , we define $A \oplus B$ as $\min \{a \cup b \mid a \in A, b \in B\}$.

This captures joint constraints because it produces an atomic constraint for each pair of atomic constraints and union preserves restrictions. As with the \otimes operator, \min eliminates redundancies.

Example 5.2. Consider the following clauses for a predicate r :

- 1 $r(X, Y, Z) :- f^{+??}(X, Y, Z).$
- 2 $r(X, Y, Z) :- g^{\{++?, ?++\}}(X, Y, Z).$
- 3 $r(X, Y, Z) :- h^{??+}(X, Y, Z).$

Use of r must reflect the constraints of all these clauses. Individually, the constraints for each clause are $\{\{1\}\}$, $\{\{1, 2\}, \{2, 3\}\}$, and $\{\{3\}\}$ respectively. The only way these three constraints are satisfied is if all three arguments are constrained; \oplus computes this:

$$\{\{1\}\} \oplus \{\{1, 2\}, \{2, 3\}\} \oplus \{\{3\}\} = \min \{\{1, 2\}, \{1, 2, 3\}\} \oplus \{\{3\}\} = \{\{1, 2\}\} \oplus \{\{3\}\} = \{\{1, 2, 3\}\}$$

LEMMA 5.3 (\oplus CONSISTENCY HOMOMORPHISM). *A binding pattern is consistent with two constraints combined with \oplus iff that binding pattern is consistent with each constraint individually, i.e.:*

$$\forall a, C_1, C_2. a \triangleleft (C_1 \oplus C_2) \iff a \triangleleft C_1 \wedge a \triangleleft C_2$$

Definition 5.4 (Whole-program analysis). Whole-program analysis is a fixpoint computation over constraint functions. In each iteration, the constraints of the clauses with a shared head are combined with the \oplus operator. Constraints of predicates without clauses, i.e., $Pr_p = \emptyset$, (such as built-in predicates) are preserved.

$$\text{analyse}_{Pr}(f) \triangleq \begin{cases} f & \text{if } f = \text{step}_{Pr}(f) \\ \text{analyse}_{Pr}(\text{step}_{Pr}(f)) & \text{otherwise} \end{cases}$$

$$\text{step}_{Pr}(f)(p) \triangleq \begin{cases} \bigoplus_{cl \in Pr_p} \text{extract}_{cl}(\text{mog}_{cl, f}) & \text{if } Pr_p \neq \emptyset \\ f(p) & \text{otherwise} \end{cases}$$

In *step*, we use the \oplus operator over constraints for each clause belonging to predicate p , extracted from the MOG by *extract* (Def. 4.14) which is defined in terms of \otimes over constraints.

We establish some monotonicity properties of inter-clausal analysis used later in termination (Theorem 5.11) and incrementality (Theorem 5.15) proofs.

LEMMA 5.5. *Every constraint function generated from a moding function gets more restrictive when *step* is applied to it.*

$$\forall Pr, mv. \llbracket mv \rrbracket_F \leq \text{step}_{Pr}(\llbracket mv \rrbracket_F)$$

LEMMA 5.6 (STEP MONOTONICITY). *The *step* function is monotonically increasing (making constraints more restrictive):*

$$\forall Pr, f, g. f \leq g \implies \text{step}_{Pr}(f) \leq \text{step}_{Pr}(g)$$

LEMMA 5.7. *Every constraint function generated from a moding function gets more restrictive by application of *analyse*.*

$$\forall Pr, mv. \llbracket mv \rrbracket_F \leq \text{analyse}_{Pr}(\llbracket mv \rrbracket_F)$$

The \oplus operation must be closed on constraints, which follows from \oplus and \otimes forming a semiring:

PROPOSITION 5.8. *For each predicate p , $(D_p, \oplus, \otimes, \{\emptyset\}, \emptyset)$ is an idempotent commutative semiring² where $\{\emptyset\}$ and \emptyset are the additive and multiplicative identities respectively.*

Generalised program adornment (Def. 3.10) requires a higher-order reordering function from clauses to functions that map binding patterns to orderings. This can be constructed as an amalgamation of local reordering functions produced in Section 4.4.

$$\text{reorderProgram}_f(cl) \triangleq \text{reordering}_{cl, f}$$

5.1 Results

We now show that inter-clausal analysis (our full analysis) is fast-failing, terminating, sound, and complete.

Fast failure for ill-moded predicates is a strength of our analysis. Ill-moded constraints quickly propagate via the *step* function.

PROPOSITION 5.9 (FAST FAILURE). *After a single application of *step*, the ill constraint propagates from the body of a clause to the entire head predicate constraint.*

$$\forall f, cl, s \in \text{body}(cl).$$

$$f(\text{pred}(s)) = \emptyset \implies \text{step}_{Pr}(f)(\text{pred}(\text{head}(cl))) = \emptyset$$

COROLLARY 5.10. *The number of *step* applications it takes to converge to the ill constraint is bounded by the static call distance between two predicates.*

THEOREM 5.11. *For all DATALOG programs, Pr , and mode functions, mv , inter-clausal analysis, $\text{analyse}_{Pr}(\llbracket mv \rrbracket)$, terminates.*

PROOF. We know that *step* function terminates because intra-clausal analysis is a function of MOG construction which terminates by Lemma 4.8. All there is left to show is that *analyse* always reaches a fixpoint. This is the case as *step* forms a chain (Lemma 5.6, Lemma 5.5) and \leq is bounded (Lemma 3.17). \square

THEOREM 5.12 (INTER-CLAUSAL SOUNDNESS). *Given a program Pr containing a query cl_q with head predicate q and a mode function mv , if the analysis yields the trivial constraint for q , Pr is well-moded with respect to query clause cl_q and mode function mv .*

$$\forall af. af = \text{analyse}_{Pr}(\llbracket mv \rrbracket_F) \wedge af(q) = \{\emptyset\}$$

$$\implies \text{wellModedProgram}(Pr, cl_q, \llbracket mv \rrbracket_F, \text{reorderProgram}_{af}(cl))$$

PROOF. Fix an arbitrary af and assume the antecedent. Let f be the constraint functions $\llbracket mv \rrbracket_F$.

We know by (Lemma 5.7) $f \leq af$, that is, for any predicate p and binding pattern a , if a is consistent with $af(p)$, then it is also consistent with $f(p)$.

Recall that $\{\emptyset\}$ is the trivial constraint with which all binding patterns are consistent. This means under all binding patterns, there is an ordering for the subgoals of cl_q , where the binding patterns

²It is known as Martelli's semiring, originally used to compute cut-sets of a graph [11].

derived for the subgoals are consistent with the constraints of af and by \leq also with f . Otherwise, we would contradict intra-clausal soundness (Lemma 4.18).

The predicate constraints of subgoals in the body of cl_q may arise from two sources. If the predicate in question is external, we know by assumption it is consistent with af and hence with f . If it is an internal predicate, then it is a combination of clause constraints via \oplus . We know by Lemma 5.3 that each of the clause constraints are consistent with the binding pattern given to the subgoal in cl_q . These clause constraints can only be generated by intra-clausal analysis. We can apply the same reasoning recursively to the body of these clauses to show that all constraints of the external predicates in the body are satisfied. Hence, all external predicate constraints according to f are satisfied as required.

The reason we need a fixpoint rather than a single application of $step$ is that intra-clausal soundness ensures soundness with respect to the input constraint function but intra-clausal analysis potentially produces a more restrictive constraint function due to (mutual) recursion of clauses. At the fixpoint, the output constraint function and the input constraint function are one and the same, hence all dataflow constraints are satisfied. \square

THEOREM 5.13 (INTER-CLAUSAL COMPLETENESS). *Given a program Pr containing a query cl_q with head predicate q and a mode function mv , if there exists a reordering for clauses in the program that produces a well-moded program with respect to the query, then the analysis yields the trivial constraint for the query predicate, q (i.e., our analysis finds a reordering):*

$$\begin{aligned} \exists gr. \text{ wellModedProgram}(Pr, cl_q, \llbracket mv \rrbracket_F, gr) \\ \implies \text{analyse}_{Pr}(\llbracket mv \rrbracket_F)(q) = \{\emptyset\} \end{aligned}$$

PROOF. We proceed with a proof-by-contradiction: we start by assuming that the antecedent is true and that $\text{analyse}_{Pr}(\llbracket mv \rrbracket_F)(q) \neq \{\emptyset\}$. Since a query only has one clause, \oplus is never invoked in analyse_{Pr} for q and thus each $step$ only ever extracts a constraint for the query from a single MOG. For the constraint to end up non-trivial, the MOG construction for cl_q must on all paths generate a non-trivial obligation in the accumulator (since a trivial obligation dominates via the definition of \otimes). Therefore, we must need to bind a variable in the head of the query q . However, by the antecedent and the definition of $adornProgram$, there is an ordering gr such that all variables in the head of query can be adorned with f (free) and all body clauses are consistent with $\llbracket mv \rrbracket_F$. This is a contradiction. Therefore, the statement of completeness holds. \square

COROLLARY 5.14 (GLOBAL REORDERING EXISTENCE). *For every program that can be well-moded with reordering, we can construct a global reordering function.*

Our analysis algorithm preserves the work that has been done on a program if it is extended by additional clauses.

THEOREM 5.15 (INCREMENTAL ANALYSIS). *For a given program, Pr , an arbitrary clause, cl , and a mode function, mv , defined on all predicates appearing in Pr and cl , inter-clausal analysis can be incrementally computed by computing a constraint function for Pr first and using this as a basis for computing constraints for $Pr \cup \{cl\}$.*

$$\text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) = \text{analyse}_{Pr \cup \{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F))$$

PROOF. By (Lemma 5.7), we have $\llbracket mv \rrbracket_F \leq \text{analyse}_{Pr}(\llbracket mv \rrbracket_F)$. We also have the following inequality:

$$\text{analyse}_{Pr}(\llbracket mv \rrbracket_F) \leq \text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F)$$

This holds because if the head of cl is not a head in Pr , the constraint of the head of cl is $\{\emptyset\}$ which is the bottom element of the constraint domain. If it appears as a head, this means at each application of $step$ there will be an additional constraint that needs to be combined using \oplus . We have $C_1 \leq C_1 \oplus C_2$ by Lemma 5.3 and Def. 3.16.

By applying $\text{analyse}_{Pr \cup \{cl\}}$ to both inequalities we obtain:

$$\begin{aligned} \text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) \\ \leq \text{analyse}_{Pr \cup \{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F)) \\ \leq \text{analyse}_{Pr \cup \{cl\}}(\text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F)) \end{aligned}$$

Additionally, analyse reaches a fixpoint (Theorem 5.11), thus

$$\begin{aligned} \text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) \\ \leq \text{analyse}_{Pr \cup \{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F)) \\ \leq \text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) \end{aligned}$$

Since \leq is anti-symmetric, then from the above inequality (of the form $x \leq y$ and $y \leq x$) the equality of the lemma holds. \square

We achieve a stronger incremental computation result if the clause extending the program has a fresh head. It allows us to perform the analysis without performing intra-clausal analysis on the original program.

COROLLARY 5.16. *When the head predicate of cl does not feature in Pr , the clauses of Pr can be ignored during analysis. That is:*

$$\text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) = \text{analyse}_{\{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F))$$

A stronger result still is achieved when the extending clause is non-recursive. Despite the restrictions on the clause, this captures queries in an interactive system. We can determine well-modedness of a query using a single application of the intra-clausal analysis.

COROLLARY 5.17 (FAST CONVERGENCE). *If cl is also non-recursive, then analyse converges to a fixpoint in a single step:*

$$\text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) = \text{step}_{\{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F))$$

6 EXTENDING DATALOG

Our analysis so far accommodates external predicates without side-effects. However, most implementations extend DATALOG. In this section, we explore these extensions in relation to our algorithm. In particular, we discuss preserving effectful-predicate order, negated subgoals, and wildcards. We do not discuss aggregates [8, Chap. 2] as they are trivially compatible.

Preserving order of effectful predicates. Since this work is motivated by incorporating external predicates, effectful predicate evaluation is a natural extension. Reordering, in general, is unsound for effectful predicates, e.g., we may reorder so that the subgoal that reads a file precedes the one that opens it.

One solution is to use a coarse-grained effect system that marks predicates as either pure or side-effectful. In such a coarse-grained effect system, all effectful predicates are assumed to interfere with each other, therefore relative order of all such predicates needs preserving. For each clause, assume a list of its effectful subgoals, l ,

matching their relative syntactic order. We modify edge-generation rules for MOGs so that the order of edge labels conforms to l .

For non-natural edges, if the subgoal selected for labelling is in l , then it has to be the head of l otherwise we do not generate the edge. For natural edges, the set of subgoals in the label can contain at most one element from l , particularly the head of l .

As this modification forces the edges to respect the syntactic order of effectful subgoals, the reordering algorithm is sound again.

The coarse-grained effect system here is only an example. So long as there is a sound partial order of effectful subgoals for each clause, a variation on the approach described here will be sound.

Negated subgoals. Negation allows a subgoal to hold when it is not satisfied. There are various ways of including it with radical effects on semantics, but dataflow-wise they behave identically.

Example 6.1. Evaluating the following naïvely, `outOfStock` consists of all finite strings but “Milk”. This compromises termination.

```
1 inStock("Milk").
2 outOfStock(X) :- not(inStock(X)).
```

Requiring variables inside a negated subgoal to be bound solves this problem. This is easy to express within our analysis by changing the generation of alternatives at the MOG root. Say we have a negated subgoal n of predicate p . If p has the ill constraint (Def. 3.3), the alternative at the root is (n, \emptyset) as before. Otherwise, we require all the variables to be bound, *i.e.*, the obligation set for n is $\{vars(n)\}$.

Wildcards. Wildcards allow the value of a subgoal argument to be ignored. For example, $p(X, _)$ ignores its second argument. This is equivalent to using an existential variable, *i.e.*, one that appears once in the body.

If a wildcard is used as an argument to a subgoal with mode + at that argument position, the dataflow requirement for that predicate cannot be satisfied. If the wildcards are eliminated through introduction of a fresh existential variable for each wildcard, no modification to our analysis is needed. However, if the analysis is performed without assigning fresh variables, we need two adjustments to the intra-clausal analysis. First, *vars* should ignore wildcards. Second, *cToOs* should add a special wildcard variable each time an atomic constraint indexes a wildcard variable in the subgoal. Since *vars* ignores wildcards and clause heads cannot have wildcards, a wildcard alternative will not be scheduled.

7 RELATED WORK

Mellish [12] introduces mode inference through abstract interpretation for PROLOG programs. Debray and Warren [6] improve on this work by precise handling of aliasing. Both perform inference on the programs as they are written without reordering of subgoals.

MERCURY [13] and HAL [5] have mode systems that are closest to ours in spirit. They both reorder subgoals to satisfy mode restrictions and both use constraint-based analysis. Both of these are higher-order languages and allow function symbols. Hence, they provide more sophisticated modes that express partial instantiations of variables, *e.g.*, a list with unbound variables is more instantiated than just a variable and less instantiated than a list with ground elements. Much of the analysis is thus concerned with precise aliasing tracking. By contrast, lack of function symbols simplifies our mode analysis. In particular, Overton et al. [13] reports

their constraint-based analysis is 10 to 100 times slower compared to their previous brute force search based algorithm on benchmark programs. Additionally, HAL only reorders subgoals during mode checking with mode specifications whereas we also do reordering in the absence of specifications. Another difference is that both of these are typed languages and their analyses rely on types for mode analysis. This is not possible for untyped DATALOG.

More recently, YEDALOG [4] and DYNA [7] were developed, inspired by DATALOG. They both add function symbols and face the same aliasing problems described above. Both provide static mode systems and refer to MERCURY as inspiration but without an explicit account of the underlying algorithm.

In addition to the order preservation method for subgoals with side-effects in Section 6, there is an alternative involving modes. Henderson et al. [9, Chap. 5] reify the external world as a value to be passed around. Use of mode constraints on external world arguments establishes mode dependencies between effectful clauses which would allow our analysis to remain sound without modification. This is similar to use of phantom types [3] in typed languages. The downside of this approach is that the external world has to be shuffled manually or a variable inserting transformation is needed.

Overall, we differ from the literature by targeting DATALOG in its standard form without function symbols and types. This simplifies the analysis and allowed us to prove soundness and completeness. Additionally, unlike other approaches, our analysis is incremental allowing performant mode checking in interactive systems.

8 CONCLUSIONS

We presented a static mode analysis for DATALOG to allow programs to be well-moded through reordering whenever possible. The combinatorial explosion of global permutation search is tackled by exploiting dataflow restrictions within the clauses and supporting incremental analysis—particularly for interactive systems. We showed that the algorithm is terminating, sound, and complete with respect to exhaustive global order search.

ACKNOWLEDGMENTS

We thank Alan Mycroft for his suggestions on the terminology and Tim Griffin for spotting errors in earlier drafts. This work was supported by the EPSRC [grant number EP/M026124/1].

REFERENCES

- [1] Molham Aref, Balder ten Cate, Todd J Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L Veldhuizen, and Geoffrey Washburn. Design and implementation of the LogicBlox system. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1371–1382. ACM, 2015.
- [2] Catriel Beeri and Raghu Ramakrishnan. On the power of magic. In *Proceedings of the sixth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 269–284. ACM, 1987.
- [3] James Cheney and Ralf Hinze. First-class phantom types. Technical report, Cornell University, 2003.
- [4] Brian Chin, Daniel von Dincklage, Vuk Ercegovac, Peter Hawkins, Mark S Miller, Franz Och, Christopher Olston, and Fernando Pereira. Yedalog: Exploring knowledge at scale. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 32. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- [5] M Garcia de la Banda, Peter J Stuckey, Warwick Harvey, Kim Marriott, et al. Mode checking in HAL. *Lecture notes in computer science*, pages 1270–1284, 2000.
- [6] Saumya K Debray and David S Warren. Automatic mode inference for logic programs. *The Journal of Logic Programming*, 5(3):207–229, 1988.
- [7] Jason Eisner and Nathaniel W Filardo. Dyna: Extending datalog for modern AI. In *Datalog Reloaded*, pages 181–220. Springer, 2011.

- [8] Todd J Green, Shan Shan Huang, Boon Thau Loo, Wenchao Zhou, et al. Datalog and Recursive Query Processing. *Foundations and Trends® in Databases*, 5(2): 105–195, 2013.
- [9] Fergus J Henderson et al. Strong modes can change the world. In *Technical Report 93/25*. Citeseer, 1993.
- [10] Herbert Jordan, Bernhard Scholz, and Pavle Subotić. Soufflé: on synthesis of program analyzers. In *International Conference on Computer Aided Verification*, pages 422–430. Springer, 2016.
- [11] Alberto Martelli. A Gaussian elimination algorithm for the enumeration of cut sets in a graph. *Journal of the ACM (JACM)*, 23(1):58–73, 1976.
- [12] Christopher S Mellish. Abstract interpretation of Prolog programs. In *International Conference on Logic Programming*, pages 463–474. Springer, 1986.
- [13] David Overton, Zoltan Somogyi, and Peter J Stuckey. Constraint-based mode analysis of Mercury. In *Proceedings of the 4th ACM SIGPLAN international conference on Principles and Practice of Declarative Programming*, pages 109–120. ACM, 2002.
- [14] Yannis Smaragdakis and Martin Bravenboer. Using Datalog for fast and easy program analysis. In *Datalog Reloaded*, pages 245–251. Springer, 2011.
- [15] Zoltan Somogyi. A System of Precise Models for Logic Programs. In *International Conference on Logic Programming*, pages 769–787. Citeseer, 1987.
- [16] Zoltan Somogyi, Fergus J Henderson, and Thomas Charles Conway. Mercury, an efficient purely declarative logic programming language. *Australian Computer Science Communications*, 17:499–512, 1995.
- [17] Jeffrey D. Ullman. Implementation of Logical Query Languages for Databases. *ACM Trans. Database Syst.*, 10(3):289–321, September 1985. ISSN 0362-5915. doi: 10.1145/3979.3980. URL <http://doi.acm.org/10.1145/3979.3980>.

A OMITTED PROOFS

LEMMA 3.15 (ILL AND TRIVIAL CONSTRAINTS). *The trivial constraint, $\{\emptyset\}$, is consistent with all binding patterns whilst the ill constraint, \emptyset , is consistent with none.*

PROOF. The \blacktriangleleft relation is universally quantified so it holds trivially for the atomic constraint \emptyset , hence $findAC$ is never empty for the trivial constraint. On the other hand, for the ill constraint, $findAC$ is always the empty set as there are no atomic constraints. \square

LEMMA 3.17. *For a fixed predicate, \leq is a bounded partial order with \emptyset as the top element (most restrictive) and $\{\emptyset\}$ as the bottom element (least restrictive). For a fixed domain, \leq is also a bounded partial order with constant functions returning \emptyset and $\{\emptyset\}$ as top and bottom elements respectively.*

PROOF. The fact that the relation \leq is a partial order follows from properties of implication. Nothing is consistent with \emptyset which confirms the top element and any adornment is consistent with \emptyset which confirms the bottom element (Lemma 3.15). Pointwise extension preserves all of these properties. \square

LEMMA 3.18. *If a constraint C_2 is more restrictive than C_1 , this means C_1 has a less restrictive atomic constraint for each atomic constraint in C_2 . That is:*

$$\forall C_1, C_2. C_1 \leq C_2 \implies \forall AC \in C_2 \exists AC' \in C_1. AC' \subseteq AC$$

PROOF. Unfolding definitions of \triangleleft and \blacktriangleleft , it is readily seen that the more restrictive constraint requires more arguments at particular locations to be *bound* than the less restrictive one, which establishes the subset relation. \square

LEMMA 4.7. *MOG construction leads to a scheduling graph.*

PROOF. The first two properties of the valid scheduling property is satisfied trivially as new vertices use $nextAcc$ and $nextAlt$ as defined in property statements. The edge labels are selected using $pickLabel$ which satisfies the third condition as it selects subgoals from the source vertex alternatives as well as the obligation it is coupled with.

Labels must not have obligation variables that do not appear in the head. When scheduling natural subgoals as they have empty obligations, that is the case. When scheduling non-natural ones, the condition is embedded inside $nonNats$ selector used by $pickLabel$, hence that too is satisfied.

The root vertex is also chosen as it is defined in the scheduling graph.

All conditions of a scheduling graph are satisfied. \square

LEMMA 4.8 (MOG TERMINATION). *For any clause and constraint function, the MOG construction terminates.*

PROOF. Apart from the root, the vertices are generated through edges. So it is sufficient to show that we can only generate finitely many edges. As a finite union of finite sets is finite, it suffices to show there is an n such that for all m bigger than to n , $E(\text{mog}_{cl,f}^m)$ is empty.

This is indeed the case since the edges can only be added using $pickLabel$ and $nextAlt$ from a given vertex. The $pickLabel$ function can only select labels from the alternatives of the preceding vertex

and $nextAlt$ removes the subgoals involved in edge labels from the succeeding alternative set. Hence, together they produce alternative sets strictly smaller than the edge's source vertex. Since we start with finite number of alternatives, we cannot generate infinite number of vertices. \square

PROPOSITION 4.11. *All paths in a MOG are greedy.*

PROOF. Natural edge generation by $pickLabel$ takes precedence during edge generation, therefore all paths satisfy greediness. \square

LEMMA 4.16. *Composition of the extract function with the MOG construction is monotonically increasing.*

$$\forall cl, f, g. f \leq g \implies \text{extract}_{cl}(\text{mog}_{cl,f}) \leq \text{extract}_{cl}(\text{mog}_{cl,g})$$

PROOF. Fix a clause cl and assume $f \leq g$ for some f and g . This implies at each point p in these functions we have $f(p) \leq g(p)$. It is sufficient to show that every terminal path that can be reached via g is also reachable with f (in the sense that it has the same number of edges sharing the same subgoals on the edges but paired possibly with different obligations) and the accumulator at the end of each these paths is a subset of that of g .

By Lemma 3.18, we know for each predicate p , for all atomic constraints of $g(p)$, there is an atomic constraint in $f(p)$ that is smaller. This property trivially transfers to obligations in clause context. This means each time we can extend a vertex using g we can extend it with f and the obligation is at most as big as that of come from g . Hence, the accumulated obligation yields the desired subset property for each path.

There may be other terminal paths generated by f . If any of those leads to smaller accumulators, then the property still holds. If they do not, it does not matter because atomic constraints are minimised, so the redundant atomic constraints are removed. \square

LEMMA 4.17. *If any predicates in the body of a clause cl has the ill constraint according to the constraint function f , the constraint for the clause as extracted from its MOG is also ill.*

$$\forall sub \in \text{body}(cl). f(\text{pred}(sub)) = \emptyset \implies \text{extract}_{cl}(\text{mog}_{cl,f}) = \emptyset$$

PROOF. The $extract$ function derives constraints using terminal vertices. However, $\text{mog}_{cl,f}$ cannot have a terminal vertex because the alternatives contain a subgoal sub with empty set of obligations. This obligation is not $\{\emptyset\}$ so we cannot schedule sub using the first case of $pickLabel$ for natural subgoals. We also cannot use the other case of $pickLabel$ which uses $nonNats$, as the side condition for the new edge includes $o \in os$ and os in this case is empty. Since there is no way of scheduling sub , there is no way of reaching a terminal vertex. \square

LEMMA 4.20 (WELL-MODED ORDERING TO TERMINAL PATH). *For a constraint function, f , clause, cl , binding pattern, \mathbf{a} , and ordering, σ , where $\text{wellModed}(cl, \mathbf{a}, f, \sigma)$, a scheduling graph $\text{schedule}(cl, \mathbf{a}, f, \sigma)$ has a terminal path p .*

PROOF. For at least one path to be terminating, it follows that E_i must be non-empty for every $0 \leq i < n$. By the above construction and the *well scheduling* properties, this would then imply that $\{(\emptyset, Acc)\} \in E_n$ for some Acc , i.e. that $\text{terminal}(p)$ for some path p .

Since the root node fills all the obligations for each subgoal from f , i.e., $V_0 = \{(s, \min cToOs_s(f(pred(s)))) \mid s \in body(cl)\}$ then each singleton obligation $\{o\}$ for a subgoal s is derived from $f(pred(s))$.

From well-modedness, we have that $adornment(s) \triangleleft f(pred(s))$ for each subgoal $s \in body(adorn(cl, a, \sigma))$, therefore, where we let $\mathbf{b} = adornment(s)$:

$$\begin{aligned} & \mathbf{b} \triangleleft f(pred(s)) \\ \iff & \mathbf{b} \triangleleft osToC_s(cToOs_s(f(pred(s)))) \\ \implies & \mathbf{b} \triangleleft osToC_s(\{o\}) \text{ where } \exists o \in cToOs_s(f(pred(s))) \\ \iff & \mathbf{b} \triangleleft osToC_s(\{o\}) \text{ where } \exists o \in \min cToOs_s(f(pred(s))) \end{aligned}$$

Therefore, from well-modedness, we can always satisfy the conditions of the E_i set comprehension for at least one pair (s, o) . Thus, $E_i \neq \emptyset$ for all $0 \leq i < n$. Therefore, there is at least one terminal path $p \in schedule(cl, a, f, \sigma)$. \square

LEMMA 4.23 (TRIVIAL OBLIGATION CONSISTENCY). *If a scheduling graph vertex v has a subgoal s with obligation $\{\emptyset\}$ in its alternative set, any ordering with a partial ordering derived from the root to v makes the predicate s consistent with respect to any binding pattern.*

PROOF. Due to the invariants of the scheduling graph, the only way v has s with the trivial obligation in its alternative set is if it had it that way at its root vertex or subgoals scheduled on p before v released variables in its obligation.

Since the variables in the obligation correspond to variables that needs to be bound to be consistent with the predicate constraint, any ordering that follows any one of the partial orderings up to v has s consistent with its constraint regardless the binding pattern of the clause head. \square

LEMMA 4.24 (SWAP PRESERVATION). *If a scheduling graph path is well-moded, performing a swap on this path preserves well-modedness and the path still belongs to some scheduling graph.*

PROOF. We first show the path is still in some scheduling graph and then show it retains well-modedness.

- (1) Structurally, swap does not change the vertices or the edges. The vertices conform with the properties of a scheduling path as they use $mkEdge$ to construct the vertices. It is used in Def. 4.5 and Lemma 4.7 establishes that the vertices produced by it are those expected by a scheduling graph. The edges also conform with scheduling graph requirements. The obligations within the labels being subset of the head variables is satisfied trivially as changing the position of the label has no effect on this. The final requirement is that the elements of the labels (subgoal obligation pairs) have to be chosen from the alternatives. By assumption the label moved to the left, has the trivial obligations in the alternatives of the preceding vertex, so the requirement is satisfied. In the new intermediate vertex the alternatives may include shrank obligations due the variables bound the subgoals moved to the left. But the label is modified to exclude these variables from the obligations, so the property holds for this new label as well.
- (2) Let the sets of subgoals in the edges p_i and p_{i+1} be P and Q respectively. Consistency depends on the adornment of the subgoal which depends on the variables bound before

adornment. We already fixed the head adornment \mathbf{a} and that does not change with subgoal swapping.

We proceed by considering different portions of the path:

$q_{j < i}$ The set of bound variables are same as before, so the subgoal prior to this point remain consistent.

$q_{j > i+1}$ When regarded atomically, P and Q together bind the same set of variables regardless the order they are scheduled in. Hence, the bound variables after the vertex q_{i+1} remain the same.

q_{i+1} The set of variables there were sufficient to make subgoals in P remain bound when the subgoals in P are moved to the right. They might be augmented by the addition of variables in Q but additional bound variables do not compromise consistency.

q_i Because we have the side condition on swap that all subgoals of Q must have the trivial obligation at p_i , we already know all that needs to be bound to achieve consistency is bound at p_i (Lemma 4.23).

Hence, both properties are preserved. \square

LEMMA 4.26 (MERGE PRESERVATION). *If a scheduling graph path is well-moded, performing a merge on this path preserves well-modedness and the path still belongs to some scheduling graph.*

PROOF. We first show the path is still in some scheduling graph and then show it retains well-modedness.

- (1) We start with a path in the scheduling graph and make no modification to the vertices before and after the edges being merged. Hence, it is enough to show that the invariants are satisfied for the edges being merged. Accumulator related invariants are trivially satisfied as the obligations on the labels of both of these edges are \emptyset (due to third invariant of the valid scheduling property), hence they do not change the accumulator (as they did not before). They subset restriction on the label obligations are also satisfied, as \emptyset is trivially a subset of the head variables. The alternative set of the destination of the merged edge is also the same as before since exactly the same subgoals are removed from the alternative set and hence the same variables (of these subgoals) are released.
- (2) Let p be the path in question and p_i & p_{i+1} be the edges being merged. Fix the binding pattern \mathbf{a} that the path is consistent with.

A merge only enables new orderings. We need to show each of these new orderings are still consistent with all of the predicate constraints when their binding patterns are derived from \mathbf{a} .

In the new orderings, the subgoals before p_i appear in the positions they did in the old orderings, in which they were already consistent. The subgoals after p_{i+1} also appear in the same positions as before and changing the positions of the subgoals in p_i and p_{i+1} do not affect the variables bound during the adornment of these subgoals, hence they too remain consistent.

The new locations subgoals of p_i can appear in the new orderings are still after the points we established their consistency, hence their consistencies are preserved.

Those in p_{i+1} can appear at locations before the points we established their consistencies, but by assumption we only merge if those subgoals appear with trivial obligations inside the source of p_i . By Lemma 4.23, we know that these predicates also retain their consistencies.

Hence, both properties are preserved. \square

LEMMA 4.27 (CONVERSION PRESERVATION). *If a scheduling graph path is well-moded, its conversion is also a well-moded path of some scheduling graph.*

PROOF. Swap and merge preserves scheduling graph structure and well-modedness by Lemma 4.24 and Lemma 4.26. \square

LEMMA 4.28 (CONVERSION GREED). *If a path is in a scheduling graph, then its conversion produces a greedy path.*

PROOF. By construction subgoals in edges are positioned such that they follow the trivial obligation, since a subgoal cannot appear in more than one edge, greediness requirement is satisfied. \square

LEMMA 4.29 (GREEDY PATH COMPLETENESS). *For a clause cl , every greedy scheduling path for cl ending in a terminal vertex and conforming to a constraint function f is present in the MOG determined by cl and f . That is:*

$$\forall p, a. \text{greedy}(p) \wedge \text{wellModedPath}(p, a, cl, f) \implies p \in \text{paths}(\text{mog}_{cl, f})$$

PROOF. We consider a more general property: that for a greedy, well-moded path p , a prefix of p of length n is a prefix path of $\text{mog}_{cl, f}$, where an empty path comprises just the root vertex. We assume a path p satisfying the antecedent of the lemma, and proceed by induction on the length n of the prefix path:

- $n = 0$. By the definition of scheduling graphs, the root node V_0 is fixed, therefore $\{V_0\} = \mathbf{V}(\text{mog}_{cl, f}^0)$ trivially; a zero-length path comprises just the start vertex.
- $n = k + 1$. Let $(s, o) = \text{src}(p_{k+1})$ and assume the inductive hypothesis: the path $p_0 \dots p_k$ is a prefix path of $\text{mog}_{cl, f}$. We consider then two cases:
 - $(\text{nats}(\text{src}(p_{k+1}))) \neq \emptyset$ therefore by greediness and the well-scheduling property, $\text{trg}(p_{k+1}) = \text{mkVertex}(\{(\text{nats}(\text{Alt}), \emptyset)\})$ which is equal to the edge constructed by $\text{mog}_{cl, f}$ give the vertex $\text{src}(p_{k+1})$;
 - $(\text{nats}(\text{src}(p_{k+1}))) = \emptyset$ therefore by well-modedness on the subgoal s we have that $\text{adornment}(s) \triangleleft f(\text{pred}(s))$. For the computed adornment to be consistent with the constraints of f , it follows that for every variable X in this clause which is bound in subgoals adornment, its corresponding index i in the constraint due to f . For X to be bound, it follows that it was bound earlier on the path, or is bound in the clause head. The former cannot be true, as if it was bound earlier on the path it would have been released from the alternative set via \ominus . Subsequently, it must be bound in the head and therefore its obligation $o \subseteq \text{vars}(\text{head}(cl))$. Therefore, by well-scheduling, we have a vertex which satisfies the requirements of nonNats in pickLabel , thus the edge p_{k+1} is equal to that constructed by $\mathbf{E}(\text{mog}_{cl, f}^k)$ at this point given vertex $\text{src}(p_{k+1})$.

Therefore, $p_0 \dots p_k p_{k+1}$ is a prefix path of $\text{mog}_{cl, f}$.

Therefore $p \in \text{paths}(\text{mog}_{cl, f})$. \square

LEMMA 4.30 (PATH EXTRACT CONNECTION). *For a fixed binding pattern a , existence of a well-moded path in a MOG implies consistency of a with the constraint extracted from the MOG.*

$$\forall a p. \text{wellModedPath}(p, a, cl, f) \wedge p \in \text{paths}(\text{mog}_{cl, f}) \\ \implies a \triangleleft \text{extract}_{cl, f}(\text{mog}_{cl, f})$$

PROOF. Fix a binding pattern a and a path p and assume the antecedent. To show the consequent, it is enough to show a stronger statement: the atomic constraint AC extracted from any well-moded MOG path p is consistent with the binding pattern a . This generalisation is valid since atomic constraints are combined via \otimes in extract which is monotonically decreasing (getting less restrictive wrt. consistency).

Assuming an arbitrary index $i \in AC$ constraining variable X in the head, it follows that X is in the terminal accumulator and was scheduling at some point in the path p by extending the accumulator with this variable. Let s be the subgoal that causes this augmentation. Since we know by the premise that p is well-moded with respect to a and f , we also know the constraint of s in f is satisfied by the binding pattern at s derived from head binding pattern a by adorn . Since scheduling s augmented the accumulator, X could not have appeared in the previous subgoals as the release operator \ominus would have eliminated X from the obligation of the alternative. Hence, X must be bound in the head. This is exactly what is required for the consistency with AC . As there is nothing particular about i and X , all indices in AC are similarly bound, hence the atomic constraint is consistent with a .

If there are no indices in the accumulator at the end of the path, the extracted constraint has to be trivial as \otimes takes the minimal elements. Every binding pattern is consistent with the trivial constraint (Lemma 3.15), so the lemma holds. \square

LEMMA 5.3 (\oplus CONSISTENCY HOMOMORPHISM). *A binding pattern is consistent with two constraints combined with \oplus iff that binding pattern is consistent with each constraint individually, i.e.:*

$$\forall a, C_1, C_2. a \triangleleft (C_1 \oplus C_2) \iff a \triangleleft C_1 \wedge a \triangleleft C_2$$

PROOF. (\implies) Assume $a \triangleleft (C_1 \oplus C_2)$. This means that for every element AC of $C_1 \oplus C_2$, we have $a \triangleleft AC$. Unfolding definition of \oplus , we have X and Y that are subsets of AC such that AC is $X \cup Y$. Unfolding definition of \triangleleft , we have $\forall i \in AC. a_i = b$. This certainly holds for all subsets of AC , so we have $\forall i \in X. a_i = b$ and similarly for Y . C_1 must be a set of such X by definition of \oplus , thus $a \triangleleft C_1$ holds. Same argument applies to C_2 as \oplus is commutative.

(\impliedby) Assume $a \triangleleft C_1$ and $a \triangleleft C_2$. $C_1 \oplus C_2$ is a subset of $C_1 \times C_2$, so it suffices to show $a \triangleleft C_1 \times C_2$. This requires showing a is consistent with every atomic constraint in this set. Since each of these atomic constraints can be represented as a union of an element from C_1 and another element from C_2 and that we know a is consistent with each of these elements, it also has to be consistent with the union. \square

LEMMA 5.5. *Every constraint function generated from a moding function gets more restrictive when step is applied to it.*

$$\forall Pr, mv. \llbracket mv \rrbracket_F \leq \text{step}_{Pr}(\llbracket mv \rrbracket_F)$$

PROOF. Unfolding definition of \leq , it is sufficient to show the equivalent pointwise property holds:

$$\forall mv, p. \llbracket mv \rrbracket_F(p) \leq \text{step}_{Pr}(\llbracket mv \rrbracket_F)(p)$$

By assumption mv only has mode requirements for external predicates. Let p be an arbitrary predicate within the domain of $\llbracket mv \rrbracket_F$. Now we consider effect of step depending on whether p is an internal or an external predicate.

If p is an external predicate, step does nothing, so the lemma holds by reflexivity of \leq (Lemma 3.17).

If p is an internal predicate, by assumption $\llbracket mv \rrbracket_F(p)$ is $\{\emptyset\}$, which is the bottom for \leq (Lemma 3.17). \square

LEMMA 5.6 (STEP MONOTONICITY). *The step function is monotonically increasing (making constraints more restrictive):*

$$\forall Pr, f, g. f \leq g \implies \text{step}_{Pr}(f) \leq \text{step}_{Pr}(g)$$

PROOF. Follows from monotonicity of extract (Lemma 4.16) and \oplus consistency homomorphism (Lemma 5.3). \square

LEMMA 5.7. *Every constraint function generated from a moding function gets more restrictive by application of analyse .*

$$\forall Pr, mv. \llbracket mv \rrbracket_F \leq \text{analyse}_{Pr}(\llbracket mv \rrbracket_F)$$

PROOF. Observe that analyse is simply repeated application of step . The lemma follows from Lemma 5.5 and Lemma 5.6. \square

PROPOSITION 5.8. *For each predicate p , $(D_p, \oplus, \otimes, \{\emptyset\}, \emptyset)$ is an idempotent commutative semiring³ where $\{\emptyset\}$ and \emptyset are the additive and multiplicative identities respectively.*

PROOF. Previously given by Martelli [11]. \square

PROPOSITION 5.9 (FAST FAILURE). *After a single application of step , the ill constraint propagates from the body of a clause to the entire head predicate constraint.*

$$\begin{aligned} \forall f, cl, s \in \text{body}(cl). \\ f(\text{pred}(s)) = \emptyset \implies \text{step}_{Pr}(f)(\text{pred}(\text{head}(cl))) = \emptyset \end{aligned}$$

PROOF. Fix f , cl , and s assume the antecedent. By Lemma 4.17, we know the clause constraint has the ill constraint. By Lemma 5.3, we know any a that is consistent with a constraint combined using \oplus with \emptyset must have $a < \emptyset$. There is no such a , thus the overall constraint for the predicate at the head of the clause is \emptyset as required. \square

COROLLARY 5.10. *The number of step applications it takes to converge to the ill constraint is bounded by the static call distance between two predicates.*

PROOF. Follows immediately from Proposition 5.9. \square

COROLLARY 5.14 (GLOBAL REORDERING EXISTENCE). *For every program that can be well-moded with reordering, we can construct a global reordering function.*

PROOF. By Theorem 5.13, we know that if a global reordering function exists, analysis will find the trivial constraint for the head. This means for each clause as a part of the analysis we construct local reordering functions that are defined for all relevant binding patterns. By combining these local reordering functions, we can construct the desired global reordering function. \square

COROLLARY 5.16. *When the head predicate of cl does not feature in Pr , the clauses of Pr can be ignored during analysis. That is:*

$$\text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) = \text{analyse}_{\{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F))$$

PROOF. Constraints of clauses are obtained by \oplus operator and intra-clausal analysis which is a function of the clause body and the constraints for those subgoals. Clauses apart from cl are unaffected by the constraint of cl as the head predicate for this clause by assumption does not feature in other clauses, hence cannot affect the overall predicate constraints. Since the step function preserves constraints that are not mentioned in the set of clauses it is parameterised over, the lemma holds. \square

COROLLARY 5.17 (FAST CONVERGENCE). *If cl is also non-recursive, then analyse converges to a fixpoint in a single step:*

$$\text{analyse}_{Pr \cup \{cl\}}(\llbracket mv \rrbracket_F) = \text{step}_{\{cl\}}(\text{analyse}_{Pr}(\llbracket mv \rrbracket_F))$$

PROOF. We know the constraints of predicates appearing in Pr are all stable and body of cl can only feature them and external predicates for which the constraints do not change. Since the clause constraint is a function of its body and the constraints at this point, single iteration is sufficient. \square

³It is known as Martelli's semiring, originally used to compute cut-sets of a graph [11].