

**A Minimal Function Graph Semantics
for Logic Programs**

by

Will Winsborough

**Computer Sciences Technical Report #711
August 1987**

A Minimal Function Graph Semantics for Logic Programs

Will Winsborough

August 14, 1987

Abstract

The problem is to devise an abstract interpretation framework for static analysis of logic programs. We seek a framework in which to couch analyses that infer stronger assertions than do previous efforts, but that remain plausibly efficient. The framework offered here is designed to efficiently accommodate large abstract domains that introduce imprecision only as diversity naturally arises in the inferences.

Our semantics is based on the view of predicates as partial functions from sets of activation instances to sets of result instances. It is essentially a *minimum function graph* (*mfg*) semantics, following [8], though it treats sets of activations rather than individuals because doing so allows us to use large abstract domains, and hence infer stronger properties, without sacrificing efficiency.

First we construct a function whose fixpoint defines a *total function graph* (*tfg*) semantics. We then define the *mfg* semantics as the *tfg* restricted to the call instances reachable from some specified query set.

Finally, we consider how to compute the *mfg* directly, without having to compute the (large) *tfg*. If reachability is added to the iteration function whose fixpoint defines the *tfg*, problems arise because we treat sets of activations. The function is no longer monotonic, and its (transfinite) Kleene's sequence doesn't converge to the *mfg*. The Kleene's sequence does converge, however, and we define our *partial function graph* (*pfg*) semantics as its limit.

For finite abstract domains the *pfg* semantics is computable and is generally much smaller and much less expensive to compute than the *tfg*. In the main technical result of the paper, we prove the *pfg* agrees exactly with the *mfg*, except that a few activation sets (though, *not* individuals) erroneously appear reachable. The difference has no meaning for most execution models. So, despite its theoretical unpleasantness, the *pfg* is very practical.

1 Introduction

A data flow analysis is an attempt to automatically infer properties of program executions with a terminating procedure. Termination of the procedure can be guaranteed for some properties that are weaker than the exact details of the program execution. The technique of abstract interpretation, based on fixpoint semantics constructions, provides a framework for constructing data flow analyses so that their correctness can be easily verified. Two essentially homomorphic interpretations of a programming language are formally related, one precise (called the concrete interpretation), the other approximate (called an abstract interpretation). The concrete interpretation induces a collecting semantics that agrees precisely with the standard semantics. An abstract interpretation induces an approximate semantics that is a data flow analysis. Given that the interpretations satisfy some relatively simple properties, powerful fixpoint theorems can be invoked to establish that the abstract semantics is a (weaker) implication of the concrete semantics.

We follow [7,8] and others by constructing a *core semantics* using a collection of domains and auxiliary functions, called an interpretation, that are assumed to have certain properties, and that act as formal parameters to the construction. Different actual interpretations induce different semantics. In this paper we construct a core semantics for pure Horn Clause logic programs and give an interpretation that induces a collecting semantics that agrees with SLD-resolution. Abstract interpretations can be provided to generate various data flow analyses.

To demonstrate that an abstract interpretation generates a safe approximate semantics, there must be an abstraction function that maps concrete domain elements to abstract domain elements and that is (almost) a homomorphism relative to the iteration functions whose least fixpoints define the respective semantics. When a core semantics is given, the homomorphism requirement is factored into requirements of the interpretations' auxiliary functions. The result of the (almost) homomorphism is that the abstract semantics is implied by the concrete semantics. This fact comes from well known theorems about fixpoints[3]. Because this concept is well understood and because of space considerations, we do not include a formal treatment of safeness of abstract interpretations for our framework.

1.1 Our Motive

This work is of general interest because it gives an abstract interpretation framework for a declarative language based on a function domain that takes sets of activation instances to sets of result instances, but that does not use power domains. Further, it defines a *minimal function graph (mfg)* semantics in such a domain and develops a practical way to compute it for finite, though possibly large, domains.

Our particular interest is to design a framework in which to couch analyses that improve the strength of the assertions inferred over the analyses that have been proposed thus far in the literature. We want a framework that will allow such an improvement while keeping overhead at a plausible level. Thus, the framework presented here is designed to accommodate large abstract domains that can represent properties with great precision, but that introduce imprecision as diversity arises in inferences. We discuss the issues involved in that accommodation in Section 1.4.

Data flow analyses of logic programs include mode inference, which infers variable instantiation patterns at call activation and result[11,6,4,5], and occur check avoidance, which finds calls in which there is no threat that unification will introduce cycles into structures[13,14].

We are primarily interested in mode inference, which is useful in sequential compilers for optimizing unification and storage management[15], and to parallel execution models, where we may execute concurrently two literals that are known to share no free variables[2,1].

The best published work on mode inference has been done by Debray. Mellish has also proposed an analysis for mode inference [11], but unfortunately that work fails to treat correctly the problem of variable aliases. The method of Debray, developed in [6,4,5], treats predicates as instantiation state transforms. A result instantiation description is associated with each activation instantiation description that could arise while executing the program. Debray's analysis is an example of the kind of analysis our framework can be used to verify. However, we have constructed the framework with larger abstract domains in mind.

Using our framework, we have developed and implemented an abstract interpretation that handles mode inference and that infers stronger properties than did previous analyses for mode inference. That will be the subject of a forthcoming paper.

1.2 Possible Frameworks

The original exposition of abstract interpretation by Cousot and Cousot gave a collecting semantics for a simple imperative language without procedure calls[3]. That semantics constructs simultaneously a description of the execution states possible at each point in the program. In their language, a program is represented by a directed graph with a node for each instruction and a directed arc from each node to the node for each possible next instruction. The arcs are the "program points" where execution state descriptions are collected. Cousot and Cousot call an execution state description a context, and a vector of these descriptions, one for each program point, a context vector.

In the Cousot framework, an iteration function (called an interpretation function in [3]) simulates one step of program execution simultaneously at all nodes in the graph. This iteration function takes one context vector and produces another. It simulates

the operation at each node by looking at the possible incoming states according to the information in the input context vector and generating a description of the possible outgoing states on that basis. The program’s collecting semantics is the least fixpoint of this iteration function.

The iteration function for a logic program will also take an input context vector. To produce an output context vector it will simulate the execution of each clause by simulating its calls according to the input context vector. The unification of calls with call results and the composition of call solutions according to the clause are the two basic operations of logic programs. Both of these are performed to complete one application of the iteration function. The context vector should represent the results of various calls. Additionally, we would like to determine which call activations are reachable from some specified query set. But what form should the context vector take?

C. S. Mellish has constructed an abstract interpretation framework for logic programs [12]. In it, procedure entry and exit are taken as the program points. Canonical literals, normalized representatives of their renaming equivalence class, are collected into two sets, called Input and Output. Input (resp. Output) contains call activation (resp. result) instances that could possibly occur during the execution of one of the programmer specified queries. In that framework, the two sets constitute the context vector. This approach does not preserve information about which result instances correspond to which activation instances. Because the semantic construction must work for either the concrete or the abstract interpretation, one cannot rely on unification to screen out inappropriate results. All result instances that have been found to occur for a given predicate are treated as possible results for every call to that predicate. Irrelevant results propagate and accumulate.

In search of a more powerful framework, we observe that (pure) logic programs are referentially transparent; the behavior of a call is entirely determined by the activation instance. If we consider *sets* of activation instances as the program points, it is possible to get a finite approximation domain, and hence termination, by selecting representative sets to constitute the approximation domain.

Thus we conclude that the context vector shall be a function from activation sets to result sets. Frequently we shall say “description” rather than “set” because an abstract domain may be any homomorphic structure. We understand the term “minimal function graph”, borrowed from [8], to suggest a convenient representation for the function.

1.3 A Minimal Function Graph Framework

As we shall see, it is a relatively straightforward matter to define $\mathbf{TFG}[[\mathbf{Pr}]]$, a *total function graph* (*tfg*) semantics for pure Horn Clause programs. It is a function that takes any set of calls to some predicate and returns the set of result instances and is defined in Section 2.4. Under the concrete interpretation, the *tfg* semantics agrees

exactly with a standard operational semantics that evaluates a query by computing *all* solutions to the first call, then, for each of those solutions, computing all solutions to the corresponding substitution instance of the second call, and so on. Such a model could be made realistic by using eager consumers to avoid problems with infinite solution sets and other non-terminating calls. That execution model is complete and is the basis for the parallel execution model of Kalé[9]. However, for our purposes it is clearer to ignore termination problems and adopt the more declarative view of the model as a composition of operations on sets. We will refer to this loosely as the *all-solutions* model.

The *tfg* semantics does not include any information about which activation sets are reachable. A *minimal function graph (mfg)* semantics is a function that is only defined on those activation sets that are reachable from a given query set. Obviously, this *mfg* semantics can be constructed from the *tfg* by constructing the set of reachable call activation sets and restricting the *tfg* to that. Thus, the *mfg* contains not only information about call results, but also about which calls are reachable. The latter information is represented by the domain over which the function is defined.

Unfortunately, we do not know how to construct the *mfg* with a single least fixpoint operation. (Reachability requires a fixpoint operation.) If one incorporates reachability of activations into the *tfg* iteration function, the function becomes non-monotonic. Which activation instance sets appear reachable at some clause body call depends on the partial results currently reported for other calls to the left in the clause. For example, suppose we are trying to analyze the following program in the concrete interpretation.

```
p(X,Y) :- s(X), r(X,Y).
s(a).
s(X) :- t(X).
t(b).
r(a,a).
r(b,b).
?- p(X,Y).
```

Further suppose we have inferred that $\{s(X)\} \mapsto \{s(a)\}$, but have not yet inferred that $s(b)$ is also a solution to $s(X)$. At this stage it appears that $\{r(a,Y)\}$ is a reachable activation set. If we repeatedly apply the iteration function, $\{r(a,Y), r(b,Y)\}$ will eventually be inferred to be reachable. But the *mfg* only includes the larger set, since the execution model collects *all* solutions to $s(X)$ before trying $r(X,Y)$.

As we will show, a limit to the Kleene's sequence¹ still exists, defining what we call here the *partial function graph (pfg)* semantics. But it is not the *mfg* because some subsets of genuinely reachable activation sets incorrectly appear reachable.

¹The Kleene's sequence of the function $f : A \rightarrow A$ is defined to be the transfinite sequence $\{f^\gamma(\perp)\}_\gamma$, where $f^0(x) =_{df} x$; $f^{\delta+1}(x) =_{df} f(f^\delta(x))$; and $f^\alpha(x) =_{df} \sqcup \{f^\delta(x) \mid \delta \sqsubset \alpha\}$, for α a limit ordinal.

Of course, most applications are for execution models in which only the reachable *individual* activations matter. On which individuals these are the *mfg* and the *pfg* agree. Furthermore, in Section 3.2 we prove that on the reachable activation descriptions the *mfg*, the *pfg*, and the *tfg* are all equal.

The other difference between *pfg* and *mfg* is that the iteration function for the *pfg* is not continuous, so the sequence whose limit defines the *pfg* will not in general converge by stage ω . However, termination is ensured when a finite approximation domain is used. And this restriction turns out to be essentially necessary if the *mfg* is to consist of finitely many pairs.

1.4 Recent Related Work

We have recently become aware of closely related work by N. Jones and H. S ndergaard, reported in [7]. Roughly speaking, that framework is based on functions from substitutions to sets of substitutions. We believe that the choice of substitutions versus canonical term-tuples is not significant.

The Jones-S ndergaard formulation is very elegant. The domains for activations and results are distinct in the core semantics and the former is a subset of the latter. In the concrete interpretation, the activations are singletons of substitutions and the results are arbitrary sets of substitutions. However, the two domains may be the same actual domain in an abstract interpretation.

Suppose B_1, \dots, B_m is a clause body. The Jones-S ndergaard core semantics specifies that, when simulating B_i , the call must be evaluated under each activation domain element less than the result domain element describing the set of substitutions resulting from B_1, \dots, B_{i-1} . So, in the concrete interpretation, each member of the set of substitutions resulting from B_1, \dots, B_{i-1} is applied to B_i in turn. If the abstract interpretation makes the activation domain and the result domain the same, then there is only one instance of B_i to evaluate. But we are interested in accommodating large abstract domains where it is possible to preserve precision until diversity arises in the inference. Such domains are relatively tall. Using such a domain, it may be necessary in the Jones-S ndergaard semantics to cycle through a large number of call activations. Our framework allays this threat to efficiency by letting the granularity of activation description for which results are collected be determined by the amount of diversity in each activation set inferred to be reachable. Rather than require the abstract interpretation to fix the granularity of activation description for which results are collected, our method allows that collecting granularity to be determined by the granularity of the activation descriptions reachable from the query in our all-solutions model.

Another feature of the Jones-S ndergaard framework that differs significantly from ours is how reachability inferences are collected. We require that a query set be provided. It provides the basis for reachability, and, for the most part, results are only collected for

reachable activation descriptions. The Jones-Søndergaard method collects a result for *every* activation, and in the process, collects a log of activations that may be reachable from each of those activations.

From a theoretical standpoint, we admire the Jones-Søndergaard construction very much. From a practical standpoint, our method appears preferable, particularly for use with large abstract domains.

1.5 Organization

The rest of this paper is organized as follows. In Section 2, we construct the *mfg* semantics and give the concrete interpretation for the formal parameters to the construction. Then, in Section 3, we present the *pfg* construction, show that it is well defined and prove that it agrees with the *mfg* wherever the latter is defined.

Because of space considerations we do not present a proof that the concrete semantics agrees with the SLD-resolution semantics.

We assume the reader is familiar with the basic ideas of abstract interpretation[3] and of logic programming[10]. Familiarity with [8] would help.

2 The MFG Construction

This section gives the core semantics construction and presents the concrete interpretation. The core semantics construction defines the *tfg*, the set of reachable activation descriptions, and the *mfg*.

2.1 Syntactic Domains

Let **Func** denote the set of ℓ -placed function symbols, $\ell \geq 0$. Let **Var** denote a countably infinite set of variables. Let **Term** denote the set of logical terms over **Func** and **Var**. Let **Pred** denote the set of k -placed predicate symbols. We assume for simplicity of presentation that one fixed k holds for all predicates in any given program. Let **Literal** denote the set of logical literals over **Pred** and **Term**.

Let **Clause** be the set of Horn clauses. Henceforth, all clauses mentioned are implicitly Horn. We write clauses $H :- B_1, \dots, B_m$, $m \geq 0$. As usual, the ‘:-’ is omitted when $m = 0$. H is called the clause head; together, B_1, \dots, B_m make up the clause body. Let $\mathbf{Prog} =_{df} \{\mathbf{Pr} | \mathbf{Pr} \subset \mathbf{Clause} \text{ and } \mathbf{Pr} \text{ is finite}\}$ be the program domain. $\mathbf{Pr} \in \mathbf{Prog}$ is a collection of Horn clauses constituting a given program. We assume that the number of (non-primitive) predicates occurring in \mathbf{Pr} is n and we refer to them as p_1, p_2, \dots, p_n . A clause for predicate i is one that has p_i in its head. $\mathbf{Queries} \in \wp(\mathbf{Literal})$ will denote the set of possible queries associated with \mathbf{Pr} . Without loss of generality, we make the simplifying assumption that elements of $\mathbf{Queries}$ are single negative literals.

2.2 Interpretations

The construction of our semantics is parameterized by two domains and several auxiliary functions. These parameters, taken together, constitute an interpretation[8]. Basic assumptions are required of an interpretation for the construction to be well defined.

Definition. An interpretation consists of the following:

Two Lattices:

$$\begin{aligned} & \text{Arg}, \top_{\text{Arg}}, \perp_{\text{Arg}}, \sqcup_{\text{Arg}}, \sqsubseteq_{\text{Arg}} \\ & \text{CArg}, \top_{\text{CArg}}, \perp_{\text{CArg}}, \sqcup_{\text{CArg}}, \sqsubseteq_{\text{CArg}} \\ & \text{Env}, \top_{\text{Env}}, \perp_{\text{Env}}, \sqcup_{\text{Env}}, \sqsubseteq_{\text{Env}} \end{aligned}$$

Monotone Functions:

$$\begin{aligned} & \text{Apply} : [0..m] \times \text{Env} \times \text{CArg} \times \text{Clause} \rightarrow \text{Env} \\ & \quad \text{Apply must be bottom-preserving and continuous} \\ & \quad \text{in its second and third arguments.} \\ & \text{Project} : [0..m] \times \text{Env} \times \text{Clause} \rightarrow \text{Arg} \\ & \quad \text{Project must be bottom-preserving and continuous} \\ & \quad \text{in its second argument.} \\ & \text{GetArg} : \wp(\text{Literal}) \rightarrow \text{Arg}^n \\ & \text{GetEnv} : \text{Clause} \rightarrow \text{Env} \\ & \text{Canon} : \text{Arg} \rightarrow \text{CArg} \\ & \text{Freelnst} : \text{CArg} \rightarrow \text{Arg} \end{aligned}$$

$x \in \text{Arg}$ represents a set of literal instances. x does not indicate which predicate symbol occurs in the literals, just which k -tuple of terms appear as arguments. In the sequel, elements of Arg will simply be said to represent sets of instances.

The elements of CArg represent sets of instances that are canonical representatives of their variable renaming equivalence class. So, $x \in \text{Arg}$ can contain two instances that are the same up to variable renaming, but, $y \in \text{CArg}$ cannot.

The context vector, which our iteration function operates on, will be a vector of functions from CArg to CArg . Recall that the iteration function takes an *input context vector* and produces an *output context vector* by simulating each clause execution based on the call results given by the input context vector. Each component of the vector will correspond to a different predicate p_i , and in the fixpoint will represent the behavior of sets of calls to p_i .

An element of Env represents a *local clause context*. The local clause context is local in the sense that its lifetime is the duration of the simulation of a single clause execution. It is used by the iteration function for recording intermediate results. But the output context vector only records what the local clause context has to say about the state of the clause *head* after clause execution simulation.

Initially, a clause's local context represents only one instance of the clause: the one that appears in the program text. But when the clause is activated with a set of call instances for the head predicate, the local context may then represent many clause instances. Each call in the clause is then simulated based on the function values in the appropriate component of the input context vector, updating the local context accordingly. This phase may also change the number of clause instances represented by the local context, as some call instances fail and others produce multiple results. When all the calls have been simulated, the head of each clause instance represented by the local context is one possible result of the set of activation instances. An element of $CArg$, representing a set of result instances, is specified by the head under the local clause context.

For a given $x \in CArg$ representing a set of instances, and for a given predicate symbol p_i , the i 'th component of the output context vector will be a function whose value at x is $y \in CArg$ where y is the meet (\sqcup_{CArg}) of the results found using each clause for p_i as described in the last paragraph. In the concrete interpretation, the meet is set union. In abstract an interpretation the meet may represent a larger set than the union.

Members of Env are not required to specify the predicate symbols for the literals that make up the clause. It only must represent the possible substitutions or clause instances that are in force.

`Apply` simulates the effect on the local context of unifying one literal in the clause with a set of instances. The integer parameter indicates which clause literal is being unified. `Apply` is used both to activate the clause and to simulate call execution. In the former case, `Apply` simulates unification of the clause head with the set of activation instances, an element of Arg , whose execution the iteration function is simulating. In the latter, `Apply` is used by the iteration function to simulate the execution of each call in turn. For a call to p_i , the iteration function evaluates the i -th component of the input context vector on the set of activation instances specified in the local context. The result is a set of instances that is supplied to `Apply`. `Apply` simulates unification of each activation instance (from the local context) with each result instance (from the input context vector), composing appropriately the successful unifiers in the local clause context that is `Apply`'s result.

In the abstract interpretation `Apply` will not in general consider each pair of instances separately. The role of the abstract interpretation is to lump many instances together so they can be regarded collectively.

`Project` is used to get the set of instances for a particular clause literal specified by the local clause environment. That is needed to be able to evaluate the components of the input context vector, and to get the the head instances from the local clause context when all the calls have been simulated.

`GetArg` generates, as a description of a set of queries, a tuple of representations of sets of instances, one for each predicate in the program. Each component represents, for

the corresponding predicate, the set of activation instances in the set of queries. Recall that we make the simplifying assumption that each query contains only one literal.

GetEnv generates a local context for the given clause.

Canon and Freelnst go back and forth from Arg to the equivalent CArg. We remark on them further in Section 2.5.

Now we begin constructing our semantics.

2.3 The Semantic Domain

The collecting semantics of $\text{Pr} \in \text{Prog}$ is an n -placed vector of transforms, $\text{TFG}[[\text{Pr}]]$, where n is the number of predicates defined in Pr . The i 'th predicate's semantics is given by $\text{TFG}[[\text{Pr}]]_i$.²

$\text{TFG}[[\text{Pr}]]$ will be constructed in the next section. We now construct and name the transform space.

$\text{Trans} =_{df} \text{CArg} \rightarrow \text{CArg}$ is the domain of predicate denotations.

Trans^n is the domain of predicate environments or context vectors.³

Trans is a complete lattice with⁴

$$\begin{aligned} \top_{\text{Trans}} &=_{df} \lambda x:\text{CArg} . \top_{\text{CArg}} \\ \perp_{\text{Trans}} &=_{df} \lambda x:\text{CArg} . \perp_{\text{CArg}} \\ f \sqsubseteq_{\text{Trans}} g &\text{ if }_{df} \forall x \in \text{CArg} . f(x) \sqsubseteq_{\text{CArg}} g(x) \\ f \sqcup_{\text{Trans}} g &=_{df} \lambda x:\text{CArg} . f(x) \sqcup_{\text{CArg}} g(x) \end{aligned}$$

So Trans^n is also a complete lattice with join and order defined pointwise.

2.4 The Total Function Graph Construction

Now we construct the iteration function, $\text{Iterate}_{\text{TFG}}$, whose fixpoint defines the program *tfg* semantics. In the definitions we use the functions *vector*, *fix* and *index*. *vector* takes a function with domain $[1..n]$ and returns the n -placed vector whose i 'th component is the value of the function applied to i . *fix* is the least fixpoint operator. *index* takes a literal and gives back the integer in $[1..n]$ that identifies its predicate symbol. It is used to select the appropriate component of the context vector.

The *tfg* semantics of $\text{Pr} \in \text{Prog}$ are given by

$$\text{TFG: Prog} \rightarrow \text{Trans}^n$$

²Subscripting denotes component selection.

³The superscript denotes standard cross product.

⁴In these typed lambda expressions the colon separates the formal parameter from its domain. As usual, the dot separates the parameter list from the expression body. In the universally quantified sentence (third line) the dot is used to separate the quantification of the variable from the rest of the formula.

$$\mathbf{TFG}[[\mathbf{Pr}]] =_{df} \text{fix } \lambda\phi:\mathbf{Trans}^n. \mathbf{Iterate}_{\mathbf{TFG}}[[\mathbf{Pr}]](\phi)$$

where **Iterate** is defined by

$$\begin{aligned} \mathbf{Iterate}: \mathbf{Prog} &\rightarrow \mathbf{Trans}^n \rightarrow \mathbf{Trans}^n \\ \mathbf{Iterate}[[\mathbf{Pr}]](\phi) &=_{df} \text{vector}(\lambda i:[1..n]. \lambda x:\mathbf{CArg}. \\ &\quad \bigsqcup_{\mathbf{CArg}} \left\{ \mathbf{UseClause}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]](x, \phi) \mid \right. \\ &\quad \left. \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \text{ is a clause for the } i\text{'th predicate in } \mathbf{Pr} \right\}) \end{aligned}$$

and where **UseClause** and **Trace** are defined by⁵

$$\begin{aligned} \mathbf{UseClause}: \mathbf{Prog} &\rightarrow \mathbf{CArg} \times \mathbf{Trans}^n \rightarrow \mathbf{CArg} \\ \mathbf{UseClause}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]](x, \phi) &=_{df} \\ &\quad \mathbf{Canon}(\mathbf{Project}(0, \mathbf{Trace}(\mathbf{Apply}(0, \mathbf{GetEnv}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]], x, \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle), \\ &\quad \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle, 1, \phi), \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle)) \\ \mathbf{Trace}: \mathbf{Env} \times \mathbf{Clause} \times \mathbf{Z}^+ \times \mathbf{Trans}^n &\rightarrow \mathbf{Env} \\ \mathbf{Trace}(E, \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle, j, \phi) &=_{df} \\ &\quad \text{if } j > m \text{ then } E \\ &\quad \text{else } \mathbf{Trace}(\mathbf{Apply}(j, E, \\ &\quad \mathbf{FreInst}(\phi_{\text{index}(\mathbf{B}_j)})(\mathbf{Canon}(\mathbf{Project}(j, E, \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle))), \\ &\quad \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle), \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle, j + 1, \phi) \end{aligned}$$

UseClause is continuous in ϕ , since it looks at a finite number of ϕ values and is a finite composition of continuous functions. So, being pointwise continuous, **Iterate**_{TFG} is also continuous.

2.5 The Concrete Interpretation

We construct the concrete interpretation. It will induce a semantics that agrees with the standard semantics.

The concrete interpretation is based on functions from sets of term-tuples to sets of term-tuples. If $x \in \mathbf{Arg}_{\text{concrete}}$ then $\vec{t} \in x$ is a k -tuple of terms.

Let $\wp(S)$ denote the powerset of S . Then

$$\mathbf{Arg}_{\text{concrete}} =_{df} \wp(\mathbf{Term}^k)$$

is a complete lattice under

$$\begin{aligned} \sqsubseteq_{\mathbf{Arg}_{\text{concrete}}} &=_{df} \subseteq \text{ (the subset relation), with} \\ \top_{\mathbf{Arg}_{\text{concrete}}} &=_{df} \mathbf{Term}^k, \\ \perp_{\mathbf{Arg}_{\text{concrete}}} &=_{df} \{\} \text{ (the empty set), and} \end{aligned}$$

⁵ \mathbf{Z}^+ represents the positive integers.

$$\sqcup_{\text{Arg}_{\text{concrete}}} =_{df} \cup \text{ (set union) }.$$

As we have discussed, the set **Term** is partitioned into variable renaming equivalence classes. Because the results of two calls in the same equivalence class are also in the same equivalence class, we select a representative of each class and only define semantics on the representatives.

Let $Normalize : \text{Term}^k \rightarrow \text{Term}^k$ return a canonical representative of its argument's equivalence class under variable renaming.⁶ For example, $Normalize$ may rename the variables using a prefix of some enumeration of **Var**, ordered by first occurrence in the argument. However, to avoid variable name conflicts, we make the requirement that the result of $Normalize$ does not contain any variables used in the program text. This could be formalized by dividing the variable space into two infinite sets, one to be used in programs, the other to be used in normalized terms.

By $Normalize(\text{Term}^k)$ we denote the function's range. Thus

$$\text{CArg}_{\text{concrete}} =_{df} Normalize(\text{Term}^k)$$

is a complete lattice under

$$\begin{aligned} \sqsubseteq_{\text{CArg}_{\text{concrete}}} &=_{df} \sqsubseteq_{\text{Arg}_{\text{concrete}}}, \text{ with} \\ \sqcup_{\text{CArg}_{\text{concrete}}} &=_{df} \sqcup_{\text{Arg}_{\text{concrete}}}, \\ \top_{\text{CArg}_{\text{concrete}}} &=_{df} \wp(Normalize(\text{Term}^k)), \text{ and} \\ \perp_{\text{CArg}_{\text{concrete}}} &=_{df} \perp_{\text{Arg}_{\text{concrete}}}. \end{aligned}$$

Recall $E \in \text{Env}_{\text{concrete}}$ represents a set of clause instances. For $E \in \text{Env}_{\text{concrete}}$, and $e \in E$, $e_i \in \text{Arg}_{\text{concrete}}$ denotes the term-tuple that is the i 'th component of e , where $i \in [1..m+1]$, and m is the number of body literals in the clause.

$$\begin{aligned} \text{Env}_{\text{concrete}} &=_{df} \bigcup_{m \in \mathbb{N}} \wp((\text{Term}^k)^{m+1}) \\ \sqsubseteq_{\text{Env}_{\text{concrete}}} &=_{df} \subseteq \\ \sqcup_{\text{Env}_{\text{concrete}}} &=_{df} \cup \\ \top_{\text{Env}_{\text{concrete}}} &=_{df} (\text{Term}^k)^m \\ \perp_{\text{Env}_{\text{concrete}}} &=_{df} \{\} \end{aligned}$$

Next we define the functions in the interpretation.

$$\begin{aligned} \text{Apply} &: [0..m] \times \text{Env} \times \text{CArg} \times \text{Clause} \rightarrow \text{Env} \\ \text{Apply}_{\text{concrete}}(i, E, x, \langle H :- B_1, \dots, B_m \rangle) &=_{df} \left\{ \sigma(e) \mid e \in E, \vec{t} \in x, \text{ and } \sigma = \right. \\ &\quad \left. mfg(e_i, \vec{t}) \text{ is not fail} \right\} \\ \text{Project} &: [0..m] \times \text{Env} \times \text{Clause} \rightarrow \text{Arg} \\ \text{Project}_{\text{concrete}}(i, E, \langle H :- B_1, \dots, B_m \rangle) &=_{df} \left\{ \vec{t} \mid e \in E \text{ and } \vec{t} = e_i \right\} \\ \text{GetArg} &: \wp(\text{Literal}) \rightarrow \text{Arg} \end{aligned}$$

⁶The scope of variables is, of course, the whole term-tuple.

$$\text{GetArg}_{\text{concrete}}(\text{Queries}) =_{df} \text{vector}(\lambda i:[1..n]. \{ \vec{t} \mid \text{there is a literal } p_i(\vec{t}) \in \text{Queries}. \})$$

$$\text{GetEnv} : \text{Clause} \rightarrow \text{Env}$$

$$\text{GetEnv}_{\text{concrete}}[[H :- B_1, \dots, B_m]] =_{df} \{e\},$$

where e_1 is the tuple of arguments to H

and e_i is the tuple of arguments to B_{i-1} , for $i \in [2..m+1]$.

$$\text{Canon} : \text{Arg} \rightarrow \text{CArg}$$

$$\text{Canon}_{\text{concrete}}(x) =_{df} \{y \mid \exists \vec{t} \in x \text{ and } y = \text{Normalize}(\vec{t})\}$$

$$\text{Freelnst} : \text{CArg} \rightarrow \text{Arg}$$

$$\text{Freelnst}_{\text{concrete}}(Y) =_{df} \{ \vec{t} \mid \exists y \in Y, y = \text{Normalize}(\vec{t}), \text{ and } \\ \text{the variables in } y \text{ have not been used yet in the computation.} \}$$

Although $\text{Apply}_{\text{concrete}}$ does not use its fourth argument, abstract interpretations may need to. The same is true of $\text{Project}_{\text{concrete}}$. In general, $E \in \text{Env}$ represents a set of substitutions.

$\text{Freelnst}_{\text{concrete}}$ is not really a function, since it returns different members of the renaming equivalence class each time it's called. But we decline to make the tedious avoidance of variable name conflict more formal. In all our constructions, any result from $\text{Freelnst}_{\text{concrete}}$ is put through $\text{Canon}_{\text{concrete}}$ in such a way as to make the final results well defined.

2.6 The Minimal Function Graph Construction

This section constructs the mfg , which is a partial function, so we must add partial functions to our semantic domain. This is accomplished by following Jones and Mycroft[8]. We augment the activation description domain with a new bottom element that is less than all other elements and that represents unreachable calls, that is, calls on which the mfg is not defined.

As an intermediate step toward constructing the mfg , we construct the set of reachable activation descriptions from the tfg and a user specified set of queries.

$$\text{CArg}_{\text{bot}} =_{df} \text{CArg} \cup \{\text{bot}\}$$

$$x \sqsubseteq_{\text{CArg}_{\text{bot}}} y \text{ if}_{df} (x = \text{bot}) \text{ or } (x, y \in \text{CArg} \text{ and } x \sqsubseteq_{\text{CArg}} y)$$

$$x \sqcup_{\text{CArg}_{\text{bot}}} y =_{df} \text{if } x = \text{bot} \text{ then } \text{bot} \text{ else } x \sqcup_{\text{CArg}} y$$

$$\perp_{\text{CArg}_{\text{bot}}} =_{df} \text{bot}$$

$$\top_{\text{CArg}_{\text{bot}}} =_{df} \text{CArg}$$

With this new domain we can construct a partial function space in which **bot** represents the undefined function value.

$$\mathbf{PTrans} =_{df} \mathbf{CArg} \rightarrow \mathbf{CArg}_{\mathbf{bot}}$$

As usual, Order, meet, bottom and top are defined pointwise for \mathbf{PTrans} and for \mathbf{PTrans}^n . In particular,

$$\perp_{\mathbf{PTrans}^n} =_{df} \mathit{vector}(\lambda i:[1..n]. \lambda x:\mathbf{CArg}. \mathbf{bot}).$$

As in [8], $\phi_i(x) = \mathbf{bot}$ indicates that the i 'th predicate has not been called on x . $\phi_i(x) = \perp$ indicates that the i 'th predicate has been called on x , but hasn't succeeded.

$$\begin{aligned} \mathbf{MFG} : \mathbf{Prog} \times \wp(\mathbf{Literal}) &\rightarrow \mathbf{PTrans}^n \\ \mathbf{MFG}[[\mathbf{Pr}, \mathbf{Queries}]] &=_{df} \mathit{vector}(\lambda i:[1..n]. \lambda x:\mathbf{CArg}. \\ &\quad \text{if } x \in (\mathbf{Reachable}[[\mathbf{Pr}, \mathbf{Queries}]])_i \\ &\quad \text{then } (\mathbf{TFG}[[\mathbf{Pr}]])_i(x) \\ &\quad \text{else } \mathbf{bot}) \end{aligned}$$

$\mathbf{Reachable}$ generates the domain on which $\mathbf{MFG}[[\mathbf{Pr}, \mathbf{Queries}]]$ is defined (not \mathbf{bot}). For each of the n predicates, this is a subset of \mathbf{CArg} . Recall that, for simplicity, we assume that each query is a single literal.

$$\begin{aligned} \mathbf{Reachable} : \mathbf{Prog} \times \wp(\mathbf{Literal}) &\rightarrow (\wp(\mathbf{CArg}))^n \\ \mathbf{Reachable}[[\mathbf{Pr}, \mathbf{Queries}]] &=_{df} \mathit{fix} \lambda R: (\wp(\mathbf{CArg}))^n. \\ &\quad \mathbf{Iterate}_{\mathbf{Reachable}}[[\mathbf{Pr}, \mathbf{Queries}]](R) \end{aligned}$$

$$\begin{aligned} \mathbf{Iterate}_{\mathbf{Reachable}} : \mathbf{Prog} \times \wp(\mathbf{Literal}) &\rightarrow (\wp(\mathbf{CArg}))^n \rightarrow (\wp(\mathbf{CArg}))^n \\ \mathbf{Iterate}_{\mathbf{Reachable}}[[\mathbf{Pr}, \mathbf{Queries}]](R) &=_{df} \mathit{vector}(\lambda i:[1..n]. \\ &\quad \{ x:\mathbf{CArg} \mid \exists y \in R_j, \\ &\quad \quad x \in (\mathbf{CallsMade}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]](y, \mathbf{TFG}[[\mathbf{Pr}]]))_i, \\ &\quad \quad \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \text{ is a clause in } \mathbf{Pr} \text{ for the } j\text{'th predicate.} \} \\ &\quad \cup \{ (\mathbf{GetArg}[[\mathbf{Queries}]])_i \}) \end{aligned}$$

$$\begin{aligned} \mathbf{CallsMade} : \mathbf{Clause} &\rightarrow \mathbf{CArg} \times \mathbf{Trans}^n \rightarrow (\wp(\mathbf{Arg}))^n \\ \mathbf{CallsMade}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]](x, \phi) &=_{df} \mathit{vector}(\lambda i:[1..n]. \\ &\quad \{ y : \mathbf{CArg} \mid y = \mathbf{Canon}(\mathbf{Project}(\ell, \mathbf{Trace} \\ &\quad \quad \mathbf{Apply}(0, \mathbf{GetEnv}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]], x, \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle), \\ &\quad \quad \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_{\ell-1} \rangle, 1, \phi)), \\ &\quad y \neq \perp, \\ &\quad \ell \in [1..m], \text{ and} \\ &\quad \mathit{index}(\mathbf{B}_\ell) = i \}) \end{aligned}$$

$\mathbf{Iterate}_{\mathbf{Reachable}}$ is continuous in R because it looks at each member of R in isolation.

3 The Partial Function Graph Semantics

We now present the *partial function graph* (*pfg*) semantics. It is the limit of an iteration function that is not monotonic, but that is sufficiently well behaved that the limit exists. In Section 3.2 we will show that this semantics essentially agrees with the *mfg* semantics. The reason for considering this semantics is that it is almost as informative as the *mfg* semantics and, for finite abstract domains, can be computed directly, obviating the necessity of computing the *tfg*, which for large abstract domains, would be impractical.

$(\mathbf{TFG}[[\text{Pr}]])_i$ is defined on all of CArg , which may be very large. Granted, if we are to guarantee termination of our analysis, we expect to have to require that $\text{CArg}_{\text{approx}}$ be finite, since otherwise it is not clear how to guarantee that $\mathbf{Reachable}[[\text{Pr}, \text{Queries}]]$ is finite. For example, if we have the clause

trouble(X) :- trouble(f(X,X)).

and pose the query

?- trouble(X)

then $\{\text{trouble}(X)\}, \{\text{trouble}(f(X,X))\}, \{\text{trouble}(f(f(X,X),f(X,X)))\}, \dots$ would all be reachable activation sets.

Since it appears that $\text{CArg}_{\text{approx}}$ must be finite, we can expect $\mathbf{TFG}[[\text{Pr}]]$ to be computable in the approximate interpretation. But we can also expect it to be impractical for large approximation domains, which would afford greater accuracy.

3.1 Construction of the PFG

The *pfg* is defined as the limit of an iteration function that is almost identical to the *tfg* iteration function, except that it incorporates reachability. As the *pfg* iteration function is applied repeatedly, the domain on which the result is defined grows to reflect new reachable activations.

The *pfg* iteration function is not monotonic. But this section proves that it is sufficiently well behaved that the limit exists.

Before we can use them in the construction of $\mathbf{Iterate}_{\mathbf{PFG}}$ we must extend the types over which several functions are defined. **UseClause**, **Trace**, and **CallsMade** all have one argument of type **Trans**. Each of these arguments will now be treated as having type **PTrans**. The change in these core-semantic functions can be handled by making one small additional requirement of one of the interpretation's basic functions, **FreeInst**, since this is the function that may now be applied to **bot**. We require that $\mathbf{FreeInst}(\text{bot}) = \perp$. With these changes in place, we can pass $\phi \in \mathbf{PTrans}$ to **CallsMade** and to **UseClause** with impunity.

$$\mathbf{PFG} : \text{Prog} \times \text{Queries} \rightarrow \mathbf{PTrans}^n$$

$\mathbf{PFG}[[\mathbf{Pr}, \mathbf{Queries}]] =_{df} \lim_{\gamma} \mathbf{Iterate}_{\mathbf{PFG}}[[\mathbf{Pr}, \mathbf{Queries}]]^{\gamma}(\perp_{\mathbf{PTrans}^n})$
 where $f^0(X) =_{df} X$,
 $f^{\delta+1}(X) =_{df} f \circ f^{\delta}(X)$, and
 $f^{\gamma} =_{df} \bigsqcup \{f^{\delta} \mid \delta < \gamma\}$ for limit ordinals, γ .

$\mathbf{Iterate}_{\mathbf{PFG}}: \mathbf{Prog} \times \wp(\mathbf{Literal}) \rightarrow \mathbf{PTrans} \rightarrow \mathbf{PTrans}$
 $\mathbf{Iterate}_{\mathbf{PFG}}[[\mathbf{Pr}, \mathbf{Queries}]](\phi) =_{df} \mathbf{vector}(\lambda i: [1..n] . \lambda x: \mathbf{CArg} .$
if $\phi_i(x) \neq \mathbf{bot}$
then $\bigsqcup_{\mathbf{Arg}} \{ \mathbf{UseClause}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]](x, \phi) \mid$
 $\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m$ is a clause for the i 'th predicate in \mathbf{Pr} $\}$
else if $(\exists y \in \mathbf{CArg})$ and $(\exists j \in [1..n])$ and
 $(\exists \langle \mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m \rangle \in \mathbf{Pr})$ such that
 $\phi_j(y) \neq \mathbf{bot}$,
 $\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m$ is a clause for the j 'th predicate and
 $x \in (\mathbf{CallsMade}[[\mathbf{H} :- \mathbf{B}_1, \dots, \mathbf{B}_m]](y, \phi))_i$ or
 $(x \in \mathbf{Queries})$
then \perp
else \mathbf{bot}

The *then* part of $\mathbf{Iterate}_{\mathbf{PFG}}$ is identical to $\mathbf{Iterate}_{\mathbf{TFG}}$. On the reachable calls $\mathbf{PFG}[[\mathbf{Pr}, \mathbf{Queries}]]$ will give the same results as will $\mathbf{TFG}[[\mathbf{Pr}]]$ (see Section 3.2).

The *else* part closely resembles $\mathbf{Iterate}_{\mathbf{Reachable}}$, though $\mathbf{Iterate}_{\mathbf{Reachable}}$ uses the fixpoint, $\mathbf{TFG}[[\mathbf{Pr}]]$, where $\mathbf{Iterate}_{\mathbf{PFG}}$ uses the partial solution (in a Kleene's sequence calculation), which is not complete. So, for example, in the precise interpretation, if the result instance set for \mathbf{B}_1 is incomplete *vis-à-vis* the fixpoint, an incomplete activation set for \mathbf{B}_2 may be reported by $\mathbf{CallsMade}$. This is how it can happen that

$$\mathbf{MFG}[[\mathbf{Pr}, \mathbf{Queries}]] \not\sqsubseteq_{\mathbf{PTrans}^n} \mathbf{PFG}[[\mathbf{Pr}, \mathbf{Queries}]].$$

As has been mentioned already, this difference is insignificant for applications that require information about the reachability of individual activations, rather than the sets of activations that arise in our particular model.

Unfortunately, $\mathbf{Iterate}_{\mathbf{PFG}}$ is not continuous. The sequence whose limit defines $\mathbf{PFG}[[\mathbf{Pr}, \mathbf{Queries}]]$ may not converge by step ω . However, as long as the approximation domain is finite, the limit is certain to be reached at a finite stage.

To ensure that $\mathbf{PFG}[[\mathbf{Pr}, \mathbf{Queries}]]$ is well defined, we demonstrate in Lemma 1 that repeated application of $\mathbf{Iterate}_{\mathbf{PFG}}[[\mathbf{Pr}, \mathbf{Queries}]]$ to $\perp_{\mathbf{PTrans}^n}$ generates an increasing sequence. Then the sequence must converge before the index reaches the successor cardinal to the cardinality of \mathbf{PTrans} . This proves that the limit exists and that \mathbf{PFG} is well defined.

We use the notation $f|_S$ to denote restriction of the function f to the set S . We also use this notation with equal sized vectors of functions and sets to denote the pointwise restriction.

For convenience, let us define

$$\begin{aligned}\psi_\gamma &=_{df} \text{Iterate}_{\mathbf{PFG}}[[\text{Pr}, \text{Queries}]]^\gamma(\perp_{\mathbf{PTrans}^n}), \text{ and} \\ RS_\gamma &=_{df} \text{vector}(\lambda i \in [1..n] . \{x:\text{CArg}(\psi_\gamma)_i(x) \neq \text{bot}\}).\end{aligned}$$

At this point we begin to omit the subscripts from $\sqcup_{\mathbf{PTrans}^n}$ and $\sqsubseteq_{\mathbf{PTrans}^n}$, and the syntactic arguments from $\text{Iterate}_{\mathbf{PFG}}[[\text{Pr}, \text{Queries}]]$ and $\text{Iterate}_{\mathbf{TFG}}[[\text{Pr}, \text{Queries}]]$, etc.

Lemma 1. $\forall \delta < \gamma . \psi_\delta \sqsubseteq_{\mathbf{PTrans}^n} \psi_\gamma$.

We show this by induction on γ . The basis is vacuous, and the result is obvious for γ a limit ordinal. Let us verify that it holds for successor ordinals, $\gamma + 1$. For this it suffices to show that $\psi_\gamma \sqsubseteq \psi_{\gamma+1}$.

We have, $\forall \delta < \gamma . \psi_\delta \sqsubseteq \psi_\gamma$ by the induction hypothesis. Let us assume that $\delta < \gamma$. $\text{Iterate}_{\mathbf{TFG}}(\psi_\delta) \sqsubseteq \text{Iterate}_{\mathbf{TFG}}(\psi_\gamma)$, by monotonicity of $\text{Iterate}_{\mathbf{TFG}}$.

Since $\text{Iterate}_{\mathbf{TFG}}(\psi_\delta)|_{RS_\delta} = \text{Iterate}_{\mathbf{PFG}}(\psi_\delta)|_{RS_\delta}$,

it follows that $\text{Iterate}_{\mathbf{PFG}}(\psi_\delta)|_{RS_\delta} \sqsubseteq \text{Iterate}_{\mathbf{PFG}}(\psi_\gamma)|_{RS_\delta}$, since $RS_\delta \subseteq RS_\gamma$.

On the other hand, $x \notin (RS_\delta)_i \Rightarrow \text{Iterate}_{\mathbf{PFG}}(\psi_\delta)_i(x) \sqsubseteq \perp$,

and by inspection of the definition of $\text{Iterate}_{\mathbf{PFG}}$ we see that

$$\text{Iterate}_{\mathbf{PFG}}(\psi_\delta)_i(x) = \perp \Rightarrow \perp \sqsubseteq \text{Iterate}_{\mathbf{PFG}}(\psi_\gamma)_i(x)$$

So, $\text{Iterate}_{\mathbf{PFG}}(\psi_\delta)|_{\overline{RS_\delta}} \sqsubseteq \text{Iterate}_{\mathbf{PFG}}(\psi_\gamma)|_{\overline{RS_\delta}}$ follows.⁷

Putting these together, we have that $\text{Iterate}_{\mathbf{PFG}}(\psi_\delta) \sqsubseteq \text{Iterate}_{\mathbf{PFG}}(\psi_\gamma) = \psi_{\gamma+1}$.

Now, if $\gamma = \delta + 1$ for some δ , we are done.

On the other hand, if γ is a limit ordinal, then $\psi_\gamma = \sqcup_{\delta < \gamma} \text{Iterate}_{\mathbf{PFG}}(\psi_\delta)$.

But, since $\forall \delta < \gamma . \text{Iterate}_{\mathbf{PFG}}(\psi_\delta) \sqsubseteq \text{Iterate}_{\mathbf{PFG}}(\psi_\gamma) = \psi_{\gamma+1}$,

it follows that $\psi_\gamma \sqsubseteq \psi_{\gamma+1}$. ■

3.2 The PFG is Safe

Now we argue that **PFG** is a safe approximation to **MFG**. It will be shown that over the calls specified by **Reachable**, all three of **MFG**, **PFG**, and **TFG** agree exactly. This is sufficient, since elsewhere **MFG**, is **bot**.

First, let us define $RS \in (\wp(\text{CArg}))^n$ to represent the reachable activations according to **PFG**:

$$RS =_{df} \text{vector}(\lambda i \in [1..n] . \{x:\text{CArg}[\mathbf{PFG}[[\text{Pr}, \text{Queries}]]]_i(x) \neq \text{bot}\}).$$

So, $RS = \lim_\gamma RS_\gamma$.

⁷ $\overline{RS_\delta}$ denotes the pointwise compliment of RS_δ .

Now we state the theorem that ensures the safeness of the *pfg* semantics as an approximation to the *mfg* semantics. We again elide the operator subscripts and syntactic arguments.

Theorem. $\mathbf{MFG}|_{\text{Reachable}} = \mathbf{PFG}|_{\text{Reachable}} = \mathbf{TFG}|_{\text{Reachable}}$

We show that $\mathbf{PFG}|_{RS} = \mathbf{TFG}|_{RS}$. Since it is obvious that $\text{Reachable} \subset RS$, the Theorem will follow from this and the definition of \mathbf{MFG} .

Claim: $\mathbf{PFG}|_{RS} \subseteq \mathbf{TFG}|_{RS}$.

Since $\mathbf{PFG} =_{df} \lim_{\gamma} \psi_{\gamma}$, it suffices to show that $\psi_{\gamma} \subseteq \mathbf{TFG}$ for all γ . This is done by transfinite induction on γ .

i. Basis: $\psi_0 \subseteq \mathbf{TFG}$. Trivial.

ii. Successor Ordinals: $\psi_{\delta} \subseteq \mathbf{TFG} \Rightarrow \psi_{\delta+1} \subseteq \mathbf{TFG}$.

$\psi_{\delta+1} = \text{Iterate}_{\mathbf{PFG}}(\psi_{\delta})$ by definition of ψ_{γ} .

$\text{Iterate}_{\mathbf{PFG}}(\psi_{\delta}) \subseteq \text{Iterate}_{\mathbf{TFG}}(\psi_{\delta})$

because they are equal on $\{x | \psi_{\delta}(x) \neq \text{bot}\}$ by inspection,
and elsewhere it is trivially true.

$\text{Iterate}_{\mathbf{TFG}}(\psi_{\delta}) \subseteq \text{Iterate}_{\mathbf{TFG}}(\mathbf{TFG})$

by the induction hypothesis and monotonicity of $\text{Iterate}_{\mathbf{TFG}}$.

But, $\text{Iterate}_{\mathbf{TFG}}(\mathbf{TFG}) = \mathbf{TFG}$ by definition.

iii. Limit Ordinals: $(\forall \delta < \gamma . \psi_{\delta} \subseteq \mathbf{TFG}) \rightarrow (\psi_{\gamma} \subseteq \mathbf{TFG})$.

This follows immediately.

Claim: $\mathbf{TFG}|_{RS} \subseteq \mathbf{PFG}|_{RS}$

If $x \in RS_i$ then for all clauses, $\text{CallsMade}[[H :- B_1, \dots, B_m]](x, \mathbf{PFG})_j \subset RS_j$, for all $j \in [1..n]$.

Let $\widehat{\mathbf{PFG}}$ be \mathbf{PFG} except values of bot are replaced by \perp .

Observe that $(\text{Iterate}_{\mathbf{TFG}}(\widehat{\mathbf{PFG}}))|_{RS} = \widehat{\mathbf{PFG}}|_{RS} = \mathbf{PFG}|_{RS}$.

Since \mathbf{TFG} is the least fixpoint of $\text{Iterate}_{\mathbf{TFG}}$, we have the desired result.

■

Of course, \mathbf{MFG} could be constructed from \mathbf{PFG} by performing reachability analysis.

This section has shown that \mathbf{PFG} is a safe approximation to \mathbf{MFG} . That is important because \mathbf{MFG} would be rather expensive to compute for large abstract domains because it entails computing \mathbf{TFG} . \mathbf{PFG} is less expensive to compute because results do not have to be collected on all $x \in \text{CArg}$, just on those that appear to be reachable.

Not only is \mathbf{PFG} safe, but it is a strong approximation. Where it is defined, it agrees exactly with \mathbf{TFG} . It is just defined on some $x \in \text{CArg}$ that \mathbf{MFG} is not defined on. However, such x always represent subsets of sets on which \mathbf{MFG} is defined. Experiments that will be reported elsewhere seem to suggest that results are collected by the *pfg* method on relatively few $x \in \text{Arg}$ that are not reachable.

4 Conclusion

This paper has constructed an abstract interpretation framework for logic programs based on a *minimal function graph (mfg)* semantics. Our framework has been designed to accommodate large abstract domains efficiently, making stronger automatic inferences practical. This is the primary basis of our argument that our framework is more practical than the elegant framework of Jones and Søndergaard[7].

We have used our framework to construct and implement a data flow analysis for mode inference that gives stronger inferences than any published analysis. That result will be reported in a forthcoming paper.

Our *mfg* cannot be computed directly. But we have shown that the *partial function graph (pfg)*, which is defined here and can be computed directly for finite approximation domains, is virtually indistinguishable from the *mfg*.

References

- [1] Jung-Herng Chang, Alvin M. Despain, Doug DeGroot, "AND-Parallelism of Logic Programs Based On a Static Data Dependency Analysis," IEEE 1985 Spring CompCon.
- [2] John S. Conery, "The AND/OR Process Model for Parallel Interpretation of Logic Programs," Ph.D. Dissertation, Department of Information and Computer Science, University of California-Irvine, TR 204, 1983.
- [3] Patrick Cousot and Radhia Cousot, "Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fix-points," *Conference Record of the Fourth Annual ACM Symposium on Principles of Programming Languages*, 1977.
- [4] Saumya K. Debray, "Synthesizing Control Strategies for AND-Parallel Logic Programs," Department of Computer Science technical report, TR 87-12, University of Arizona, 1987.
- [5] Saumya K. Debray, "Approximation Domains for Efficient Flow Analysis of Logic Programs," Department of Computer Science technical report, TR 87-9, University of Arizona, 1987.
- [6] Saumya K. Debray, David S. Warren, "Automatic Mode Inference for Prolog Programs," *Proceedings of the 1986 Symposium on Logic Programming*.
- [7] Neil D. Jones, Harald Søndergaard, "A Semantics-Based Framework for the Abstract Interpretation of Prolog," To appear in S. Abramsky and C. Hankin (eds.), *Abstract Interpretation of Declarative Languages*, Ellis Horwood.

- [8] Neil D. Jones, Alan Mycroft, "Data Flow Analysis of Applicative Programs Using Minimal Function Graphs: Abridged Version," *Conference Record of the Thirteenth Annual ACM Symposium on Principles of Programming Languages*, 1986.
- [9] Laxmikant Vasudeo Kalé, "Parallel Architectures for Problem Solving," Ph.D. Dissertation, University of New York at Stony Brook, 1985.
- [10] J. W. Lloyd, *Foundations of Logic Programming*, Springer-Verlag, 1984.
- [11] C. S. Mellish, "The Automatic Generation of Mode Declarations for Prolog Programs," Research Report 163, Department of Artificial Intelligence, University of Edinburgh, 1981.
- [12] C. S. Mellish, "Abstract Interpretation of Prolog Programs," *Third International Conference on Logic Programming*, LNCS 225, Springer-Verlag, 1986.
- [13] David A. Plaisted, "The Occur-Check Problem in Prolog," *1984 International Symposium on Logic Programming*, 1984.
- [14] H. Søndergaard, "An Application of Abstract Interpretation of Logic Programs: Occur Check Reduction," *ESOP 86*, LNCS 213, Springer-Verlag, 1986.
- [15] D.H.D. Warren, "Implementing Prolog - Compiling Predicate Logic Programs," Research Reports 39 and 40, Department of Artificial Intelligence, University of Edinburgh, 1977.