Section 4 illustrates the need to interpret programs over domains other than abstractions of the accumulating semantics. Some program analyses not naturally expressed by abstracting either an accumulating or an *instrumented* semantics are exemplified, showing the need for more sophisticated analysis techniques, and an overview is given of some alternative approaches including tree grammars.

The idea of an "instrumented" semantics is introduced and correctness is discussed. This section is problem-oriented, with simulation techniques chosen ad hoc to fit the analysis problem and the language being analysed. It thus centers more around programs' operational behaviour than the structure of their domains, with particular attention to describing the set of program states reachable in computations on given input data, and to finite description of the set of all computations on given input. The section ends by describing approaches to abstract interpretation of Prolog.

2 Basic Concepts and Problems to be Solved

We begin with parity analysis of a very simple example program, and introduce basic concepts only as required. We discuss imperative programs without procedure calls since this familiar program class has a simple semantics and is most often treated in the analysis algorithms found in compiling textbooks. Later sections will discuss functional and logic programs, but many of their analysis problems are also visible, and usually in simpler form, in the imperative context. Throughout this section the reader is encouraged to ask himself "what is the analogue of this concept in a functional or logic programming framework?".

An example program, where $\div$ stands for integer division (and program points A,...,G have been indicated for future reference):

A: **while** $n \neq 1$ **do**
    B: **if** n even
          then (C: $n := n \div 2$; D:)
          else (E: $n := 3 * n + 1$; F:)
        fi
    od
G:

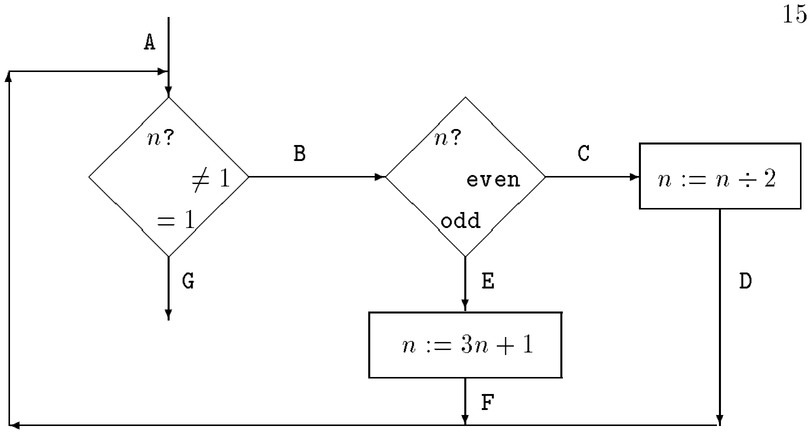

Side remark: Collatz' problem in number theory amounts to determining whether this program terminates for all positive initial n. To our knowledge it is still unsolved.

2.1 A Naive Analysis of the Simple Program

Abstraction of a single execution If this program is run with initial value $n = 5$, then n takes on values $5, 16, 8, 4, 2$ at point B, values $16, 8, 4, 2$ at C, etc. Using $\top$ to represent "either even or odd" the results of this single run can be abstracted as:

n at A	n at B	n at C	n at D	n at E	n at F	n at G
odd	$\top$	even	$\top$	odd	even	odd

Extension to all possible executions This result was obtained by performing *one* execution completely, and then abstracting its outcome. Such an analysis may of course not terminate, and it does not as wished describe all executions. The question is: how to obtain even-odd information valid for *all possible computations*? A natural way is to simulate the computation, but to do the computation using the abstract values

$$\text{Abs} = \{\bot, \text{even}, \text{odd}, \top\}$$

instead of natural numbers, each representing a set of possible values of n; and to ensure that all possible control flow paths are taken.

Doing this informally, we can see that if n is odd at program entry, it will always be even at points C and F, always odd at point E, sometimes even and sometimes odd at points B and D, and odd at G, provided control ever reaches G. Individual operations can be simulated by known properties of numbers, e.g. $3n+1$ is even if n is odd and odd if n is even, while $n \div 2$ can be either even or odd.

Simulating the whole program is not as straightforward as simulating a single execution. The reason was mentioned before: execution over abstract values cannot in general be deterministic, since it must take account of all possible execution sequences on real data satisfying the abstract data description.

Towards a less naive analysis procedure The very earliest data-flow analysis algorithms amounted to glorified interpreters, and proceeded by executing the program symbolically, keeping a record of the desired flow information (abstract values) as the interpretation proceeded. Such algorithms, which in essence traced all possible control paths through the program, were very slow and often incorrect. They further suffered from a number of problems of semantic nature, for example difficulties in seeing how to handle nondeterminism due to tests with insufficient information to recognize their truth or falsity, convergence and divergence of control paths, loops and nontermination.

Better methods were soon developed to solve these problems, including

- putting a partial order on the abstract data values, so they always change in the same direction during abstract interpretation, thus reducing termination problems
- storing flow information in a separate data structure, usually bound to program points (such as entry points to "basic blocks", i.e. maximal linear program segments)
- constructing from the program a system of "data-flow equations", one for each program point
- solving the data-flow equations (usually by computing their greatest fixpoint or least fixpoint).

Much more efficient algorithms were developed and some theoretical frameworks were developed to make the new methods more precise; [Hecht, 1977], [Kennedy, 1981] and [Aho, Sethi and Ullman, 1986] contain good overviews.

None of the "classical" approaches to program analysis can, however, be said to be formally related to the semantics of the language whose programs were being analysed. Rather, they formalized and tightened up methods used in existing practice. In particular none of them was able to include *precise execution as a special case of abstract interpretation* (albeit an uncomputable one). This was first done in [Cousot, 1977a], the seminal paper relating abstract interpretation to program semantics.

2.2 Accumulating Semantics for Imperative Programs

The approach of [Cousot, 1977a] is appealing because of its generality: it expresses a large number of special program analyses in a common framework. In particular, this makes questions of safety (*i.e.* correctness) much easier to formulate and answer, and sets up a framework making it possible to relate and compare the precision of a range of different program analyses. It is solidly based in semantics, and precise execution of the program is included as a special case. This implies program verification may also be based on the accumulating semantics, a theme developed further in [Cousot, 1977b] and several subsequent works.

The ideas of [Cousot, 1977a] have had a considerable impact on later work in abstract interpretation, for example [Mycroft, 1981], [Muchnick, 1981], [Burn, 1986], [Donzeau-Gouge, 1978], [Nielson, 1982], [Nielson, 1984], [Mycroft, 1987]).

2.2.1 Overview of the Cousot Approach

The article [Cousot, 1977a] begins by presenting an operational semantics for a simple flow chart language. It then develops the concept of what we call the *accumulating semantics* (the same as Cousots' static semantics and some others' collecting semantics). This associates with each program point the set of all memory stores that can ever occur when program control reaches that point, as the program is run on data from a given initial data space. It was shown in [Cousot, 1977a] that a wide variety of flow analyses (but not all!) may be realized by finding finitely computable approximations to the accumulating semantics.

The (sticky) accumulating semantics maps program points to sets of program stores. The set $\wp(\text{Store})$ of all sets of stores forms a *lattice* with set inclusion $\subseteq$ as its partial order, so any two store sets A, B have least upper bound $A \cup B$ and greatest lower bound $A \cap B$. The lattice $\wp(\text{Store})$ is *complete*, meaning that any collection of sets of stores has a least upper bound in $\wp(\text{Store})$, namely its union.

Various approximations can be expressed by simpler lattices, connected to $\wp(\text{Store})$ by an *abstraction* function $\alpha : \wp(Store) \rightarrow \text{Abs}$ where Abs is a lattice of descriptions of sets of stores. Symbol $\sqcup$ is usually used for the least upper bound operation on Abs, $\sqcap$ for the greatest lower bound, and $\top, \bot$ for the least, resp. greatest elements of Abs.

An abstraction function is most often used together with a dual *concretization* function $\gamma : \text{Abs} \rightarrow \wp(\text{Store})$, and the two are required to satisfy natural conditions (given later).

For a one-variable program we could use as Abs the lattice with elements

$$\{\bot, \top, \text{even}, \text{odd}\},$$

where the abstraction of any nonempty set of even numbers is lattice element "even", and the concretization of lattice element "even" is the set of all even numbers. Abstract interpretation may thus be thought of as executing the program over a lattice of imprecise but computable *abstract store descriptions* instead of the precise and uncomputable accumulating semantics lattice.

In practice computability is often achieved by using a *noetherian* lattice, i.e. one without infinite ascending chains. More general lattices can, however, be used, cf. the Cousots' "widening" techniques, or the use of grammars to describe infinite sets finitely.

Let p_0 be the program's initial program point and let p be another program point. The set of store configurations that can be reached at program point p, starting from a set S_0 of possible initial stores is defined by:

$$\mathrm{acc}_p = \{s \ \mid (\mathrm{p},\, s) = \mathrm{next}^n((\mathrm{p}_0,\, s_0)) \text{ for some } s_0 \in S_0, n \geq 0 \ \}$$

The accumulating semantics thus associates with each program point the set $\mathrm{acc}_p \subseteq \mathrm{Store}$.

2.2.2 Accumulating Semantics of the Example Program

For the example program there is only one variable, so a set of stores has form

$$\{[n \mapsto\ a_1], [n \mapsto\ a_2], [n \mapsto\ a_3], \ldots\}$$

For notational simplicity we can identify this with the set $\{\mathrm{a}_1,\, \mathrm{a}_2,\, \mathrm{a}_3,\, \ldots\}$ (an impossible simplification if the program has more than one variable). Given initial set $S_0 = \{5\}$ the sets of stores reachable at each program point are:

acc_A	acc_B	acc_C	acc_D	acc_E	acc_F	acc_G
{5}	{5,16,8,4,2}	{16,8,4,2}	{8,4,2,1}	{5}	{16}	{1}

The following *data-flow* equations have a unique least fixpoint by completeness of $\wp(\mathrm{Store})$, and it is easy to see that their fixpoint is exactly the tuple of sets of reachable stores as defined above.

$$\begin{array}{lcl}
\mathrm{acc}_A & = & S_0 \\
\mathrm{acc}_B & = & (\mathrm{acc}_A \cup\ \mathrm{acc}_D \cup\ \mathrm{acc}_F) \cap \{n \mid\ n \in \{0,1,2,\ldots\} \setminus \{1\}\} \\
\mathrm{acc}_C & = & \mathrm{acc}_B\ \cap\ \{n \mid n \in \{0,2,4,\ldots\}\} \\
\mathrm{acc}_D & = & \{n \div 2 \mid n \in \mathrm{acc}_C\} \\
\mathrm{acc}_E & = & \mathrm{acc}_B \cap\ \{n \mid n \in \{1,3,5,\ldots\}\} \\
\mathrm{acc}_F & = & \{3n+1 \mid n \in \mathrm{acc}_E\}
\end{array}$$

$$\mathrm{acc}_G = (\mathrm{acc}_A \cup \mathrm{acc}_D \cup \mathrm{acc}_F) \cap \{1\}$$

The equation set can be derived mechanically from the given program's syntax, e.g. as seen in [Cousot, 1977a] or [Nielson, 1982].

2.2.3 Abstract Interpretation of the Example Program

The *abstraction function* $\alpha : \wp(\mathrm{Store}) \to \mathrm{Abs}$ below may be used to abstract a set of stores, where $\mathrm{Abs} = \{\perp, \mathrm{even}, \mathrm{odd}, \top\}$:

$$\alpha(S) = \begin{cases} \perp & \text{if S} = \{\}, \text{ else} \\ \mathrm{even} & \text{if S} \subseteq \{0, 2, 4, \ldots\}, \text{ else} \\ \mathrm{odd} & \text{if S} \subseteq \{1, 3, 5, \ldots\}, \text{ else} \\ \top & \end{cases}$$

Defining $\perp \sqsubseteq \mathrm{even} \sqsubseteq \top$ and $\perp \sqsubseteq \mathrm{odd} \sqsubseteq \top$ makes Abs into a partially ordered set. Least upper and greatest lower bounds $\sqcup, \sqcap$ exist so it is also a lattice.

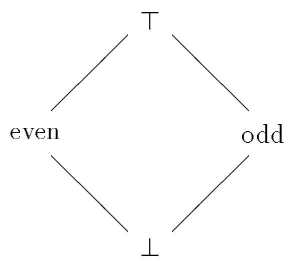

Applying α to the sets of reachable stores yields the following:

abs_A	abs_B	abs_C	abs_D	abs_E	abs_F	abs_G
odd	$\top$	even	$\top$	odd	even	odd

Abstraction of the set of all runs This method is still unsatisfactory for describing all computations since the value sets involved are unbounded and possibly infinite. But we may model the equations above by applying α to the sets involved. The abstraction function α just given is easily seen to be monotone, so set inclusion $\subseteq$ in the world of actual computations is modelled by $\sqsubseteq$ in the world of simulated computations over Abs. Union is the least upper bound over sets, so it is natural to model $\cup$ by $\sqcup$, and similarly to model $\cap$ by $\sqcap$.

The arithmetic operations are faithfully modelled as follows, using familiar properties of natural numbers:

$$f_{n\div 2}(abs) = \begin{cases} \perp & \text{if } abs = \perp \\ \top & \text{else} \end{cases}$$

$$f_{3n+1}(abs) = \begin{cases} \perp & \text{if } abs = \perp\text{, else} \\ \text{even} & \text{if } abs = \text{odd, else} \\ \text{odd} & \text{if } abs = \text{even, else} \\ \top & \text{if } abs = \top \end{cases}$$

This yields the following system of *approximate data-flow equations*, describing the program's behaviour on Abs:

$$\begin{array}{lcll} abs_A & = & \alpha(S_0) & \\ abs_B & = & (abs_A \sqcup abs_D \sqcup abs_F) \sqcap \top & (\text{“}\sqcap \top\text{” may be omitted}) \\ abs_C & = & abs_B \sqcap \text{even} & \\ abs_D & = & f_{n\div 2}(abs_E) & \\ abs_E & = & abs_B \sqcap \text{odd} & \\ abs_F & = & f_{3n+1}(abs_E) & \\ abs_G & = & (abs_A \sqcup abs_D \sqcup abs_F) \sqcap \text{odd} & \end{array}$$

Remark Here $f_{n\div 2}$ and f_{3n+1} were defined ad hoc; a systematic way to define them will be seen in section 2.3.

The lattice Abs is also complete. The operators $\sqcap$, $\sqcup$, $f_{n\div 2}$ and f_{3n+1} are monotone, so the equation system has a (unique) least fixpoint. The abstraction function α is easily seen to be monotone, so if it also were a homomorphism with respect to $\cup$, $\sqcup$ and $\cap$, $\sqcap$, the least solution to the approximate flow equations would be exactly

$$abs_A = \alpha(\text{acc}_A), \ldots, abs_G = \alpha(\text{acc}_G).$$

It is, however, *not* homomorphic since for example

$$\alpha(\{2\}) \sqcap \alpha(\{4\}) = \text{even} \neq \perp = \alpha(\{2\} \cap \{4\})$$

On the other hand the following *do* hold:

$$\begin{array}{rclcrcl}
\alpha(A) \sqcup \alpha(B) & = & \alpha(A \cup B) & \quad & f_{n \div 2}(\alpha(A)) & \sqsupseteq & \alpha(\{n \div 2 \mid n \in A\}) \\
\alpha(A) \sqcap \alpha(B) & \sqsupseteq & \alpha(A \cap B) & \quad & f_{3n+1}(\alpha(A)) & \sqsupseteq & \alpha(\{3n+1 \mid n \in A\})
\end{array}$$

Using these, it is easy to see by inspection of the two equation systems (more formally: a simple fixpoint induction) that their least fix points are related by:

$$\begin{array}{rcl}
abs_A & \sqsupseteq & \alpha(\mathrm{acc}_A), \\
abs_B & \sqsupseteq & \alpha(\mathrm{acc}_B), \\
 & \vdots & \\
abs_G & \sqsupseteq & \alpha(\mathrm{acc}_G)
\end{array}$$

Following is the iterative computation of the least fixpoint, assuming $S_0 = \{5\}$:

abs_A	abs_B	abs_C	abs_D	abs_E	abs_F	abs_G	iteration
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	0
odd	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	1
odd	odd	$\bot$	$\bot$	$\bot$	$\bot$	odd	2
odd	odd	$\bot$	$\bot$	odd	$\bot$	odd	3
odd	odd	$\bot$	$\bot$	odd	even	odd	4
odd	$\top$	$\bot$	$\bot$	odd	even	odd	5
odd	$\top$	even	$\bot$	odd	even	odd	6
odd	$\top$	even	$\top$	odd	even	odd	7, 8, ...

The conclusion is that n is always even at points C and F, and always odd at E and G.

2.2.4 An Optimization Using the Results of the Analysis

The flow analysis reveals that the program could be made somewhat more efficient by "unrolling" the loop after F. The reason is that tests "$n \neq 1$" and "n even" must be both be true in the iteration after F, so they need not be performed. The result is

> **while** $n \neq 1$ **do if** n even **then** $n := n \div 2$ **else** $n := (3*n+1) \div 2$
> **fi od**

which avoids the two tests every time n is odd. In practice, one of the main reasons for doing abstract interpretation is to find out when such optimizing transformations may be performed.

2.2.5 Termination

The least fixed point may (as usual) be computed by beginning with $[\mathrm{pp}_1 \mapsto \bot, \ldots, \mathrm{pp}_m \mapsto \bot]$ (every program point is mapped to the least element of Abs), and repeatedly replacing the value currently assigned to pp_i by the value of the right side of pp_i's equation. By monotonicity of $\sqcap$, $f_{n\div 2}$ etc., these values can only grow or remain unchanged, so the iterations terminate provided the approximation lattice has no ascending chains of infinite height, as is the case here.

[Cousot, 1977a] describes ways to achieve termination even when infinite chains exist, by inserting so-called *widening* operators in the data-flow equations at each junction point of a loop. To explain the basic idea consider the problem of finding the fixed point of a continuous function f. The usual Kleene iteration sequence is $d_0 = \bot, \cdots, d_{n+1} = f(d_n), \cdots$ and is known to converge to the least fixed point of f but the sequence need not stabilize, i.e. it need not be the case that $d_{n+1} = d_n$ for some n. To remedy this one may introduce a *widening* operator $\bigtriangledown$ that dominates the least upper bound operation, i.e. $d' \sqcup d'' \sqsubseteq d' \bigtriangledown d''$, and such that the chain $d_0 = \bot, \cdots, d_{n+1} = d_n \bigtriangledown f(d_n)$ always stabilizes. This leads to overshooting the least fixed point but always gives a safe solution. By iterating down from the stabilization-value (perhaps by using the technique of *narrowing*) one may then be able to recover some of the information lost.

Constant propagation This is an example of a lattice which is infinite but has finite height (three). It is used for detecting variables that don't vary, and has $\mathrm{Abs} = \{\top, \bot, 0, 1, 2, \ldots\}$ where $\bot \sqsubseteq n \sqsubseteq \top$ for $n = 0, 1, 2, \ldots$.

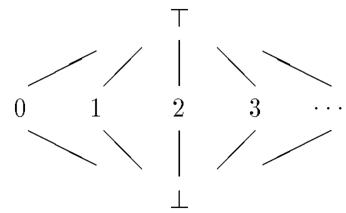

The corresponding abstraction function is:

$$\alpha(V) = \begin{cases} \bot & \text{if V} = \{\ \} \\ n & \text{if V} = \{n\} \\ \top & \text{otherwise} \end{cases}$$

There also exist lattices in which all ascending chains have finite height, even though the lattice as a whole has unbounded vertical extent. An example: let $\mathrm{Abs} = (N, \geq)$.

2.2.6 Safety: First Discussion

The analysis of the Collatz-sequence program is clearly "safe" in the following sense: if control reaches point C then the value of n will be even, and similarly for the other program points and abstract values. Correctness (or soundness) of the even-odd analysis for *all possible* programs and program points is also fairly easy to establish, given the close connection of the flow equations to those defining the accumulating semantics.

Reachable program points A similar but simpler *reachability* analysis (e.g. for dead code elimination) serves to illustrate a point concerning safety. It uses Abs $= \{\top, \perp\}$ with $\perp \sqsubseteq \top$ and abstraction function α defined as follows (where $a \in$ Abs and $S \subseteq$ Store):

$$\begin{array}{lcll} \alpha(S) & = & \perp \text{ if } S = \{\} & \text{else } \top \\ f_{n\div 2}(a) & = & \perp \text{ if } a = \perp & \text{else } \top \\ f_{3n+1}(a) & = & \perp \text{ if } a = \perp & \text{else } \top \end{array}$$

Intuitively, $\perp$ abstracts only the empty set of stores and so appropriately describes unreachable program points, while $\top$ describes reachable program points. Computing the fixpoint as above we get:

abs_A	abs_B	abs_C	abs_D	abs_E	abs_F	abs_G
$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$

This might be thought to imply that *all* program points including G are reachable, regardless of the initial value of n. On the other hand, reachability of G for input n implies termination, and it is a well-known open question whether the program does in fact terminate for all n.

A more careful analysis reveals that $\perp$ at program point p represents "p *cannot* be reached", while $\top$ represents "p *might* be reached" and so does not necessarily imply termination. The example shows that we must examine the questions of correctness and safety more carefully, which we now proceed to do.

2.3 Correctness and Safety

In this and remaining parts of section 2, we describe informally several different approaches to formulating safety and correctness, and discuss some advantages and disadvantages. A more detailed domain-based framework will be set up in section 3.

2.3.1 Desirable Properties of the Abstract Value Set Abs

In order to model the accumulating semantics equations, Abs could be a *complete lattice*: a set with a partial order $\sqsubseteq$, with least upper and greatest lower bounds $\sqcup$ and $\sqcap$ to model $\cup$ and $\cap$, and such that any collection of sets of stores has a least upper bound in Abs. Note: any lattice of finite height is complete. In the following we sometimes write $a \sqsupseteq a'$ in place of $a' \sqsubseteq a$.

2.3.2 Desirable Properties of the Abstraction Function

Intuitively "even" represents the set of all even numbers. This viewpoint is made explicit in [Cousot, 1977a] by relating complete lattices Conc and Abs to each other by a pair α, γ of *abstraction and concretization* functions with types

$$\begin{array}{rlcl} \alpha: & \text{Conc} & \rightarrow & \text{Abs} \\ \gamma: & \text{Abs} & \rightarrow & \text{Conc} \end{array}$$

In the even-odd example above the lattice of concrete values is Conc = $\wp$(Store), and the natural concretization function is

$$\begin{array}{rcl} \gamma(\bot) & = & \{\} \\ \gamma(even) & = & \{0, 2, 4, \ldots\} \\ \gamma(odd) & = & \{1, 3, 5, \ldots\} \\ \gamma(\top) & = & \{0, 1, 2, 3, \ldots\} = \mathbb{N} \end{array}$$

Cousot and Cousot impose natural conditions on α and γ (satisfied by the examples):

1. α and γ are monotonic
2. $\forall a \in \text{Abs},\ a = \alpha(\gamma(a))$
3. $\forall c \in \text{Conc},\ c \sqsubseteq_{Conc} \gamma(\alpha(c))$

For the accumulating semantics, larger abstract values represent larger sets of stores by condition 1. Condition 2 is natural, and condition 3 says that $S \subseteq \gamma(\alpha(S))$ for any $S \subseteq$ Store.

The conditions can be summed up as: (α, γ) form a Galois insertion of Abs into $\wp$(Store), a special case of an adjunction in the sense of category theory. It is easy to verify the following

Lemma 1 If conditions 1-3 hold, then

- $\forall c \in$ Conc, a $\in$ Abs: $c \sqsubseteq_{Conc} \gamma(a)$ if and only if $\alpha(c) \sqsubseteq_{Abs} a$, and

- α is continuous

□

Thus the abstract flow equations converge to a fixpoint. If α is semihomomorphic on union, intersection and base functions, then the abstract flow equations' fixpoint will be pointwise larger than or equal to the abstraction of the fixpoint of the accumulating semantics' equations.

Again, note that stores are unordered, so α and γ need only preserve the subset ordering. The more complex situation that arises when modelling nonflat domains is investigated in [Mycroft, 1983].

2.3.3 Safety: Second Discussion

Recalling the program of section 2.2, we can define the solution $(\mathrm{abs}_A,\ldots,\mathrm{abs}_G) \in \mathrm{Abs}^7$ to the abstract flow equations to be *safe* with respect to the accumulating semantics $(\mathrm{acc}_A,\ldots,\mathrm{acc}_G) \in \wp(\mathrm{Store})^7$ if the reachable sets of stores are represented by the abstract values:

$$\begin{array}{rcl} acc_A & \subseteq & \gamma(abs_A), \\ acc_B & \subseteq & \gamma(abs_B), \\ & \vdots & \\ acc_G & \subseteq & \gamma(abs_G) \end{array}$$

This is easy to verify for the even-odd abstraction given before.

Returning to the question raised after the "reachable program points" example, we see that safety at point G only requires that $\mathrm{acc}_G \subseteq \gamma(\mathrm{abs}_G)$, i.e. that every store that can reach G appears in $\gamma(\mathrm{abs}_G)$. This also holds if acc_G is empty, so $\gamma(\mathrm{abs}_G) = \top$ does *not* imply that G is reachable in any actual computation. For any program point X, $\mathrm{abs}_X = \bot$ implies $\mathrm{acc}_X = \{\}$, which signifies that control cannot reach X. Thus abstract value $\bot$ can be used to eliminate dead code.

Safe approximation of base functions Consider a base function op : $\mathbb{N} \rightarrow \mathbb{N}$, and extend it, by "pointwise lifting" to sets of numbers, yielding $\wp\mathrm{op} : \wp(\mathbb{N}) \rightarrow \wp(\mathbb{N})$ where

$$\wp\mathrm{op}(N) = \{\mathrm{op}(n) \mid n \in N\}$$

Suppose α, γ satisfy conditions 1-3. It is natural to define $\mathbf{op} : \mathrm{Abs} \rightarrow \mathrm{Abs}$ to be a *safe approximation* to op if the following holds for all $N \subseteq \mathbb{N}$:

$$\wp\mathrm{op}(N) \subseteq \gamma(\mathbf{op}(\alpha(N)))$$

or, diagrammatically:

$$\begin{tikzcd}
\wp(\mathbb{N}) \arrow[r, "\wp\mathrm{op}"] \arrow[dd, "\alpha"] & \wp(\mathbb{N}) \\
 & \rotatebox[origin=c]{90}{$\supseteq$} \\
\mathrm{Abs} \arrow[r, "\mathbf{op}"] & \mathrm{Abs} \arrow[u, "\gamma"']
\end{tikzcd}$$

By the conditions and lemma this is equivalent to

$$\alpha(\wp\mathrm{op}(N)) \sqsubseteq \mathbf{op}(\alpha(\mathbf{N}))$$

corresponding to diagram:

$$\begin{tikzcd}
\wp(\mathbb{N}) \arrow[r, "\wp\mathrm{op}"] \arrow[dd, "\alpha"] & \wp(\mathbb{N}) \arrow[d, "\alpha"] \\
 & \rotatebox[origin=c]{90}{$\sqsupseteq$} \\
\mathrm{Abs} \arrow[r, "\mathbf{op}"] & \mathrm{Abs}
\end{tikzcd}$$

Intuitively, for any subset $N \subseteq \mathbb{N}$, applying the induced abstract operation **op** to the abstraction of N represents at least all the values obtainable by applying op to members of N.

Induced approximations to base functions We now show how the best possible approximation **op** can be extracted from op (at least in principle, although perhaps not computably; a more detailed discussion appears in section 3.4). Recall that smaller elements of Abs abstract smaller sets of concrete values and so are less approximate, i.e. more precise descriptions.

Lemma 2 Given $\alpha : \wp(\mathbb{N}) \to \text{Abs}$ and $\gamma : \text{Abs} \to \wp(\mathbb{N})$ satisfying the three conditions above, define the operator *induced* by op to be $\mathbf{op} : \text{Abs} \to \text{Abs}$ where

$$\mathbf{op} = \alpha \circ \wp\text{op} \circ \gamma$$

Then $\mathbf{op}$ is the most precise function on Abs satisfying $\alpha(\wp\text{op}(N)) \sqsubseteq \mathbf{op}(\alpha(N))$ for all N.

Proof Suppose $f : \text{Abs} \to \text{Abs}$ with $\alpha(\wp\ \text{op}(N)) \sqsubseteq f(\alpha(N))$ for all N. Then for any a,

$$\mathbf{op}(a) = \alpha(\wp\text{op}(\gamma(a))) \sqsubseteq f(\alpha(\gamma(a))) = f(a)$$

□

The definition of $\mathbf{op}$ as a diagram:

$$\begin{tikzcd}
\wp(\mathbb{N}) \arrow[r, "\wp\text{op}"] & \wp(\mathbb{N}) \arrow[d, "\alpha"] \\
\text{Abs} \arrow[u, "\gamma"'] \arrow[r, "\mathbf{op}"] & \text{Abs}
\end{tikzcd}$$

For example, if $\text{op}(n) = n \div 2$ then $\mathbf{op}$ is $f_{n\div 2}$ as seen above, e.g.

$$\begin{array}{llll}
\mathbf{op}(\bot) & = \alpha(\{n \div 2 \mid n \in \gamma(\bot)\}) & = \alpha(\{\}) & = \bot \\
\mathbf{op}(\text{even}) & = \alpha(\{n \div 2 \mid n \in \gamma(\text{even})\}) & = \alpha(\{0, 1, 2, \ldots\}) & = \top
\end{array}$$

Unfortunately the definition of $\mathbf{op}$ does not necessarily give a terminating algorithm for computing it, even if op is computable. In practice the problem is solved by approximating from above, i.e. choosing $\mathbf{op}$ to give values in Abs that may be larger (less informative) than implied by the above equation. We will go deeper into this in Subsection 3.5.

A local condition for safe approximation of transitions A safety condition on one-step transitions can be formulated analogously. Define for any two control points p, q the function $\text{next}_{p,q}:\wp(\text{Store}) \rightarrow \wp(\text{Store})$:

$$\text{next}_{p,q}(S) = \{s' \mid (q, s') = \text{next}((p, s)) \text{ for some } s \in S\}$$

This is the earlier transition function, extended to include all transitions from p to q on a set of stores. Exactly as above we can define the abstract transition function *induced* by α and γ to be

$$\mathbf{next}_{p,q} = \alpha \circ \text{next}_{p,q} \circ \gamma$$

This is again the most precise function satisfying

$$\mathbf{next}_{p,q}(\alpha(S)) \sqsubseteq \alpha(\text{next}_{p,q}(S))$$

for all S.

2.3.4 An Example: the Rule of Signs

Consider the abstract values +, - and 0 with the natural concretization function

$$\begin{array}{rcl} \gamma(0) & = & \{0\} \\ \gamma(+) & = & \{1, 2, 3, ...\} \\ \gamma(-) & = & \{-1, -2, -3, \ldots\} \end{array}$$

This can be made into a complete lattice by adding greatest lower and least upper bounds in various ways. Assuming $\sqcap$, $\sqcup$ should model $\cap$, $\cup$ respectively, the following is obtained:

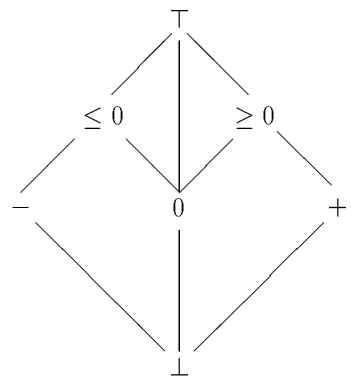

with

$$\begin{aligned}\gamma(\perp) &= \{\}\\ \gamma(\geq 0) &= \{0, 1, 2, 3, ...\}\\ \gamma(\leq 0) &= \{0, -1, -2, -3, ...\}\\ \gamma(\top) &= \{\ldots, -2, -1, 0, 1, 2, 3, \ldots\} = \mathbb{Z}\end{aligned}$$

and abstraction function

$$\alpha(S) = \begin{cases} \perp & \text{if } S = \{\} \text{ else}\\ + & \text{if } S \subseteq \{1,\ 2,\ 3,...\} \text{ else}\\ \geq 0 & \text{if } S \subseteq \{0,\ 1,\ 2,\ 3, \ldots\} \text{ else}\\ - & \text{if } S \subseteq \{\text{-1,- 2, -3,}\ldots\} \text{ else}\\ \leq 0 & \text{if } S \subseteq \{0\text{, -1, -2, -3,}\ldots\} \text{ else}\\ \top & \end{cases}$$

The induced approximation for operator $+ : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$ is:

$+'$	$\perp$	$-$	0	$+$	≥ 0	≤ 0	$\top$
$\perp$	$\perp$	$\perp$	$\perp$	$\perp$	$\perp$	$\perp$	$\perp$
$-$	$\perp$	$-$	$-$	$\top$	$\top$	$-$	$\top$
0	$\perp$	$-$	0	$+$	≥ 0	≤ 0	$\top$
$+$	$\perp$	$\top$	$+$	$+$	$+$	$\top$	$\top$
≥ 0	$\perp$	$\top$	≥ 0	$+$	≥ 0	$\top$	$\top$
≤ 0	$\perp$	$-$	≤ 0	$\top$	$\top$	≤ 0	$\top$
$\perp$	$\perp$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$

2.3.5 Composition of Safety Diagrams

Suppose we have two diagrams for safe approximation of two base functions op and op':

$$\begin{tikzcd} \wp(\mathbb{N}) \arrow[r, "\wp\mathrm{op}"] \arrow[d, "\alpha"] & \wp(\mathbb{N}) \arrow[d, "\alpha"] \\ \mathrm{Abs} \arrow[r, "\mathbf{op}"] & \mathrm{Abs} \end{tikzcd}$$

and

$$
\begin{tikzcd}
\wp(\mathbb{N}) \arrow[r, "\wp\mathrm{op'}"] \arrow[dd, "\alpha"] & \wp(\mathbb{N}) \arrow[d, "\alpha"] \\
 & \sqcap\!| \\
\mathrm{Abs} \arrow[r, "\mathbf{op}'"] & \mathrm{Abs}
\end{tikzcd}
$$

It is easy to see that $\mathbf{op}' \circ \mathbf{op}$ is a safe approximation to $\wp\ \mathrm{op}' \circ \wp\ \mathrm{op}$, so the two may be composed:

$$
\begin{tikzcd}
\wp(\mathbb{N}) \arrow[r, "\wp\mathrm{op'} \circ \ \wp\mathrm{op}"] \arrow[dd, "\alpha"] & \wp(\mathbb{N}) \arrow[d, "\alpha"] \\
 & \sqcap\!| \\
\mathrm{Abs} \arrow[r, "\mathbf{op}' \circ \mathbf{op}"] & \mathrm{Abs}
\end{tikzcd}
$$

On the other hand the diagrams for the *induced* approximations to base functions *cannot* be so composed, since the best approximation to $\wp\mathrm{op}' \circ \wp\mathrm{op}$ may be better than the composition of the best approximations to $\wp\mathrm{op}$ and $\wp\mathrm{op}'$. (This is precisely because α is a semihomomorphism, not a homomorphism.) For a concrete example, let op and op$'$ respectively describe the effects of the two assignments

$$n := 4 * n + 2;\ n := n \div 2$$

Then

$$\alpha(\wp\mathrm{op}' \circ \wp\mathrm{op}(\{0, 1, 2, ...\})) = \alpha(\{1, 3, 5, ...\}) = odd$$

whereas

$$\mathbf{op'} \circ \mathbf{op}\ (\alpha(\{0,1,2,...\}) = \mathbf{op'}(\mathrm{even}) = \top.$$

2.4 Scott Domains, Lattice Duality, and Meet versus Join

Relation to Scott-style domains The partial order $\sqsubseteq$ on Abs models the set inclusion order $\subseteq$ used for $\wp$(Store) in the accumulating semantics. In abstract interpretation, larger elements of Abs correspond to *more approximate* descriptions, so if $a \sqsubseteq a'$ then a' describes a *larger set* of concrete values. For example, "even" describes any set of even numbers, and $\top$ describes the set of all numbers.

In contrast, Scott domains as used in denotational semantics use an ordering by "information content", where a larger domain element describes *a single value that is more completely calculated.* During a computation $\perp$ means "not yet calculated", intuitively a slot to be filled later in with the final value. Appearance of $\perp$ in a program's final result signifies "was never filled in", and so represents nontermination (at least in languages with eager evaluation).

A value in a Scott domain represents perhaps incomplete knowledge about a *single* program value, for example a finite part of an infinite function f. The partial order $f \sqsubseteq f'$ signifies that f' is more completely defined than f, and that f' agrees with f where ever it is defined. $\top$, if used at all, indicates inconsistent values.

Clearly this order is not the same as the one used in abstract interpretation, and the difference is more than just one of duality.

Least or Greatest Fixpoints? Literature on data-flow analysis as used in compilers [Aho, Sethi and Ullman, 1986,Hecht, 1977,Kennedy, 1981] often uses abstract value lattices which are dual to the ones we consider, so larger elements represent more precise descriptions rather than more approximate. This is mainly a matter of taste; but has the consequence that *greatest* fixpoints are computed instead of least ones, and that the $\cup$ and $\cap$ of the accumulating semantics are modelled by $\sqcap$ and $\sqcup$, respectively. We prefer least fixpoints due to their similarity to those naturally used in defining the accumulating semantics.

Should $\sqcup$ or $\sqcap$ be Used on Converging Paths? We have argued that $\sqcup$ naturally models the effect of path convergence because it corresponds to $\cup$ in the accumulating semantics. On the other hand, there exist abstract interpretations that are *not* approximations to the accumulating semantics, and for some of these path convergence is properly modelled by $\sqcap$. To see this, consider the two dependence analyses mentioned in section 1.1. For analysis I, path convergence should be modelled by $\sqcup$ since a variable dependence is to be recorded if it occurs along *at least one* path. For analysis II it should be modelled by $\sqcap$ since a dependence is recorded only if it occurs along *all* paths. So the choice between $\sqcup$ and $\sqcap$ on converging paths is just another incarnation of the modality distinction encountered in section 1.

2.5 Abstract Values Viewed as Relations or Predicates

The accumulating semantics binds to each program point a set of stores. Suppose the program's variables are $V_1, \ldots, V_n$, so a store is an element of Store = $\{V_1, \ldots, V_n\} \to$ Value. In the examples above there was only one variable, so a set of stores was essentially a set of values, which simplified the discussion considerably. The question arises: how can we abstract a set of stores when $n > 1$?

2.5.1 Independent Attribute Analyses

Suppose value sets are abstracted by $\alpha_{val} : \wp(\text{Value}) \to A$. The *independent attribute method* models a set of stores S at program point p by mapping each variable V_i to an abstraction of the set of values it takes in all the stores of S. This abstract value is thus independent of all other variables, hence the term "independent attribute". For example, $\{[X \mapsto 1, Y \mapsto 2], [X \mapsto 3, Y \mapsto 1]\})$ would be modelled by $[X \mapsto odd, Y \mapsto \top]$.

Formally, we model

$$S \in \wp(\text{Store}) = \wp(\{V_1, \ldots, V_n\} \to \text{Value})$$

by a function

$$\text{abs}_p \in \text{Abs} = \{V_1, \ldots, V_n\} \to A$$

The store abstraction function $\alpha_{sto} : \wp(\text{Store}) \to \text{Abs}$ is defined by

$$\alpha_{sto}(S) = [V_i \mapsto \alpha_{val}(\{s(V_i) \mid s \in S\})]_{i=1,\ldots,n}$$

For example, consider an even-odd analysis of a program with variables X, Y, Z. The independent attribute method would abstract a set of two stores as follows:

$$\begin{array}{c}
\alpha_{sto}(\{[X \mapsto 1, Y \mapsto 2, Z \mapsto 1], [X \mapsto 2, Y \mapsto 2, Z \mapsto 1]\}) \;\; = \\
{[X \mapsto \alpha_{val}(\{1,2\}), Y \mapsto \alpha_{val}(\{2\}), Z \mapsto \alpha_{val}(\{1\})]} \;\; = \\
{[X \mapsto \top, Y \mapsto even, Z \mapsto odd]}
\end{array}$$

The independent attribute method abstracts each variable independently of all others, and so allows "cross over" effects. An example:

$$\alpha(\{[X \mapsto 1, Y \mapsto 1], [X \mapsto 2, Y \mapsto 2]\}) \;= [X \mapsto \top, Y \mapsto \top] \;= \\ \alpha(\{[X \mapsto 1, Y \mapsto 2], [X \mapsto 2, Y \mapsto 1]\})$$

This loses information about relationships between X's and Y's values, e.g. whether or not they always have the same parity.

2.5.2 Relational Analyses

Relations and predicates Abstract value abs_p is an abstraction of the set of stores acc_p, so the question arises as to how to represent it by a lattice element. An approach used in [Cousot, 1977a], [Cousot, 1977b] is to describe acc_p and its approximations abs_p by predicate calculus formulas. For instance the set of two stores $\{[X \mapsto 1, Y \mapsto 1], [X \mapsto 2, Y \mapsto 2]\}$ above could be approximately described by the formula:

$$(odd(X) \wedge odd(Y)) \vee (even(X) \wedge even(Y))$$

More generally, suppose Store = $\{\text{V}_1,...,\text{V}_n\} \rightarrow$ Value. Clearly Store is isomorphic to Value^n, the set of all n-tuples of values. Thus any set of stores i.e. any element of $\wp$(Store)) can be interpreted as a set of n-tuples. For example, store set $\{[X \mapsto 1, Y \mapsto 1], [X \mapsto 2, Y \mapsto 2]\}$ corresponds to $\{(1,1), (2,2)\}$. Thus a store set is essentially a set of n-tuples or, in other words, an n-ary *predicate* or *relation*.

For program point p, the accumulating semantics defines relation $\text{acc}_p(\text{v}_1,\ldots,\text{v}_n)$ to be true just in the case that $(\text{v}_1,...,\text{v}_n)$ is a tuple of values which can occur at p in one or more computations on the given initial input data. This is the weakest possible relation among variables that always holds at point p.

Relational Analyses These use more sophisticated methods to approximate $\wp$(Store), which can give more precise information. Examples of practically motivated program analysis problems that require relational information include aliasing analysis in Pascal, the recognition of possible substructure sharing in Lisp or Prolog, and interference analysis.

For an example not naturally represented by independent attributes, suppose we wish to find out which of a program's variables always assume the same value at a given program point p. A suitable abstraction of a set of stores is a *partition* π_p that divides the program's variables into equivalence classes, so any one class of π_p contains all variables that have the same value at p. The effect of an assignment such as "p: X:=Y; **goto** q" is that π_q is obtained from π_p by removing X from its previous equivalence class and adding it to Y's class.

Intensional versus extensional descriptions Above we represented store set

$$\{[X \mapsto 1, Y \mapsto 1],\ [X \mapsto 2, Y \mapsto 2]\}$$

by the binary relation {(1,1), (2,2)}, and approximated it by the superset {(x,y) | x and y are both even or both odd}, denoted by the predicate calculus formula

$$(odd(X) \wedge odd(Y)) \vee (even(X) \wedge even(Y))$$

The view of "predicate as a set of tuples" and "predicate as a formula" is exactly the classical distinction between the *extensional* and the *intension*al views of a predicate.

Descriptions by predicate calculus formulas must of necessity be only approximate, since there are only countably many formulas but uncountably many sets of stores (if we assume an infinite variable value set). In terms of predicate calculus formulas, for each program point p the appropriate formulation of a safe approximation is that acc_p *logically implies* abs_p. In terms of sets of n-tuples: each acc_p is a subset of the set of all tuples satisfying abs_p.

2.5.3 Abstract Interpretation and Predicate Transformers

The new view of the accumulating semantics is: given a program and a predicate describing its input data, the accumulating semantics maps every program point to the smallest *relation among variables that holds whenever control reaches that point.*

From this viewpoint, the function $\mathrm{next}_{p,q} : \wp(\mathrm{Store}) \rightarrow \wp(\mathrm{Store})$ is clearly the *forward predicate transformer* [Dijkstra, 1976] associated with transitions from p to q. Further, acc_p is clearly *the strongest postcondition* holding at program point p over all computations on input data satisfying the program's input precondition.

Program verification amounts to proving that each acc_p logically implies a user-supplied program assertion for point p. Note however that this abstract interpretation framework says nothing at all about program termination. This approach is developed further in [Cousot, 1977b] and their subsequent works.

Backwards analyses All this can easily be dualised: the *backward predicate transformer* $\mathrm{next}_{p,q}^{-1} : \wp(\mathrm{Store}) \rightarrow \wp(\mathrm{Store})$ is just the inverse of $\mathrm{next}_{p,q}$, and given a *program postcondition* one may find the *weakest precondition* on program input sufficient to imply the postcondition at termination. For the simple imperative language, a backward accumulating semantics is straightforward to construct. For the example program

A: **while** $n \neq 1$ **do**
  B: **if** n even
    then (C: $n := n \div 2$; D:)
    else (E: $n := 3 * n + 1$; F:)
   fi
  od
G:

the appropriate equations are:

$$\begin{array}{lcl}
\mathrm{acc}_A & = & (\{1\} \cap \mathrm{acc}_G) \cup (\{0,2,3,4,\ldots\} \cap \mathrm{acc}_B) \\
\mathrm{acc}_B & = & (\mathrm{acc}_C \cap Evens) \cup (\mathrm{acc}_E \cap \mathrm{Odds}) \\
\mathrm{acc}_C & = & \{n \mid n \div\ 2 \in\ \mathrm{acc}_D\} \\
\mathrm{acc}_D & = & (\{1\} \cap \mathrm{acc}_G) \cup (\{0,2,3,4,\ldots\} \cap \mathrm{acc}_B) \\
\mathrm{acc}_E & = & \{n \mid 3n+1 \in \mathrm{acc}_F\} \\
\mathrm{acc}_F & = & (\{1\} \cap \mathrm{acc}_G) \cup (\{0,2,3,4,\ldots\} \cap \mathrm{acc}_B) \\
\mathrm{acc}_G & = & S_{final}
\end{array}$$

where acc_p is the set of all stores at point p that cause control to reach point G with a final store in S_{final}.

Such a backward accumulating semantics can, for example, provide a basis for an analysis that detects the set of states that may lead to an error. More generally backward analyses (although not the one shown here) may provide a basis for "future sensitive" analysis such as *live variables*, where variable X is "semantically live" at point p if there is a computation sequence starting at p and later referencing X's value. This is approximated by: X is "syntactically live" if there is a program path from p to a use of X's value. Section 3 contains an example of live variable analysis for functional programs.

Many analysis problems can be solved by either a forwards or a backwards analysis. There can, however, be significant differences in efficiency.

Backwards analysis of functional programs The backwards accumulating semantics is straightforward for imperative programs, partly because of its close connections with the well studied weakest preconditions [Dijkstra, 1976], and because the state transition function is monadic. It is semantically less well understood, however, for functional programs, where recent works include [Hughes, 1987], [Dybjer, 1987], [Wadler, 1987], and [Nielson, 1989]. Natural connections between backwards analyses and continuation semantics are seen in [Nielson, 1982] and [Hughes, 1987].

2.6 Important Points from Earlier Sections

In the above we have employed a rather trivial programming language so as to motivate and illustrate one way to approximate real computations by computations over a domain of abstract values: Cousot's accumulating semantics. Before proceeding to abstract interpretation of more interesting languages we recapitulate what has been learned so far.

- Computations in the abstract world are at best *semihomomorphic* models of corresponding computations in the world of actual values.
- *Safety* of an abstract interpretation is analogous to *reliability* of a numerical analyst's results: the obtained results must always lie within specified confidence intervals (usually "one-sided intervals" in the case of program analysis).
- To obtain safe results for specific applications it is essential to understand the interpretation of the abstract values and their relation to actual computational values. One example is *modality*, e.g. "all computations" versus "some computations".
- Abstract values often do not contain enough information to determine the outcome of tests, so abstract interpretation must achieve the effect of simulating a *set* of real computations.
- Computations on *complete lattices* of abstract values appropriately model computations on real values.
- The partial order on these lattices expresses the degree of precision in an approximate description, and is *quite different* from the traditional Scott-style ordering based on filling in incomplete information.
- Termination can be achieved by choosing lattices without infinite ascending chains.
- *Best* approximations to real computations exist in principle, but may be uncomputable.
- There are close connections between the "accumulating semantics" and the predicates and predicate transformers (both forwards and backwards) used in program verification.

2.7 Towards Generalizing the Cousot Framework

Abstract interpretation is a semantics-based approach to program analysis, but so far we have only dealt with a single, rather trivial language.