


Not all Data Flow Problems are Distributive

Constant propagation is *not* distributive.

Consider the following (with variables (x,y,z)):



Now $f(t) = t'$ where

$t'(y) = t(y)$, $t'(z) = t(z)$,

$t'(x) =$ if $t(y) = \perp$ or $t(z) = \perp$
 then \perp
 elseif $t(y) = T$ or $t(z) = T$
 then T
 else $t(y) + t(z)$

Now $f(t_1 \wedge t_2) = f(T, \perp, \perp) = (\perp, \perp, \perp)$

$f(t_1) = (4, 1, 3)$

$f(t_2) = (4, 2, 2)$

$f(t_1) \wedge f(t_2) = (4, \perp, \perp) \geq (\perp, \perp, \perp)$

Why does it Matter if a Data Flow Problem isn't Distributive?

Consider actual program execution paths from b_0 to (say) b_k .

One path might be $b_0, b_{i_1}, b_{i_2}, \dots, b_{i_n}$ where $b_{i_n} = b_k$.

At b_k the Data Flow information we want is

$f_{i_n}(\dots f_{i_2}(f_{i_1}(f_0(T)))) \equiv f(b_0, b_{i_1}, \dots, b_{i_n})$

On a different path to b_k , say $b_0, b_{j_1}, b_{j_2}, \dots, b_{j_m}$, where $b_{j_m} = b_k$

the Data Flow result we get is

$f_{j_m}(\dots f_{j_2}(f_{j_1}(f_0(T)))) \equiv f(b_0, b_{j_1}, \dots, b_{j_m})$.

Since we can't know at compile time which path will be taken, we must *combine* all possible paths:

$$\bigwedge_{p \in \text{all paths to } b_k} f(p)$$

This is the *meet over all paths* (MOP) solution. It is the *best possible* static solution. (Why?)

As we shall see, the meet over all paths solution can be computed efficiently, using standard Data Flow techniques, if the problem is Distributive.

Other, non-distributive problems (like Constant Propagation) can't be solved as precisely.

Explicitly computing and meeting all paths is prohibitively expensive.

Conditional Constant Propagation

We can extend our Constant Propagation Analysis to determine that some paths in a CFG aren't executable. This is *Conditional Constant Propagation*.

Consider

```
i = 1;
if (i > 0)
    j = 1;
else j = 2;
```

Conditional Constant Propagation can determine that the else part of the if is unreachable, and hence *j* must be 1.

The idea behind Conditional Constant Propagation is simple. Initially, we mark all edges out of conditionals as "not reachable."

Starting at b_0 , we propagate constant information *only* along edges considered reachable.

When a boolean expression $b(v_1, v_2, \dots)$ controls a conditional branch, we evaluate $b(v_1, v_2, \dots)$ using the $t(v)$ mapping that identifies the "constant status" of variables.

If $t(v_i) = \perp$ for any v_i , we consider all out edges unreachable (for now).

Otherwise, we evaluate $b(v_1, v_2, \dots)$ using $t(v)$, getting true, false or \perp .

Note that the short-circuit properties of boolean operators may yield true or false even if $t(v_i) = \perp$ for some v_i .


If $b(v_1, v_2, \dots)$ is true or false, we mark only one out edge as reachable.

Otherwise, if $b(v_1, v_2, \dots)$ evaluates to \perp , we mark all out edges as reachable.

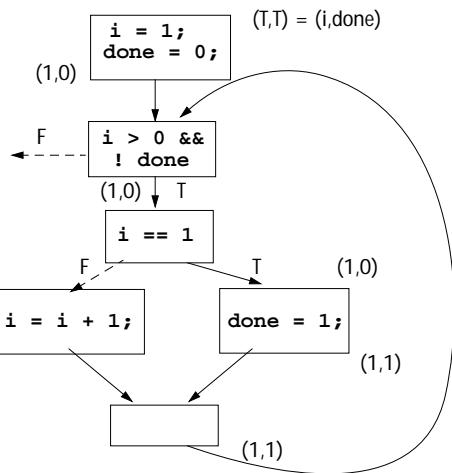
We propagate constant information only along reachable edges.

Example


```
i = 1;
done = 0;
while ( i > 0 && ! done ) {
    if ( i == 1 )
        done = 1;
    else i = i + 1;
}
```



Pass 1:



Pass 2:



Reading Assignment

- Read pages 63-end of "Automatic Program Optimization," by Ron Cytron. (Linked from the class Web page.)

Iterative Solution of Data Flow Problems


This algorithm will use DFO numbering to determine the order in which blocks are visited for evaluation. We iterate over the nodes until convergence.

```

EvalDF{
  For (all n ∈ CFG) {
    soln(n) = T
    ReEval(n) = true }
  Repeat
    LoopAgain = false
    For (all n ∈ CFG in DFO order){
      If (ReEval(n)) {
        ReEval(n) = false
        OldSoln = soln(n)
        In =  $\bigwedge_{p \in \text{Pred}(n)} \text{soln}(p)$ 
        soln(n) =  $f_n(\text{In})$ 
        If (soln(n) ≠ OldSoln) {
          For (all s ∈ Succ(n)) {
            ReEval(s) = true
            LoopAgain = LoopAgain OR
              IsBackEdge(n,s)
          }
        }
      }
    }
  Until (! LoopAgain)
}

```

Example: Reaching Definitions



We'll do this as a set-valued problem (though it really is just three bit-valued analyses, since each analysis is independent).

L is the power set of Basic Blocks

\cup is set union

T is ϕ ; \perp is the set of all blocks

$a \leq b \equiv b \subseteq a$

$f_3(\text{in}) = \{3\}$

$f_6(\text{in}) = \{6\}$

$f_7(\text{in}) = \{7\}$

For all other blocks, $f_b(\text{in}) = \text{in}$

We'll track soln and ReEval across multiple passes

	0	1	2	3	4	5	6	7	8	Loop-Again
Initial	ϕ	ϕ	ϕ	ϕ	ϕ	ϕ	ϕ	ϕ	ϕ	true
	true	true	true	true	true	true	true	true	true	
Pass 1	ϕ	ϕ	ϕ	{3}	{3}	{3}	{6}	{7}	{7}	true
	false	true	false	false	true	false	false	false	false	
Pass 2	ϕ	{3}	{3}	{3}	{3,7}	{3,7}	{6}	{7}	{7}	true
	false	true	false	false	false	false	false	false	false	
Pass 3	ϕ	{3,7}	{3,7}	{3}	{3,7}	{3,7}	{6}	{7}	{7}	false
	false	false	false	false	false	false	false	false	false	

Properties of Iterative Data Flow Analysis

- If the height of the lattice (the maximum distance from T to \perp) is finite, then termination is *guaranteed*.

Why?

Recall that transfer functions are assumed monotone ($a \leq b \Rightarrow f(a) \leq f(b)$).

Also, \wedge has the property that $a \wedge b \leq a$ and $a \wedge b \leq b$.

At each iteration, some solution value must change, else we halt. If something changes it must “move down” the lattice (we start at T). If the lattice has finite height, each block’s value can change only a bounded number of times. Hence termination is guaranteed.

- If the iterative data flow algorithm terminates, a valid solution *must* have been computed. (This is because data flow values flow forward, and any change along a backedge forces another iteration.)

How Many Iterations are Needed?

Can we bound the number of iterations needed to compute a data flow solution?

In our example, 3 passes were needed, but why?

In an “ideal” CFG, with no loops or backedges, only 1 pass is needed.

With backedges, it can take several passes for a value computed in one block to reach a block that depends upon the value.

Let p be the maximum number of backedges in any acyclic path in the CFG.

Then $(p+1)$ passes suffice to propagate a data flow value to any other block that uses it.

Recall that any block’s value can change only a bounded number of times. In fact, the height of the lattice (maximum distance from top to bottom) is that bound.

Thus the maximum number of passes in our iterative data flow evaluator = $(p+1) * \text{Height of Lattice}$

In our example, $p = 2$ and lattice height really was 1 (we did 3 independent bit valued problems).

So passes needed = $(2+1)*1 = 3$.

Rapid Data Flow Frameworks

We still have the concern that it may take many passes to traverse a solution lattice that has a significant height.

Many data flow problems are *rapid*. For rapid data flow problems, extra passes to feed back values along cyclic paths aren't needed.

For a data flow problem to be rapid we require that:


$$(\forall a \in A)(\forall f \in F) \quad a \wedge f(T) \leq f(a)$$

This is an odd requirement that states that using $f(T)$ as a very crude approximation to a value computed by F is OK when joined using the \wedge operator. In effect the term "a" rather than $f(T)$ is dominant).

(Recall that $a \wedge f(a) \leq f(a)$ always holds.)

How does the Rapid Data Flow Property Help?

Consider a direct feedback loop (the idea holds for indirect loops too):



a is an input from outside the loop. Our concern is how often we'll need to reevaluate f , as new values are computed and fed back into f . Initially, we'll use T to model the value on the backedge.

Iteration 1: Input = $a \wedge T = a$
Output = $f(a)$

Iteration 2: Input = $a \wedge f(a)$
Output = $f(a \wedge f(a))$

Iteration 3: Input = $a \wedge f(a \wedge f(a))$

Now we'll exploit the rapid data flow property: $b \wedge f(T) \leq f(b)$

Let $b \equiv a \wedge f(a)$

Then $a \wedge f(a) \wedge f(T) \leq f(a \wedge f(a))$ (*)

Note that $x \leq y \Rightarrow a \wedge x \leq a \wedge y$ (**)

To prove this, recall that

$$(1) \quad p \wedge q = p \Rightarrow p \leq q$$

$$(2) \quad x \leq y \Rightarrow x \wedge y = x$$

Thus $(a \wedge x) \wedge (a \wedge y) = a \wedge (x \wedge y) = (a \wedge x)$
(by 2) $\Rightarrow (a \wedge x) \leq (a \wedge y)$ (by 1).

From (*) and (**) we get
 $a \wedge a \wedge f(a) \wedge f(T) \leq f(a \wedge f(a)) \wedge a$ (***)

Now $a \leq T \Rightarrow f(a) \leq f(T) \Rightarrow$

$f(a) \wedge f(T) = f(a)$.

Using this on (***) we get

$a \wedge f(a) \leq f(a \wedge f(a)) \wedge a$

That is, $\text{Input}_2 \leq \text{Input}_3$

Note too that

$a \wedge f(a) \leq a \Rightarrow f(a \wedge f(a)) \leq f(a) \Rightarrow$

$a \wedge f(a \wedge f(a)) \leq a \wedge f(a)$

That is, $\text{Input}_3 \leq \text{Input}_2$

Thus we conclude $\text{Input}_2 = \text{Input}_3$,
 which means we can stop after two
 passes *independent* of lattice height!
 (One initial visit plus one reevaluation
 via the backedge.)

Many Important Data Flow Problems are Rapid

Consider reaching definitions, done as sets. We may have many definitions to the same variable, so the height of the lattice may be large.

L is the power set of Basic Blocks

\wedge is set union

T is ϕ ; \perp is the set of all blocks

$a \leq b \equiv a \supseteq b$

$f_b(\text{in}) = (\text{In} - \text{Kill}_b) \cup \text{Gen}_b$

where Gen_b is the last definition to a variable in b ,

Kill_b is all defs to a variable except the last one in b ,

Kill_b is empty if there is no def to a variable in b .

The Rapid Data Flow Property is

$a \wedge f(T) \leq f(a)$

In terms of Reaching Definitions this is

$a \cup f(\phi) \supseteq f(a) \equiv$


$a \cup (\phi - \text{Kill}) \cup \text{Gen} \supseteq (a - \text{Kill}) \cup \text{Gen}$

Simplifying,

$a \cup \text{Gen} \supseteq (a - \text{Kill}) \cup \text{Gen}$

which always holds.

Recall




Here it took two passes to transmit the def in b_7 to b_1 , so we expect 3 passes to evaluate *independent* of the lattice height.

Constant Propagation isn't Rapid

We require that

$$a \wedge f(T) \leq f(a)$$

Consider



Look at the transfer function for the second (bottom) block.

$$f(t) = t' \text{ where}$$

$$t'(v) = \text{case}(v) \{$$

$$k: t(j);$$

$$j: t(i);$$

$$i: 2; \}$$

$$\text{Let } a = (\perp, 1, 1).$$

$$f(T) = (2, T, T)$$

$$a \wedge f(T) = (\perp, 1, 1) \wedge (2, T, T) = (\perp, 1, 1)$$

$$f(a) = f(\perp, 1, 1) = (2, \perp, 1).$$

Now $(\perp, 1, 1)$ is not $\leq (2, \perp, 1)$
so this problem isn't rapid.

Let's follow the iterations:

$$\text{Pass 1: In} = (1, 1, 1) \wedge (T, T, T) = (1, 1, 1)$$

$$\text{Out} = (2, 1, 1)$$

$$\text{Pass 2: In} = (1, 1, 1) \wedge (2, 1, 1) = (\perp, 1, 1)$$

$$\text{Out} = (2, \perp, 1)$$

$$\text{Pass 3: In} = (1, 1, 1) \wedge (2, \perp, 1) = (\perp, \perp, 1)$$

$$\text{Out} = (2, \perp, \perp)$$

This took 3 passes. In general, if we had N variables, we could require N passes, with each pass resolving the constant status of one variable.

How Good Is Iterative Data Flow Analysis?

A single execution of a program will follow some path

$$b_0, b_{i_1}, b_{i_2}, \dots, b_{i_n}$$

The Data Flow solution along this path is

$$f_{i_n}(\dots f_{i_2}(f_{i_1}(f_0(T)))) \equiv f(b_0, b_{i_1}, \dots, b_{i_n})$$

The best possible static data flow solution at some block b is computed over all possible paths from b_0 to b.

Let P_b = The set of all paths from b_0 to b.

$$\text{MOP}(b) = \bigwedge_{p \in P_b} f(p)$$

Any particular path p_i from b_0 to b is included in P_b .

Thus $MOP(b) \wedge f(p_i) = MOP(b) \leq f(p_i)$.

This means $MOP(b)$ is *always* a safe approximation to the "true" solution $f(p_i)$.

If we have the distributive property for transfer functions,

$$f(a \wedge b) = f(a) \wedge f(b)$$

then our iterative algorithm *always* computes the MOP solution, the best static solution possible.

To prove this, note that for trivial path of length 1, containing only the start block, b_0 , the algorithm computes $f_0(T)$ which is $MOP(b_0)$ (trivially).

Now assume that the iterative algorithm for paths of length n or less to block c *does* compute $MOP(c)$.

We'll show that for paths to block b of length $n+1$, $MOP(b)$ is computed.

Let P be the set of all paths to b of length $n+1$ or less.

The paths in P end with b .

$$MOP(b) = f_b(f(P_1)) \wedge f_b(f(P_2)) \wedge \dots$$

where P_1, P_2, \dots are the prefixes (of length n or less) of paths in P with b removed.

Using the distributive property,

$$f_b(f(P_1)) \wedge f_b(f(P_2)) \wedge \dots = f_b(f(P_1) \wedge f(P_2) \wedge \dots).$$

But note that $f(P_1) \wedge f(P_2) \wedge \dots$ is just the input to f_b in our iterative algorithm, which then applies f_b .

Thus $MOP(b)$ for paths of length $n+1$ is computed.

For data flow problems that aren't distributive (like constant propagation), the iterative solution is \leq the MOP solution.

This means that the solution is a safe approximation, but perhaps not as "sharp" as we might wish.