

Dominators in Graphs

Ming-Ho Yee

May 8, 2020

Definitions

d dominates n

$d \text{ dom } n, \quad d \in \text{Dom}(n)$

- If every path from s_0 (start node) to n contains d .
- Every node dominates itself

d strictly dominates n

- If d dominates n and $d \neq n$

d immediately dominates n

$d = \text{idom}(n)$

- If d strictly dominates n and every other dominator of n dominates d

Dominator tree

- Every child is immediately dominated by its parent
- Root node is the start node
- Ancestor a of node $n \rightarrow a \text{ dom } n$




FIGURE 18.3. (a) A flow graph; (b) its dominator tree.

Applications

- Static Single Assignment
 - Compute the *dominance frontier* first
 - Use the frontier to insert phi nodes
- Finding loops
- Hoisting instructions

- And many other optimizations...

Anecdotal Performance in Ř

	Before		After	
	<i>Total</i>	<i>Average</i>	<i>Total</i>	<i>Average</i>
Benchmark (5 iterations)	89,205 ms	17,841 ms	82,425 ms	16,485 ms
Constructing the dominator tree	4,988 ms	80 μ s	90 ms	1.5 μ s
$a \text{ dom? } b$	42 ms	46.1 ns	169 ms	182 ns
$a \text{ idom? } b$	186 ms	24 ns	153 ms	19 ns

History

- 1959: Definition of *dominance* [Prosser 1959]
 - 1969: First sketch of algorithm [Lowry and Medlock, 1969]
 - Complexity is at least quadratic
 - 1970: Data-flow equations [Allen, 1970]
 - 1972: Iterative data-flow algorithm [Allen and Cocke, 1972]
- ... more quadratic algorithms ...*
- 1979: Almost linear complexity [Lengauer and Tarjan, 1979]
- ... more almost linear algorithms ...*

Data-flow Equations

Let $Dom(n)$ be the set of nodes that dominate n . Then:

$$Dom(s_0) = \{s_0\}$$

$$Dom(n) = \{n\} \cup \left(\bigcap_{p \in \text{preds}(n)} Dom(p) \right)$$

Iterate until you reach a fixed point.

$O(n^2)$ time complexity, and very slow in practice

Another Algorithm [Aho and Ullman, 1972]

For each node $v \neq s_0$:

- Remove v from the graph
- Consider the set of nodes S that are now unreachable
- Then $\text{Dom}(v) = S$

$O(n^2)$ time complexity

Iterative Algorithm, revisited

- Set intersection is the bottleneck
- Idea: use a consistent ordering for sets
 - Perform intersection by walking through both sets in order, with pairwise comparisons
 - Order encodes a path through the dominator tree
- Very simple implementation [Cooper, Harvey, and Kennedy, 2006]
- But slower in practice than the Lengauer-Tarjan algorithm [Georgiadis, Tarjan, and Werneck, 2006]

Lengauer-Tarjan Algorithm

- Simple version: $O(m \log n)$
- Sophisticated version: $O(m \alpha(m, n))$
 - Where $\alpha(m, n)$ is the inverse of the Ackerman function

The simple Lengauer-Tarjan algorithm is faster in practice, and less sensitive to pathological graphs.

This explanation is adapted from Appel and Palsberg [2004].

Depth-First Spanning Tree

- Use DFS to compute a spanning tree
 - Assign a *dfnum* to each node

a is an *ancestor* of *b*

- If $a = b$ or there is a path from a to b in the spanning tree
- *i.e.* $\text{dfnum}(a) \leq \text{dfnum}(b)$

a is a *proper ancestor* of *b*

- If a is an ancestor of b and $a \neq b$
- *i.e.* $\text{dfnum}(a) < \text{dfnum}(b)$

Note: Can test ancestor relation by comparing *dfnums*





FIGURE 19.8. A control-flow graph and trees derived from it. The numeric labels in part (b) are the *dfnums* of the nodes.

Dominators and *dfnums*

- If $\text{idom}(n) = d$, then d must be an ancestor of n
 - *i.e.*, $\text{dfnum}(d) < \text{dfnum}(n)$
- Therefore: ancestors of n are idom candidates!
- If an ancestor x does not dominate n , there must be some “detour” starting above x .
 - Nodes on the detour are not ancestors of n
 - *i.e.* their *dfnums* must be greater than n 's




Semidominators

s semidominates n

$$s = \text{semi}(n)$$

- If s is the highest ancestor with a path to n , using non-ancestor nodes
 - Highest ancestor \rightarrow smallest $dfnum$
 - $dfnum(s) < dfnum(n)$
 - Path $p = s, u_1, \dots, u_k, n$ using non-ancestor nodes
 - $dfnum(u_i) > dfnum(n)$




$\text{semi}(n)$ is a candidate for $\text{idom}(n)$

- Often, $\text{semi}(n) = \text{idom}(n)$
- An exception: $\text{semi}(n)$ itself is bypassed

Semidominator Theorem

Consider all predecessors v of n in the CFG. Then:

- If v is a proper ancestor of n $\text{dfnum}(v) < \text{dfnum}(n)$
 - v is a candidate for $\text{semi}(n)$
- If v is a non-ancestor of n $\text{dfnum}(v) > \text{dfnum}(n)$
 - For each u that is an ancestor of v , (and not an ancestor of n)
 $\text{semi}(u)$ is a candidate for $\text{semi}(n)$



$\text{semi}(n)$ is the candidate with the lowest dfnum

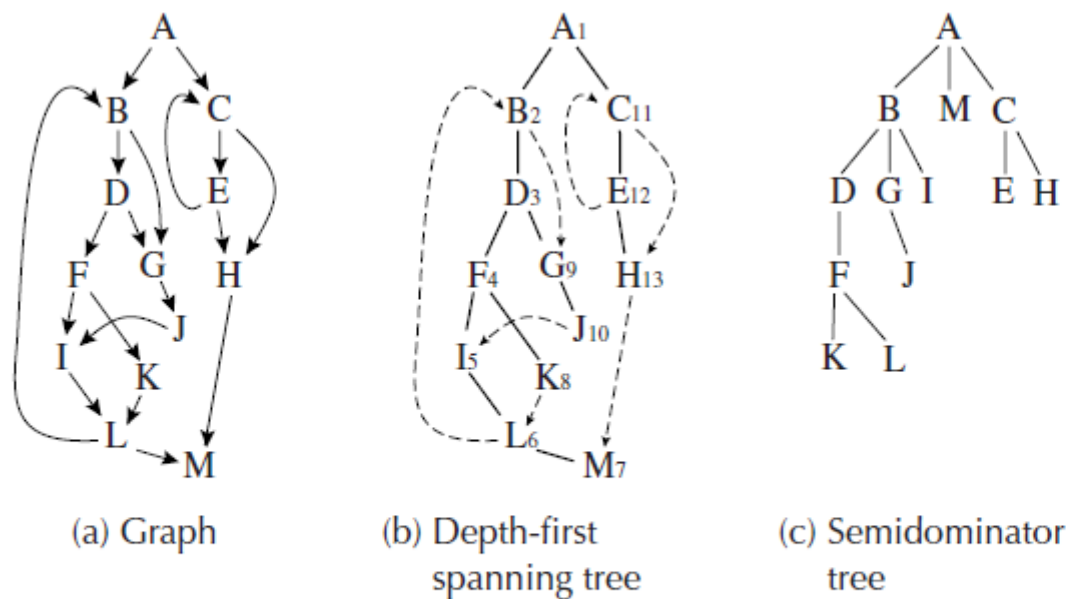


FIGURE 19.8. A control-flow graph and trees derived from it. The numeric labels in part (b) are the *dfnums* of the nodes.


Dominator Theorem

Consider the spanning tree path from $s = \text{semi}(n)$ to n .

Let y be the node with smallest numbered semidominator, *i.e.* minimum $\text{dfnum}(\text{semi}(y))$.

Then:

$$\text{idom}(n) = \begin{cases} \text{semi}(n) & \text{if } \text{semi}(y) = \text{semi}(n) \\ \text{idom}(y) & \text{if } \text{semi}(y) \neq \text{semi}(n) \end{cases}$$







FIGURE 19.8. A control-flow graph and trees derived from it. The numeric labels in part (b) are the *dfnums* of the nodes.

Lengauer-Tarjan Algorithm

1. Perform DFS to number nodes and create the depth-first spanning tree
2. For each node n (in decreasing $dfnum$ order):
 - Use the Semidominator Theorem to compute $semi(n)$
 - Insert n into the spanning forest
3. Implicitly define the idom by applying the first clause of the Dominator Theorem
4. For each node n (in increasing $dfnum$ order):
 - Explicitly define the idom by applying the second clause of the Dominator Theorem

Spanning Forest

- Build a spanning forest as the CFG is traversed
 - When n is processed, only non-ancestors of n are in the forest
- `link(p, n)`
 - Add the edge (n, p) to the spanning forest
- `ancestorWithLowestSemi(v)`
 - Search upwards in the forest, starting from v
 - Find the ancestor of v whose semidominator has the lowest *dfnum*






FIGURE 19.11. Path compression. (a) *Ancestor* links in a spanning tree; *AncestorWithLowestSemi*(v) traverses three links. (b) New nodes a_2 , a_3 are linked into the tree. Now *AncestorWithLowestSemi*(w) would traverse 6 links. (c) *AncestorWithLowestSemi*(v) with path compression redirects ancestor links, but *best*[v] remembers the best intervening node on the compressed path between v and a_1 . (d) Now, after a_2 and a_3 are linked, *AncestorWithLowestSemi*(w) traverses only 4 links.

Dominators() =

$N \leftarrow 0; \forall n. \text{bucket}[n] \leftarrow \{\}$

$\forall n. \text{dfnum}[n] \leftarrow 0, \text{semi}[n] \leftarrow \text{ancestor}[n] \leftarrow \text{idom}[n] \leftarrow \text{samedom}[n] \leftarrow \text{none}$

DFS(*none*, *r*)

for $i \leftarrow N - 1$ **downto** 1

$n \leftarrow \text{vertex}[i]; p \leftarrow \text{parent}[n]; s \leftarrow p$

for each predecessor *v* of *n*

if $\text{dfnum}[v] \leq \text{dfnum}[n]$

$s' \leftarrow v$

else $s' \leftarrow \text{semi}[\text{AncestorWithLowestSemi}(v)]$

if $\text{dfnum}[s'] < \text{dfnum}[s]$

$s \leftarrow s'$

$\text{semi}[n] \leftarrow s$

$\text{bucket}[s] \leftarrow \text{bucket}[s] \cup \{n\}$

Link(*p*, *n*)

for each *v* in $\text{bucket}[p]$

$y \leftarrow \text{AncestorWithLowestSemi}(v)$

if $\text{semi}[y] = \text{semi}[v]$

$\text{idom}[v] \leftarrow p$

else $\text{samedom}[v] \leftarrow y$

$\text{bucket}[p] \leftarrow \{\}$

for $i \leftarrow 1$ **to** $N - 1$

$n \leftarrow \text{vertex}[i]$

if $\text{samedom}[n] \neq \text{none}$

$\text{idom}[n] \leftarrow \text{idom}[\text{samedom}[n]]$

DFS(*node p*, *node n*) =

if $\text{dfnum}[n] = 0$

$\text{dfnum}[n] \leftarrow N; \text{vertex}[N] \leftarrow n; \text{parent}[n] \leftarrow p$

$N \leftarrow N + 1$

for each successor *w* of *n*

DFS(*n*, *w*)

AncestorWithLowestSemi(*node v*) =

$a \leftarrow \text{ancestor}[v]$

if $\text{ancestor}[a] \neq \text{none}$

$b \leftarrow \text{AncestorWithLowestSemi}(a)$

$\text{ancestor}[v] \leftarrow \text{ancestor}[a]$

if $\text{dfnum}[\text{semi}[b]] <$

$\text{dfnum}[\text{semi}[\text{best}[v]]]$

$\text{best}[v] \leftarrow b$

return $\text{best}[v]$

Link(*node p*, *node n*) =

$\text{ancestor}[n] \leftarrow p; \text{best}[n] \leftarrow n$

Implementation in Ě

Straightforward translation from pseudocode to C++

- About 100 LOC, without comments

<https://github.com/reactorlabs/rir/blob/b265f9e/rir/src/compiler/util/cfg.cpp#L52>