

SSA (2 of 2)

15-411/15-611 Compiler Design

Ben L. Titzer and Seth Goldstein

January 30, 2025


Today

- Thinking in SSA: Local reasoning
- Maintaining SSA properties during optimization
- Deconstructing SSA
 - Two issues: critical edges and ordering moves
- Syntax-directed SSA generation
- Bonus: Virgil compiler's SSA CFG

Recap From Last Time

- SSA is an IR for analysis and optimization
- Every modern optimizing compiler uses it!
- Invariant: every variable defined statically once
- Φ -functions are a notational fiction for merging dataflow at joins
- Can build SSA using the dominator tree
- Iterated dominance frontier guides Φ placement
- Renaming step walks down dominator tree
- General algorithm handles all possible CFGs


What is a Φ anyway?



A Φ merges multiple versions of a variable at a **WHAT?** in the CFG.

Inputs *positionally correspond* with predecessor edges.


What is a Φ anyway?



A Φ merges multiple versions of a variable at a join point in the CFG.

Inputs *positionally correspond* with predecessor edges.

What is a Φ (for a loop) anyway?




Φ s at loop headers relate the dataflow on a loop with the control flow.


Allows finding induction variables really easily.

WHAT?

Dominance Frontier



The *dominance frontier* can be determined by looking at CFG edges that *leave a dominator subtree* to non-dominated nodes.



Thinking in SSA: Local Reasoning

- Local Φ simplifications

- $y = \Phi(x, x) \rightsquigarrow$ WHAT?

- $y = \Phi(x, y) \rightsquigarrow$ WHAT?

Thinking in SSA: Local Reasoning

- Local Φ simplifications
 - $y = \Phi(x, x) \rightsquigarrow x$
 - $y = \Phi(x, y) \rightsquigarrow x$

A Φ can be simplified if all inputs are the same, or all inputs are the same or the Φ itself.

Thinking in SSA: Local Reasoning

- Constant propagation

- $x = 13;$

- $y = x + 2;$

- $z = \text{func}(x);$

- $w = x - x;$

→



Thinking in SSA: Local Reasoning

- Constant propagation

- `x = 13;`

- `y = x + 2;`

- `z = func(x);`

- `w = x - x;`

→

- `y = 13 + 2;`

- `z = func(13);`

- `w = 13 - 13;`

After SSA construction, any variable which is assigned a constant can be substituted *everywhere it occurs*.

Thinking in SSA: Local Reasoning

- Copy propagation

- $x = a;$

- $y = x + 2;$

- $z = \text{func}(x);$

- $w = x - x;$

→



WHAT?

Thinking in SSA: Local Reasoning

- Copy propagation

- $x = a;$

- $y = x + 2;$

- $z = \text{func}(x);$

- $w = x - x;$

→

- $y = a + 2;$

- $z = \text{func}(a);$

- $w = a - a;$

After SSA construction, any variable which is assigned a simple copy can be substituted *everywhere it occurs*.

Thinking in SSA: Local Reasoning

- Copy propagation

- $x = a;$

- $y = x + 2;$

- $z = \text{func}(x);$

- $w = x - x;$

→

- $y = a + 2;$

- $z = \text{func}(a);$

- $w = a - a;$

Copies are vestigial.

After SSA construction, any variable which is assigned a simple copy can be substituted *everywhere it occurs*.

Thinking in SSA: Local Reasoning

- Reasoning about value ranges of Φ s
 - $x = \Phi(3, 7);$
 $y = x > 0;$
 $z = x == 12; \rightsquigarrow$
 $w = 33 / x;$



Thinking in SSA: Local Reasoning

- Reasoning about value ranges of Φ s

- $x = \Phi(3, 7);$
 $y = x > 0;$
 $z = x == 12;$
 $w = 33 / x;$
 - \rightsquigarrow
 - $y = \text{true};$
 $z = \text{false};$
 $w = 33 / x \text{ (safe)};$

Φ inputs represent all the possible values the phi could take on across all control flow. Therefore some conditions are statically decidable.

Thinking in SSA: Local Reasoning

- Dominating conditions

- if (x == 0) return x;

→ if (x == 0) return WHAT?

- if (x != null) return x.f;

→ if (x != null) return WHAT?

- if (x != 0) return 1000 / x;

→ if (x != 0) return WHAT?

Thinking in SSA: Local Reasoning

- Dominating conditions

- if (x == 0) return x;

→ if (x == 0) return 0;

- if (x != null) return x.f;

→ if (x != null) return x.f (safe);

- if (x != 0) return 1000 / x;

→ if (x != 0) return 1000 / x;

Any branch can assume the condition is *true or false* on the respective output control flow edges.

Thinking in SSA: Local Reasoning

- Dominating conditions

- if (x == 0) return x;

→ if (x == 0) return 0;

- if (x != null) return x.f;


→ if (x != null) return x.f (safe);

- if (x != 0) return 1000 / x;


→ if (x != 0) return 1000 / x;

Any branch can assume the condition is *true or false* on the respective output control flow edges.

Thinking in SSA: Local Reasoning




Dominator Tree




Any branch can assume the condition is true or false on the respective output control flow edges *and their dominated blocks.*

Thinking in SSA: Local Reasoning




Dominator Tree




Any branch can assume the condition is true or false on the respective output control flow edges *and their dominated blocks.*

Thinking in SSA: Local Reasoning





Dominator Tree





Any branch can assume the condition is true or false on the respective output control flow edges *and their dominated blocks.*

Thinking in SSA: Local Reasoning



Any branch can assume the condition is true or false on the respective output control flow edges *and their dominated blocks.*

Thinking in SSA: Local Reasoning



Any branch can assume the condition is true or false on the respective output control flow edges *and their dominated blocks.*

Thinking in SSA: Local Reasoning

- Branch folding

- if (x == x) {
 $y_1 = a + 3;$
} else {
 $y_2 = a + 5;$
}
 $y_3 = \Phi(y_1, y_2)$

→



Thinking in SSA: Local Reasoning

- Branch folding


```
○ if (x == x) {  
    y1 = a + 3;  
} else {  
    y2 = a + 5;  
}  
y3 = Φ(y1, y2)
```

→


```
if (true) {  
    y1 = a + 3;  
} else {  
    y2 = a + 5;  
}  
y3 = y1
```

Any branch with a constant condition can be folded and edges removed from the CFG.

Thinking in SSA: Local Reasoning




Dominator Tree




Any branch with a constant condition can be folded and edges removed from the CFG.

Thinking in SSA: Local Reasoning




Dominator Tree




The entire dominator subtree will become dead and can be removed.

Thinking in SSA: Local Reasoning




Dominator Tree




The entire dominator subtree will become dead and can be removed.

Thinking in SSA: Local Reasoning





Dominator Tree



The entire dominator subtree will become dead and can be removed.


Thinking in SSA: Local Reasoning

Dominator Tree



To maintain SSA, any Φ s at join points with dead predecessors need to be edited and simplified.


Editing and simplifying Φ s



A Φ merges multiple versions of a variable at a join point in the CFG.

Inputs *positionally correspond* with predecessor edges.


Editing and simplifying Φ s



A Φ merges multiple versions of a variable at a join point in the CFG.

Dead predecessors mean dead input values.

Editing and simplifying Φ s



Dead predecessors mean dead input values.

Simply edit the input values out.


May lead to further simplification.

Deconstructing SSA

- Real machines don't have Φ functions.
- Have to insert moves at predecessors.
- Mentioned earlier, but with huge caveats.
- We resolve those caveats today.


Deconstructing SSA

- When during compilation to deconstruct SSA?
- There are two common choices: before or after regalloc.
- Regalloc before deconstruction is relatively new (2010s).



Deconstructing SSA


- When during compilation to deconstruct SSA?
- There are two common choices: before or after regalloc.
- Regalloc before deconstruction is relatively new (2010s).



Nice chordal interference graphs


Deconstructing SSA

- When during compilation to deconstruct SSA?
- There are two common choices: before or after regalloc.
- Regalloc before deconstruction is relatively new (2010s).



Deconstructing SSA


- When during compilation to deconstruct SSA?
- There are two common choices: before or after regalloc.
- Regalloc before deconstruction is relatively new (2010s).



Deconstruction is more or less the same either way.


Deconstructing SSA

- Insert moves according to the positional correspondence of inputs.




Deconstructing SSA

- Insert Φ -resolution moves and remove Φ s.




Deconstructing SSA

- Insert moves according to the positional correspondence of inputs.



Deconstructing SSA


- Insert moves according to the positional correspondence of inputs.



Notice the alignment of *data flow* and *control flow*. The Φ nodes represent this explicitly in the IR.

Deconstructing SSA

- Insert moves according to the positional correspondence of inputs.



Deconstructing SSA

- Insert moves according to the positional correspondence of inputs.

```
a1 ← x + y
b1 ← a1 + x
a2 ← a1
c3 ← c1
```

```
a2 ← b + 2
c2 ← y + 1
a3 ← a2
c3 ← c2
```

Remove Φ s after inserting moves.

```
a3 ←  $\Phi(a_1, a_2)$ 
c3 ←  $\Phi(c_1, c_2)$ 
a4 ← c3 + a3
```

Deconstructing SSA

- Insert moves according to the positional correspondence of inputs.

```
a1 ← x + y  
b1 ← a1 + x  
a3 ← a1  
c3 ← c1
```

```
a2 ← b + 2  
c2 ← y + 1  
a3 ← a2  
c3 ← c2
```


Removing all Φ s after deconstruction gives a completely valid non-SSA program.

```
a4 ← c3 + a3
```

The program is now directly executable again.

Issue 1: Critical Edges

- Consider a simple triangle CFG.




Issue 1: Critical Edges

- Consider a simple triangle CFG.

But wait, there's a bug in the SSA for this program!

What's the bug?




Issue 1: Critical Edges

- Consider a simple triangle CFG.


But wait, there's a bug in the SSA for this program!

What's the bug?




Issue 1: Critical Edges

- Consider a simple triangle CFG.




Issue 1: Critical Edges

- Consider a simple triangle CFG.
- We insert moves in both predecessors and remove the Φ .




Issue 1: Critical Edges

- Consider a simple triangle CFG.
- We insert moves in both predecessors and remove the Φ .



Issue 1: Critical Edges


- Consider a simple triangle CFG.
- We insert moves in both predecessors and remove the Φ .



Issue 1: Critical Edges

- Consider a simple triangle CFG.
- We insert moves in both predecessors and remove the Φ .


Naïve insertion can introduce redundant code on some execution paths.



Issue 1: Critical Edges

- Consider a simple triangle CFG.
- We insert moves in both predecessors and remove the Φ .


Naïve insertion can introduce redundant code on some execution paths.



Issue 1: Critical Edges

- Consider a *more complicated* CFG.
- We insert moves in *all* predecessors and remove the Φ .


Naïve insertion can introduce redundant code on some execution paths.



Issue 1: Critical Edges

- Consider a *more complicated* CFG.
- We insert moves in *all* predecessors and remove the Φ .

Naïve insertion can introduce redundant code on some execution paths.




Issue 1: Critical Edges

- Consider a *more complicated* CFG.
- We insert moves in *all* predecessors and remove the Φ .


Naïve insertion can introduce redundant code on some execution paths.

Can actually get *really* bad.




Splitting Critical Edges

- To avoid redundant moves, split *critical edges* by inserting an empty block between.




Splitting Critical Edges

- To avoid redundant moves, split *critical edges* by inserting an empty block between.



Splitting Critical Edges


- To avoid redundant moves, split *critical edges* by inserting an empty block between.
- This block is the proper place for Φ -resolution moves.



Splitting Critical Edges

- To avoid redundant moves, split *critical edges* by inserting an empty block between.
- This block is the proper place for Φ -resolution moves.


A *critical edge* is any edge that connects a block with multiple successors to a block with multiple predecessors.



Splitting Critical Edges

- To avoid redundant moves, split *critical edges* by inserting an empty block between.
- This block is the proper place for Φ -resolution moves.

Splitting all critical edges prior to SSA deconstruction is easy.




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- For CFGs without loops, *no*.
- Let's convince ourselves.


Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.
- **By SSA invariants, the definition of the RHS of each move dominates the move.**




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.
- **By SSA invariants, the definition of the RHS of each move dominates the move.**




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.
- By SSA invariants, the definition of the RHS of each move dominates the move.
- **It cannot be the case that the LHS is live, because previously there was only one definition, below.**




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.
- By SSA invariants, the definition of the RHS of each move dominates the move.
- It cannot be the case that the LHS is live, because previously there was only one definition, below.
- **Therefore we are only assigning to fresh variables, and not overwriting anything.**




Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.
- By SSA invariants, the definition of the RHS of each move dominates the move.
- It cannot be the case that the LHS is live, because previously there was only one definition, below.
- Therefore we are only assigning to fresh variables, and not overwriting anything.
- **Therefore any order is fine.**



Issue 2: Ordering Moves


- Does the order of Φ -resolution moves matter?
- Consider a join with at least two Φ s.
- Moves are inserted into predecessors.
- By SSA invariants, the definition of the RHS of each move dominates the move.
- It cannot be the case that the LHS is live, because previously there was only one definition, below.
- Therefore we are only assigning to fresh variables, and not overwriting anything.
- **Therefore any order is fine.**



Issue 2: Ordering Moves

- Does the order of Φ -resolution moves matter?
- For CFGs without loops, *no*.
- But what about loops?


Issue 2: Ordering Moves



Φ s at loop headers relate the dataflow on a loop backedge with the control flow.


A loop Φ can be defined in terms of itself.

Issue 2: Ordering Moves



Like any other join, we insert Φ -resolution moves at predecessors.


Issue 2: Ordering Moves



Like any other join, we insert Φ -resolution moves at predecessors.

With only one Φ , there is no problem yet.


Issue 2: Ordering Moves



Like any join, a loop header can have multiple Φ s.


Because Φ s can use inductively defined versions of themselves, they can be recursive or even *mutually recursive*.

Issue 2: Ordering Moves




A simple example:
swap of variables in
a loop.

Issue 2: Ordering Moves




Issue 2: Ordering Moves



After optimizations such as copy propagation, the Φ s can be mutually recursive.


This is totally legal and cool.

Issue 2: Ordering Moves




SSA deconstruction using the naïve move insertion will always generate incorrect code, regardless of the order.


Issue 2: Ordering Moves




SSA deconstruction using the naïve move insertion will always generate incorrect code, regardless of the order.



Issue 2: Ordering Moves



The reason is that phi resolution moves have *parallel move* semantics.



Implementing Parallel Moves

- Φ resolution moves must be done in parallel, without overwriting old versions.
- One simple solution: introduce new temps again.

$$\begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{pmatrix} \leftarrow \begin{pmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \end{pmatrix}$$

generates

$$\begin{aligned} \mathbf{t}_0 &\leftarrow \mathbf{y}_0 \\ \mathbf{t}_1 &\leftarrow \mathbf{y}_1 \\ \mathbf{t}_2 &\leftarrow \mathbf{y}_2 \\ \mathbf{t}_3 &\leftarrow \mathbf{y}_3 \\ \mathbf{x}_0 &\leftarrow \mathbf{t}_0 \\ \mathbf{x}_1 &\leftarrow \mathbf{t}_1 \\ \mathbf{x}_2 &\leftarrow \mathbf{t}_2 \\ \mathbf{x}_3 &\leftarrow \mathbf{t}_3 \end{aligned}$$

Works every time.

Generates *a lot* of temporaries, but maybe the register allocator / copy propagation can clean them up?

Implementing Parallel Moves

- Φ resolution moves must be done in parallel, without overwriting old versions.
- Better solution: order moves more intelligently.

$$\begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{pmatrix} \leftarrow \begin{pmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \end{pmatrix}$$

Implementing Parallel Moves

- Φ resolution moves must be done in parallel, without overwriting old versions.
- Better solution: order moves more intelligently.

$$\begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{pmatrix} \leftarrow \begin{pmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \end{pmatrix}$$

Notice that because parallel moves originate from SSA deconstruction, variables on the LHS appear only once on the LHS.

$$\mathbf{x}_0 \neq \mathbf{x}_1 \neq \mathbf{x}_2 \neq \mathbf{x}_3$$

Implementing Parallel Moves


- Φ resolution moves must be done in parallel, without overwriting old versions.
- Better solution: order moves more intelligently *using LTG*.

$$\begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{pmatrix} \leftarrow \begin{pmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \end{pmatrix}$$

We can build a graph where each node in the parallel moves gets a node, and directed edges represent moves.

$$\mathbf{x}_0 \neq \mathbf{x}_1 \neq \mathbf{x}_2 \neq \mathbf{x}_3$$

Location Transfer Graph



Implementing Parallel Moves


- Φ resolution moves must be done in parallel, without overwriting old versions.
- Better solution: order moves more intelligently *using LTG*.

$$\begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{pmatrix} \leftarrow \begin{pmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_1 \end{pmatrix}$$

Variables may appear *multiple times* on the RHS, and may appear on both LHS and RHS.



$$\mathbf{x}_0 \neq \mathbf{x}_1 \neq \mathbf{x}_2 \neq \mathbf{x}_3$$

Location Transfer Graph




Location Transfer Graphs

- A location transfer graph represents a set of parallel moves.
- It can be traversed to generate a legal move ordering.
- It's constrained:
 - Every node in the graph has at most one incoming edge.
 - That implies the graph can only have simple cycles.





Location Transfer Graphs

- A location transfer graph represents a set of parallel moves.
- It can be traversed to generate a legal move ordering.
- It's constrained:
 - Every node in the graph has at most one incoming edge.
 - That implies the graph can only have simple cycles.



Location Transfer Graphs



- A location transfer graph represents a set of parallel moves.
- It can be traversed to generate a legal move ordering.
- It's constrained:
 - Every node in the graph has **at most one incoming edge**.
 - That implies the graph can **only have simple cycles**.



impossible to
have nested
cycles



Location Transfer Graphs

- A location transfer graph represents a set of parallel moves.
- It can be traversed to generate a legal move ordering.
- It's constrained:
 - Every node in the graph has at most one incoming edge.
 - That implies the graph can only have simple cycles.



Location Transfer Graphs



- A location transfer graph represents a set of parallel moves.
- It can be traversed to generate a legal move ordering.
- It's constrained:
 - Every node in the graph has at most one incoming edge.
 - That implies the graph can only have simple cycles.



Can leave a cycle though!

Location Transfer Graphs

- A location transfer graph represents a set of parallel moves.
- It can be traversed to generate a legal move ordering.
- It's constrained:
 - Every node in the graph has at most one incoming edge.
 - That implies the graph can only have simple cycles.




Also known as
“windmill graphs” [1]

[1] See “SSA Elimination after Register Allocation” by Pereira and Palsberg, 2009.

Location Transfer Graphs


- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])



[1] See “SSA Elimination after Register Allocation” by Pereira and Palsberg, 2009.


Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])





Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])




Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])



Location Transfer Graphs



- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])



Location Transfer Graphs




- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])

Break cycle by saving in a temp, **overwriting**, and using the temp instead.






Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])






Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])






Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])






Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])



Location Transfer Graphs

- Post-order depth-first search (DFS) on this graph yields a legal move ordering.
- Must break cycles with a temporary (or use swaps [1])



SSA Elimination After Register Allocation

- Breaking cycles in LTG is easy with a temporary
- Before regalloc, temporaries are plentiful
- If deconstructing after register allocation, all regs may be used
 - How do we know? May need to save something from regalloc
- It's possible to implement LTG with swaps and/or sequences of xor's
- Another option is to use a temporary slot
- See paper for details:
 - [“SSA Elimination after Register Allocation” by Pereira and Palsberg, 2009.](#)

Syntax-directed SSA Construction

- Is it possible to skip dominators/dominance frontier?
- “Single-Pass Generation of Static Single-Assignment Form for Structured Languages” by Brandis and Mössenbock (1994).
 - <https://dl.acm.org/doi/10.1145/197320.197331>
- Also done in my Virgil compiler
 - <https://github.com/titzer/virgil>