

SSA

15-411/15-611 Compiler Design

Seth Copen Goldstein

September 16, 2021

Using DF to compute minimal SSA

- place all $\Phi()$
- Rename all variables

minimal, but not pruned

Using DF to Place $\Phi()$

- Gather all the defsites of every variable
- Then, for every variable
 - foreach defsite
 - foreach node in DF(defsite)
 - if we haven't put $\Phi()$ in node put one in
 - If this node didn't define the variable before: add this node to the defsites
- This essentially computes the Iterated Dominance Frontier on the fly, creating minimal SSA

Using DF to Place $\Phi()$

```
foreach node n {
  foreach variable v defined in n {
    orig[n]  $\cup$ = {v}
    defsites[v]  $\cup$ = {n}
  }
  foreach variable v {
    W = defsites[v]
    while W not empty {
      foreach y in DF[n]
      if y  $\notin$  PHI[v] {
        insert " $v \leftarrow \Phi(v,v,...)$ " at top of y
        PHI[v] = PHI[v]  $\cup$  {y}
        if v  $\notin$  orig[y]: W = W  $\cup$  {y}
      }
    }
  }
}
```

Renaming Variables

- Walk the D-tree, renaming variables as you go
- Replace uses with more recent renamed def
 - For straight-line code this is easy
 - If there are branches and joins?

Renaming for Straight-Line Code

- Need to extend for ϕ -functions.
- Need to maintain property that definitions dominate uses.

for each variable a :

Count[a] = 0

Stack[a] = [0]

renameBasicBlock(B):

for each instruction S in block B :

for each use of a variable x in S :

$i = \text{top}(\text{Stack}[x])$

replace the use of x with x_i

for each variable a that S defines

count[a] = Count[a] + 1

$i = \text{Count}[a]$

push i onto Stack[a]

replace definition of a with a_i

Renaming in CFG

rename(n):

renameBasicBlock(n)

for each successor Y of n, **where** n is the j^{th} predecessor of Y:

for each phi-function f in Y, **where** the operand of f is 'a'

$i = \text{top}(\text{Stack}[a])$

replace j^{th} operand with a_i

for each child of n in D-tree, X:



rename(X)

for each instruction $S \in n$:

for each variable v that S defines:



pop Stack[v]

Compute D-tree



D-tree


Compute Dominance Frontier



1	{}
2	{2}
3	{2}
4	{}
5	{7}
6	{7}
7	{2}

DFs

Insert $\Phi()$



		orig[n]
1	{}	1 {i,j,k}
2	{2}	2 {}
3	{2}	3 {}
4	{}	4 {}
5	{7}	5 {j,k}
6	{7}	6 {j,k}
7	{2}	7 {}

	defsites[v]
i	{1}
j	{1,5,6}
k	{1,5,6}

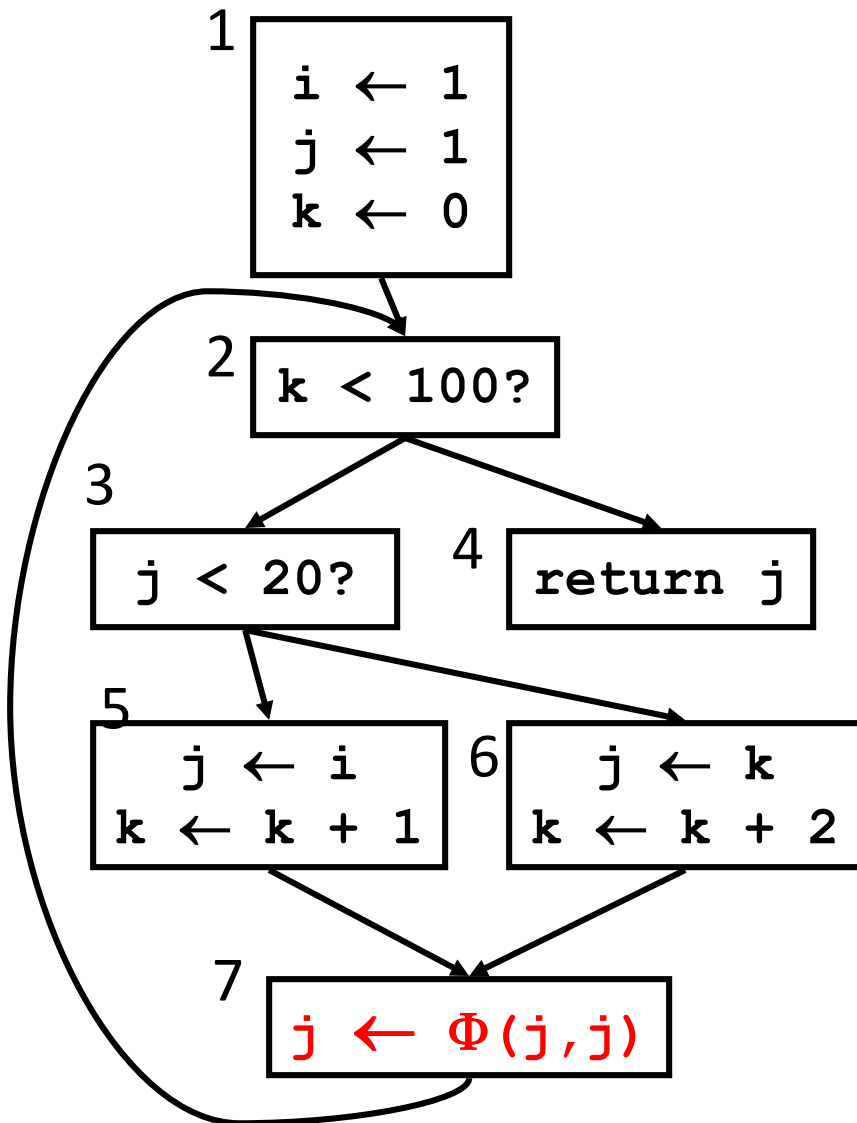
DFs

var i: $W=\{1\}$

var j: $W=\{1,5,6\}$

DF{1}, DF{5}

Insert $\Phi()$



DFs

1	{}
2	{2}
3	{2}
4	{}
5	{7}
6	{7}
7	{2}

orig[n]

1	{ i,j,k }
2	{}
3	{}
4	{}
5	{j,k}
6	{j,k}
7	{}


defsites[v]

i	{1}
j	{1,5,6}
k	{1,5,6}

var j: W={1,5,6}

DF{1}, DF{5}

Insert $\Phi()$



		orig[n]
1	{}	1 {i,j,k}
2	{2}	2 {}
3	{2}	3 {}
4	{}	4 {}
5	{7}	5 {j,k}
6	{7}	6 {j,k}
7	{2}	7 {}

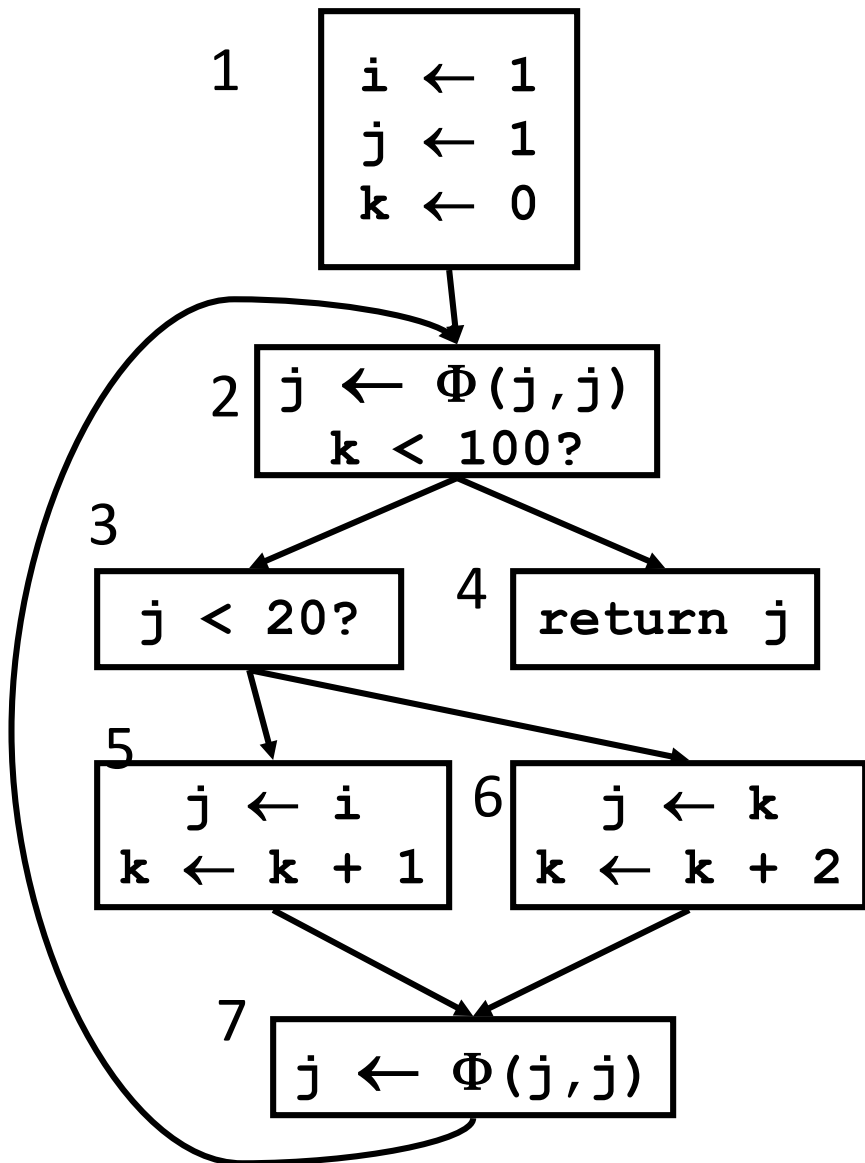
	defsites[v]
i	{1}
j	{1,5,6}
k	{1,5,6}

DFs

var j: W={1,5,6}

DF{1}, DF{5}

Insert $\Phi()$



		orig[n]
1	{}	1 {i,j,k}
2	{2}	2 {}
3	{2}	3 {}
4	{}	4 {}
5	{7}	5 {j,k}
6	{7}	6 {j,k}
7	{2}	7 {}

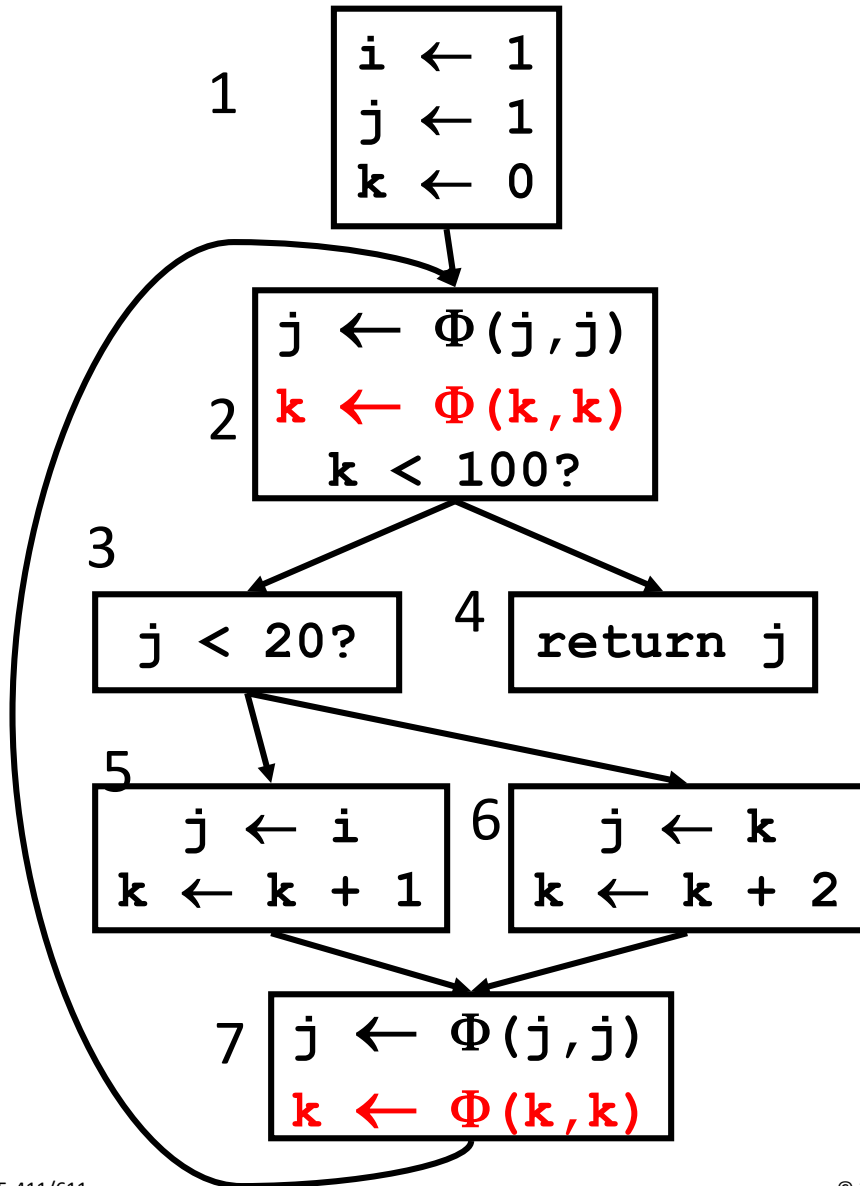
	defsites[v]
i	{1}
j	{1,5,6}
k	{1,5,6}

DFs

var j: W={1,5,6}

DF{1}, DF{5}, **DF{6}**

Insert $\Phi()$




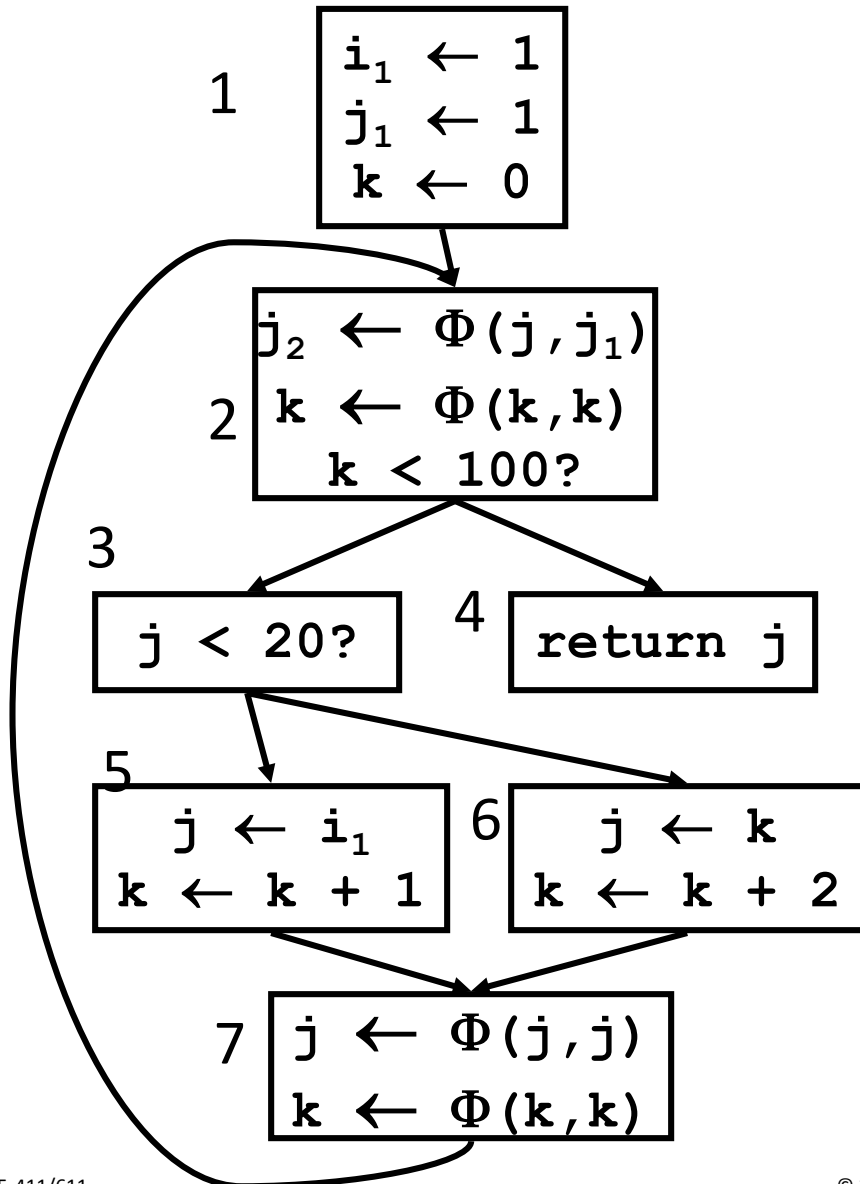
		orig[n]
1	{}	1 {i,j,k}
2	{2}	2 {}
3	{2}	3 {}
4	{}	4 {}
5	{7}	5 {j,k}
6	{7}	6 {j,k}
7	{2}	7 {}

	defsites[v]
i	{1}
j	{1,5,6}
k	{1,5,6}


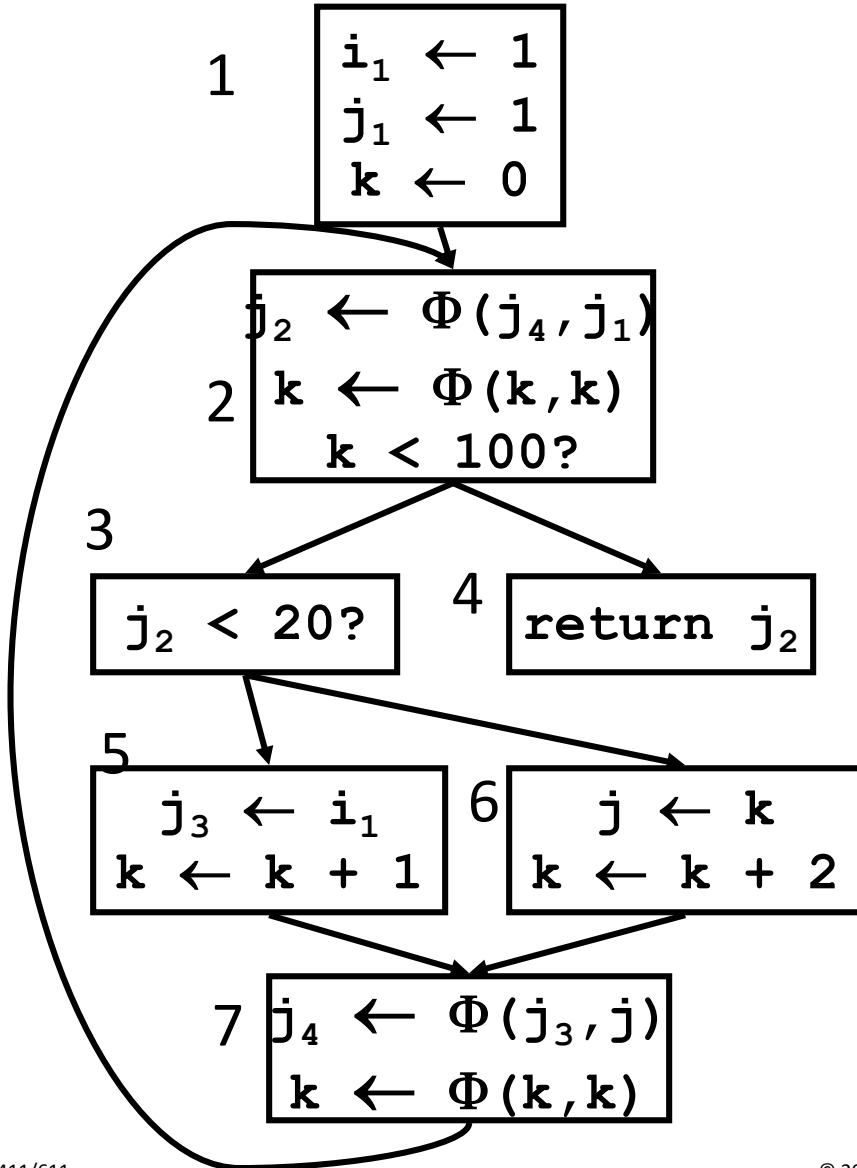
DFs

var k: $W=\{1,5,6\}$


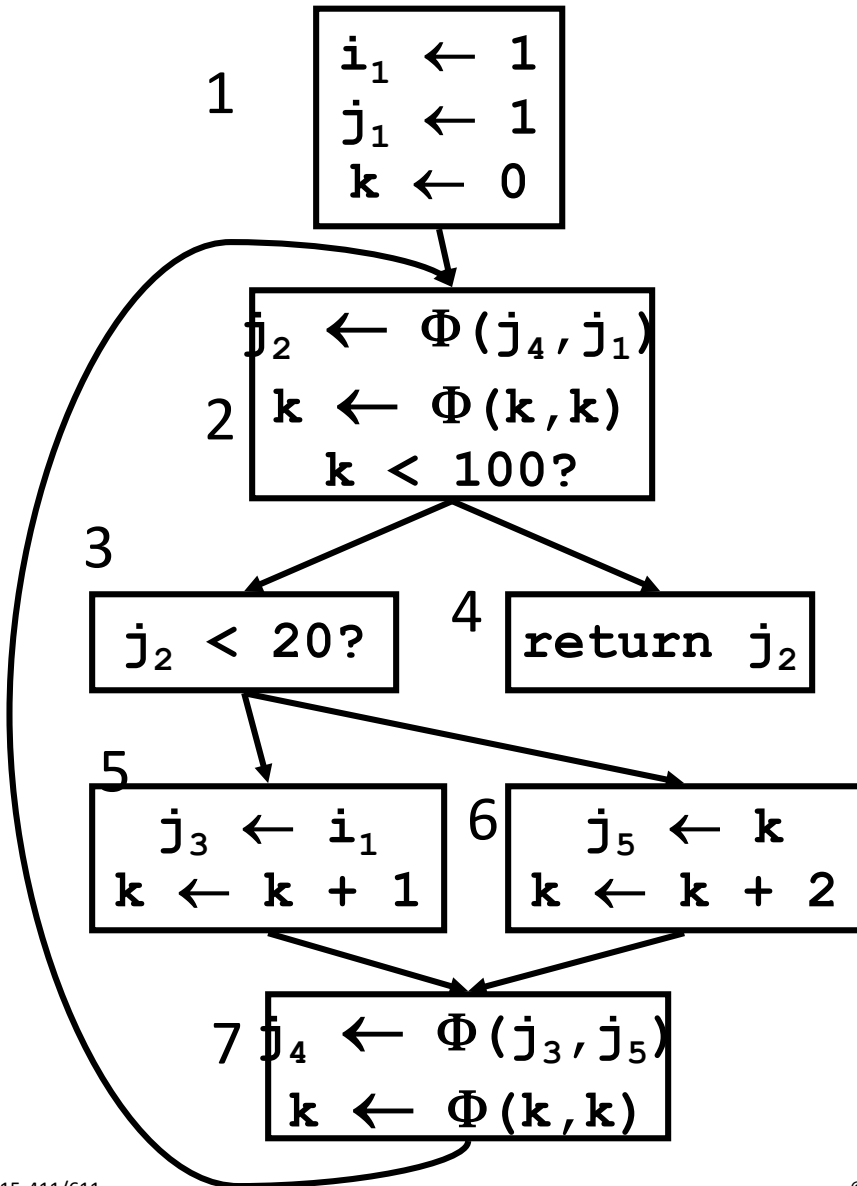
Rename Vars




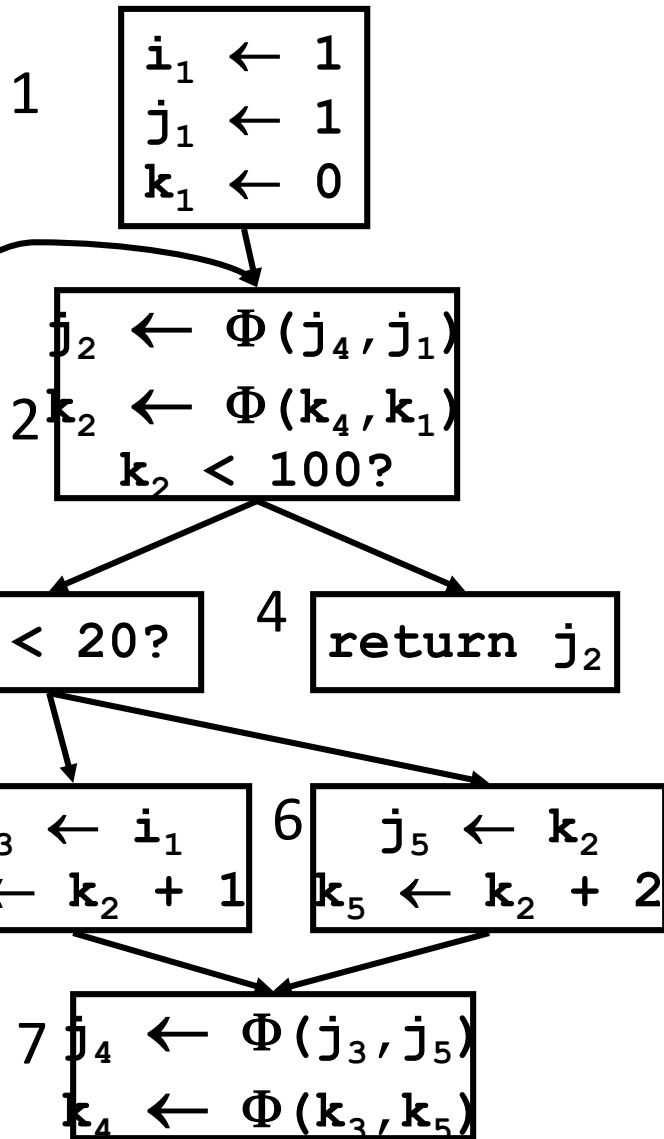
Rename Vars



Rename Vars



Rename Vars



Flavors of SSA

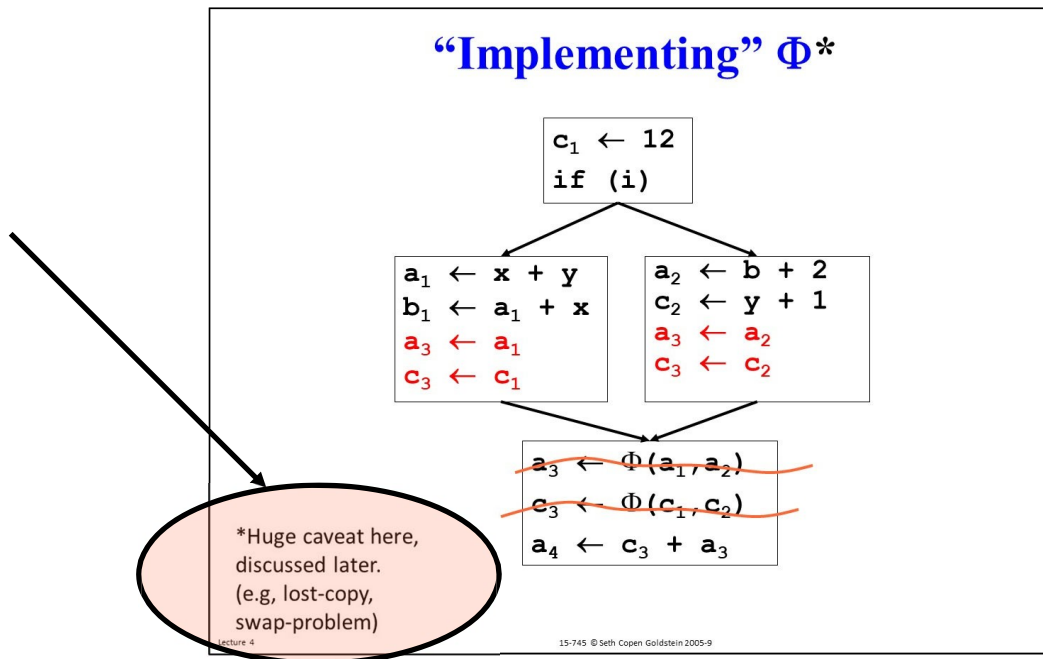
- Minimal SSA
 - at each join point with >1 outstanding definition insert a ϕ -function
 - Some may be dead
- Pruned SSA
 - only add live ϕ -functions
 - must compute LIVEOUT
- Semi-pruned SSA
 - Same as minimal SSA, but only on names live across more than 1 basic block

Summary – getting into

- SSA is a useful and efficient IR.
- Definitions dominate Uses
- Constructing SSA can be efficient
(No need to do Lengaur-Tarjan Algorithm,
instead see [A Simple, Fast Dominance
Algorithm by Cooper, Harvey, and Kennedy](#))
- Don't do any optimizations yet!


But, ...

- Eventually, have to get out of SSA and deconstruct all the Φ -functions
- Recall from Lecture 4




But, ...

- Eventually, have to get out of SSA and deconstruct all the Φ -functions
- Two choices:
 - Before register allocation



– After register allocation



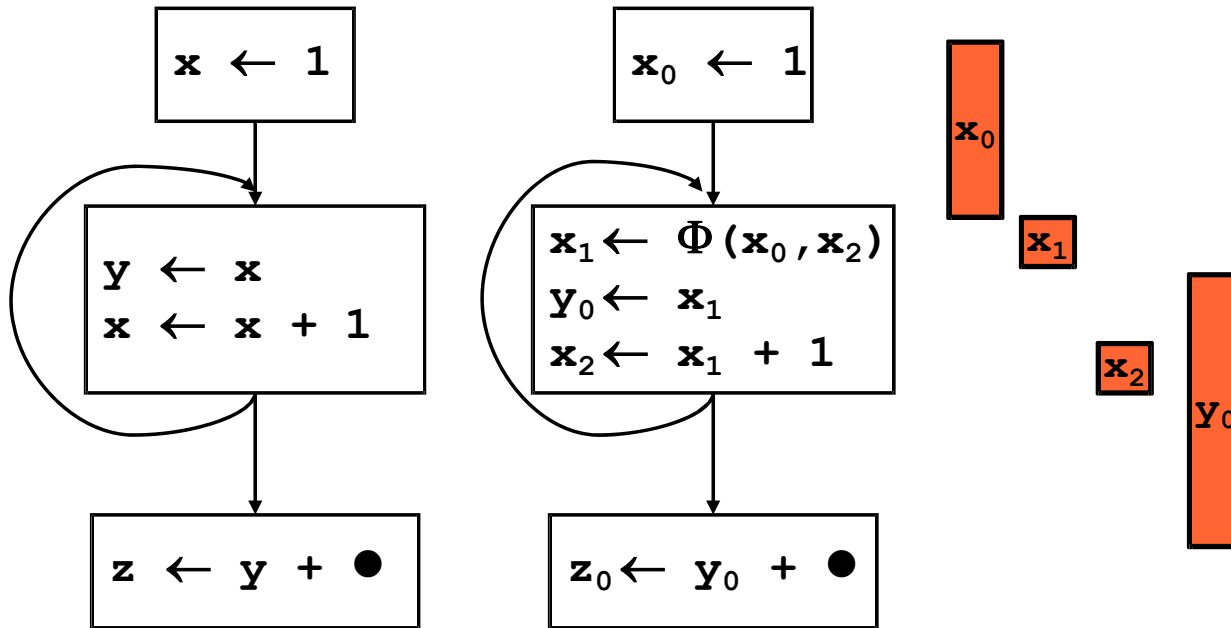
When to deconstruct SSA?

- Before register allocation
 - deconstructing SSA can introduce lots of copies which are easier to eliminate without register constraints
- After register allocation 😊
 - Enables decoupled register allocation
 - spill, color, coalesce
 - Φ -functions may have sources which are registers and memory.
 - Complicated by code-motion optimizations

Conventional-SSA

- Initial Conversion to SSA creates
“Conventional SSA”
- Main feature:
 - variables involved in a Φ -function never interfere
 - Thus, can allocate to a single **resource**
 - the same register
 - the same frame slot (in case of spill)
- However, code motion can destroy this property


→ SSA




Original

In SSA


Leaving SSA



Original

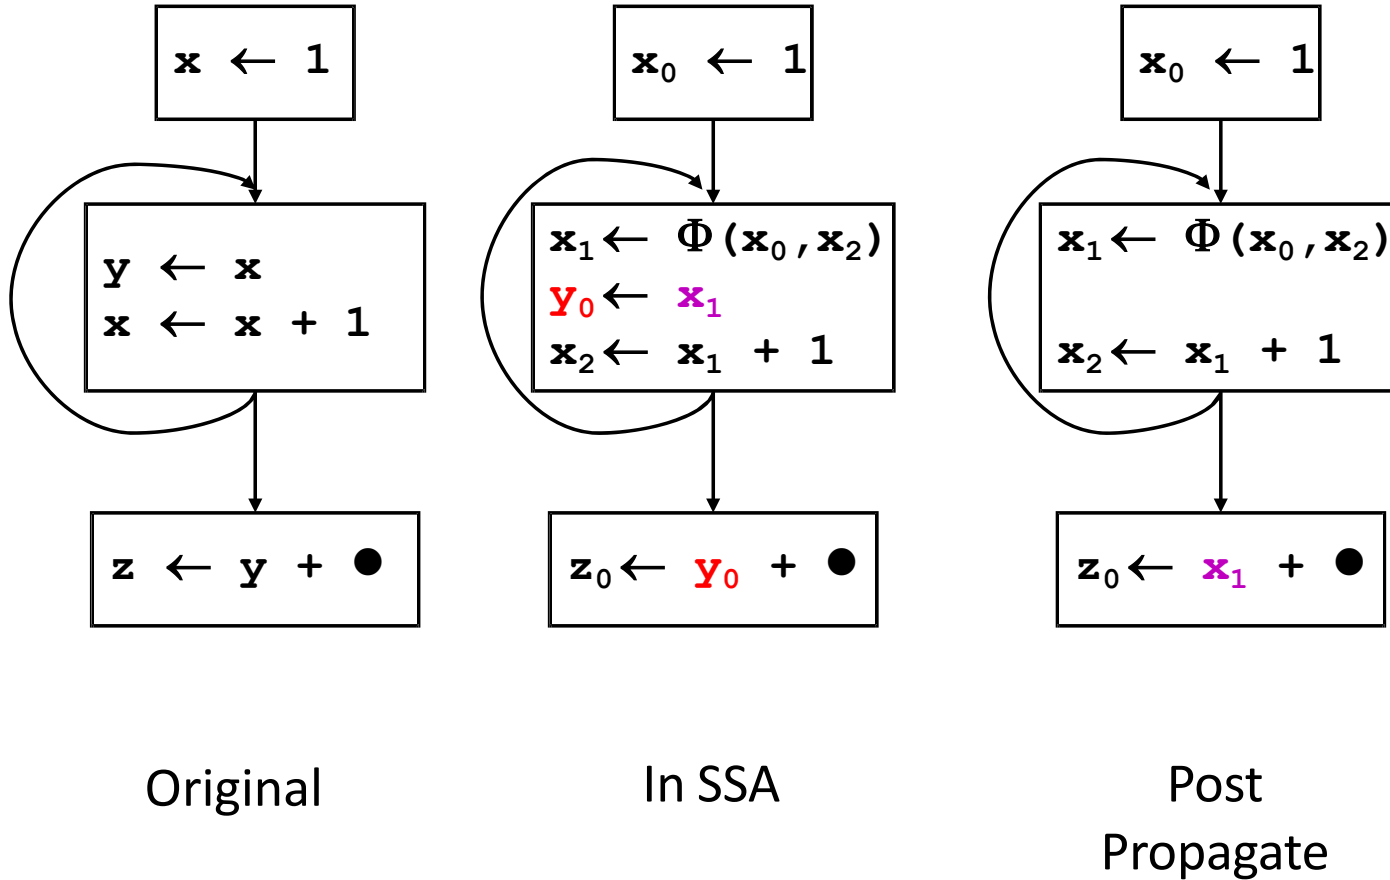


In SSA




Out of SSA

Copy Propagation



Leaving SSA After Copy Folding



Original


Post Prop

Out of SSA

"Lost Copy" Problem

Critical Edges

critical edge



Copy
Propagated


INCORRECT

A critical edge is an edge $a \rightarrow b$ where


- a has > 1 successor and
- b has > 1 predecessor.

Critical Edges


critical edge



Copy Folded




INCORRECT




Inserting block on critical edge

Critical Edges


critical edge



Copy Folded




INCORRECT



Inserting block on critical edge,
Then deconstruct


Semantics of Φ -functions

- Semantics of Φ -functions requires copies to be done in parallel.




Semantics of Φ -functions

- Semantics of Φ -functions requires copies to be done in parallel.



Semantics of Φ -functions

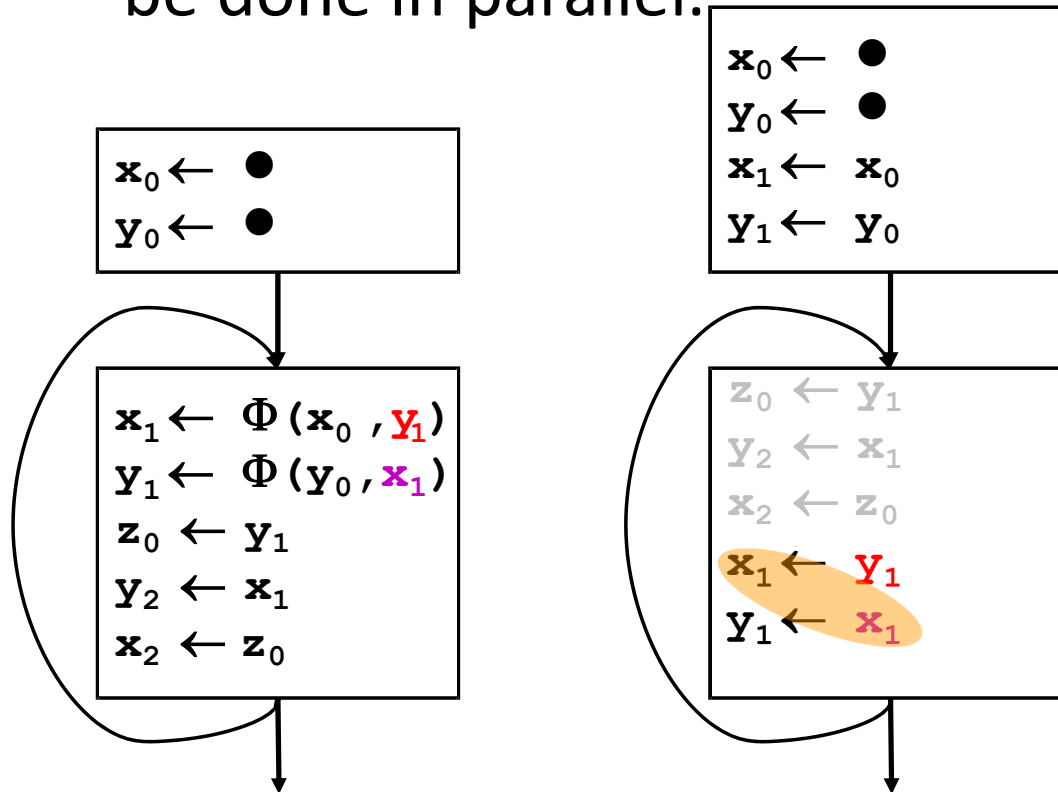
- Semantics of Φ -functions requires copies to be done in parallel.



INCORRECT

Semantics of Φ -functions

- Semantics of Φ -functions requires copies to be done in parallel.

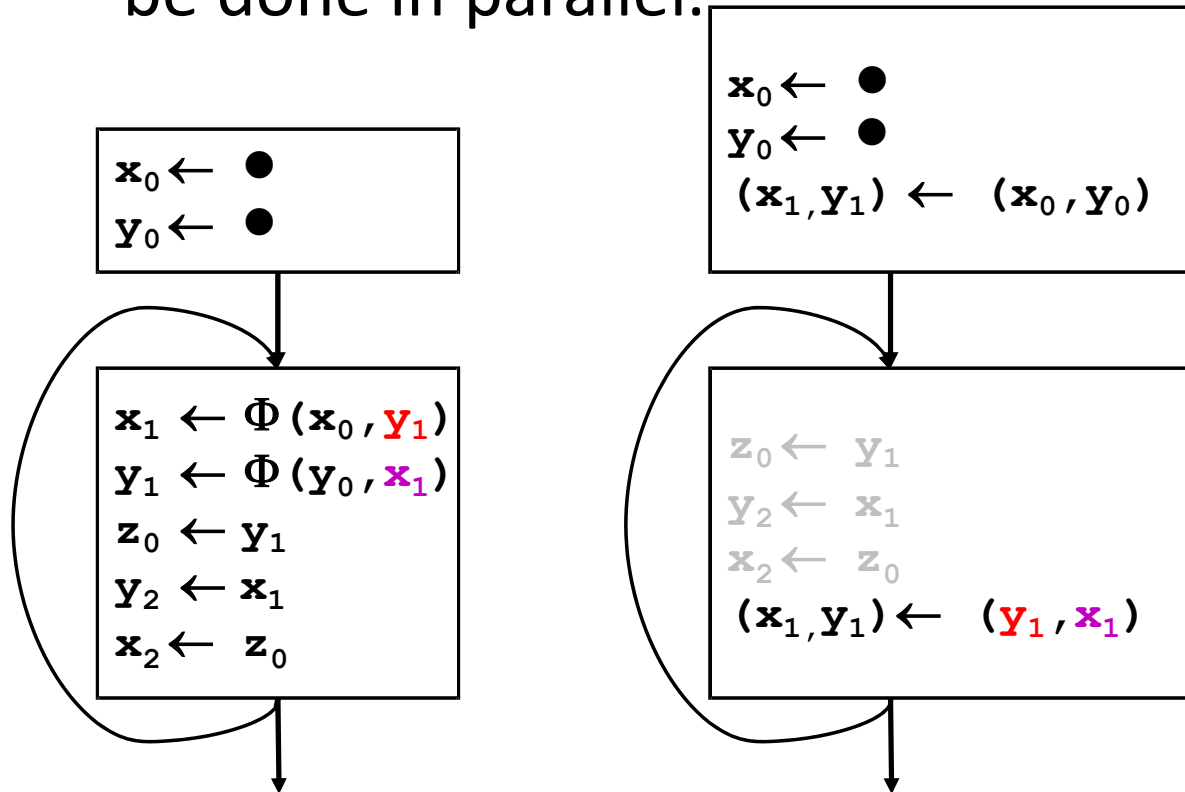


Lost value of x_1 !
because did Φ
assignments
sequentially

INCORRECT

Semantics of Φ -functions

- Semantics of Φ -functions requires copies to be done in parallel.



Using Parallel copies

Impact of Spilling

- What happens when we spill a Φ related variable?
- For example:
 - $\mathbf{r} \leftarrow \Phi(\mathbf{r}, \mathbf{m}_0)$
 - $\mathbf{m}_1 \leftarrow \Phi(\mathbf{r}, \mathbf{m}_0)$
- Could require memory-memory move after deconstructing SSA


Solution

- Critical Edge Splitting
- Convert back to Conventional-SSA (CSSA)
- Register Allocation
 - Build interference graph
 - pre-spilling
 - coloring
- Deconstruct SSA
 - put parallel-copies in predecessors
 - Eliminate parallel copies
- Coalescing

Note: we changed traditional register allocation sequence


Removing a Critical Edge

- A critical edge is an edge $a \rightarrow b$ where
 - a has > 1 successor and
 - b has > 1 predecessor.
- For each edge (a,b) in CFG where $a > 1$ succ and $b > 1$ pred
 - Insert new block Z
 - replace (a,b) with
 - (a,z) and (z,b)




Removing a Critical Edge

- A critical edge is an edge $a \rightarrow b$ where
 - a has > 1 successor and
 - b has > 1 predecessor.
- For each edge (a,b) in CFG where $a > 1$ succ and $b > 1$ pred
 - Insert new block Z
 - replace (a,b) with
 - (a,z) and (z,b)




Removing a Critical Edge

- A critical edge is an edge $a \rightarrow b$ where
 - a has > 1 successor and
 - b has > 1 predecessor.
- For each edge (a,b) in CFG where $a > 1$ succ and $b > 1$ pred
 - Insert new block Z
 - replace (a,b) with
 - (a,z) and (z,b)



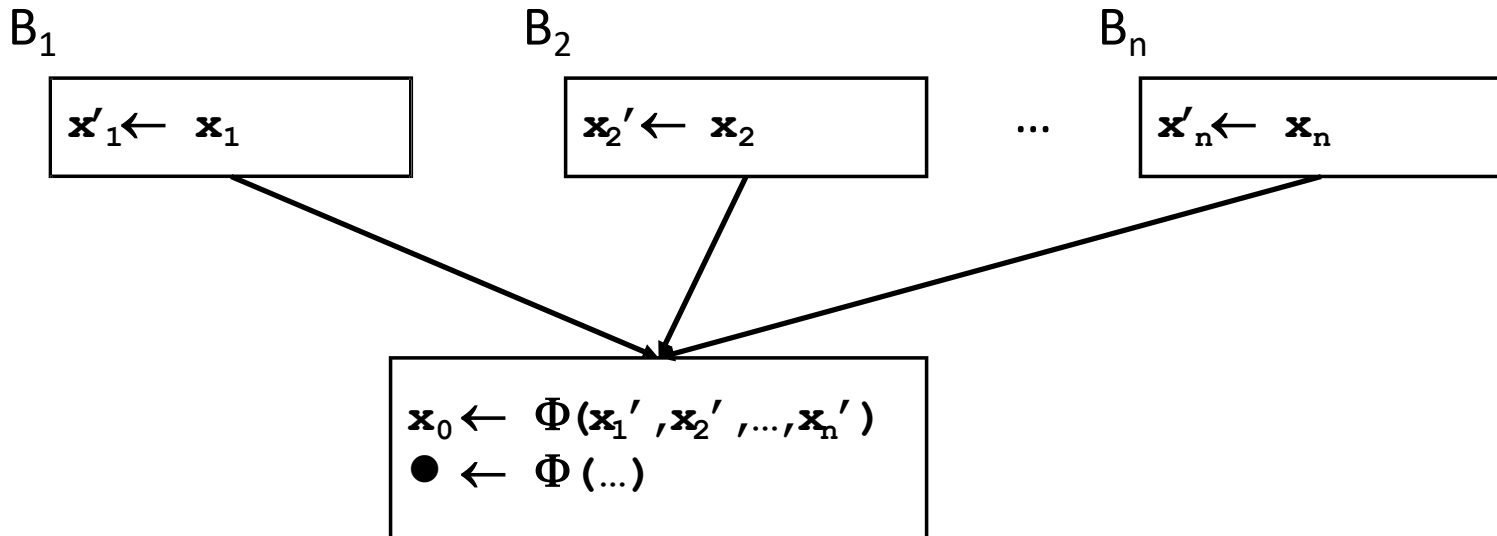
Converting to CSSA

- Goal is to ensure that all Φ related variables do NOT interfere
 - insert copies to (possibly) split live ranges



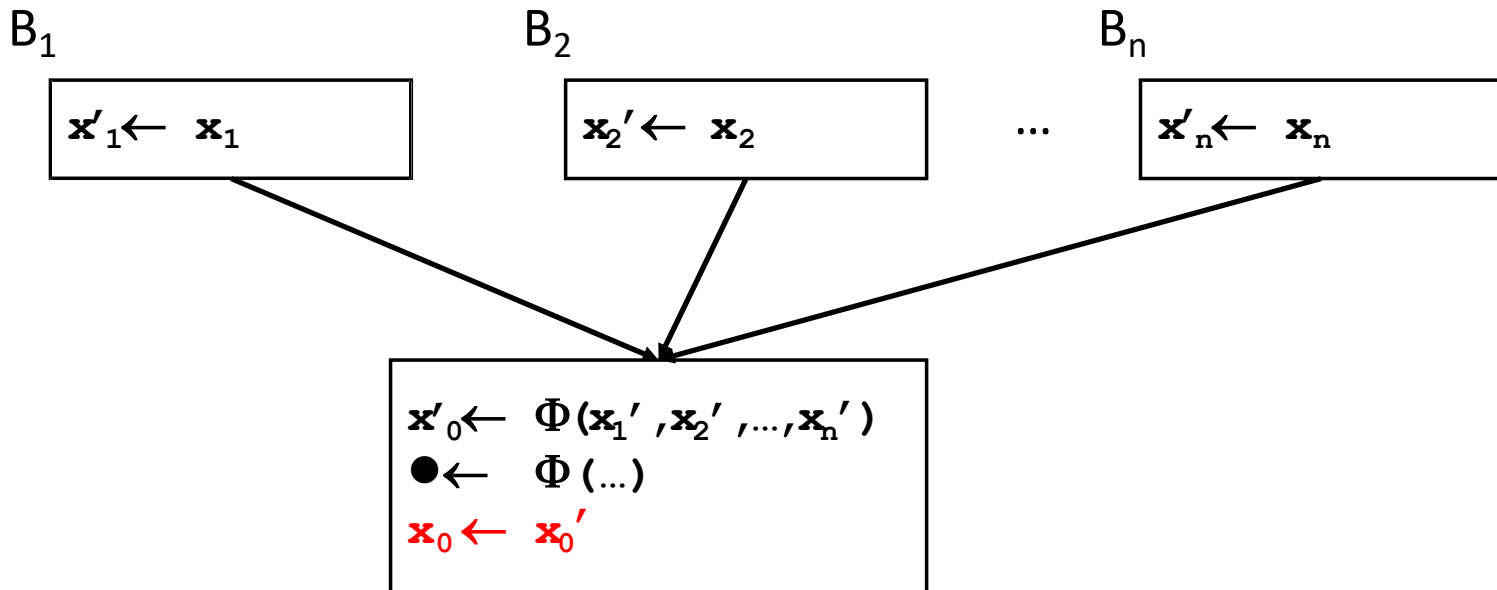
Converting to CSSA

- Goal is to ensure that all Φ related variables do NOT interfere
 - For each argument, insert copy at end of predecessor block and use copy in Φ -function



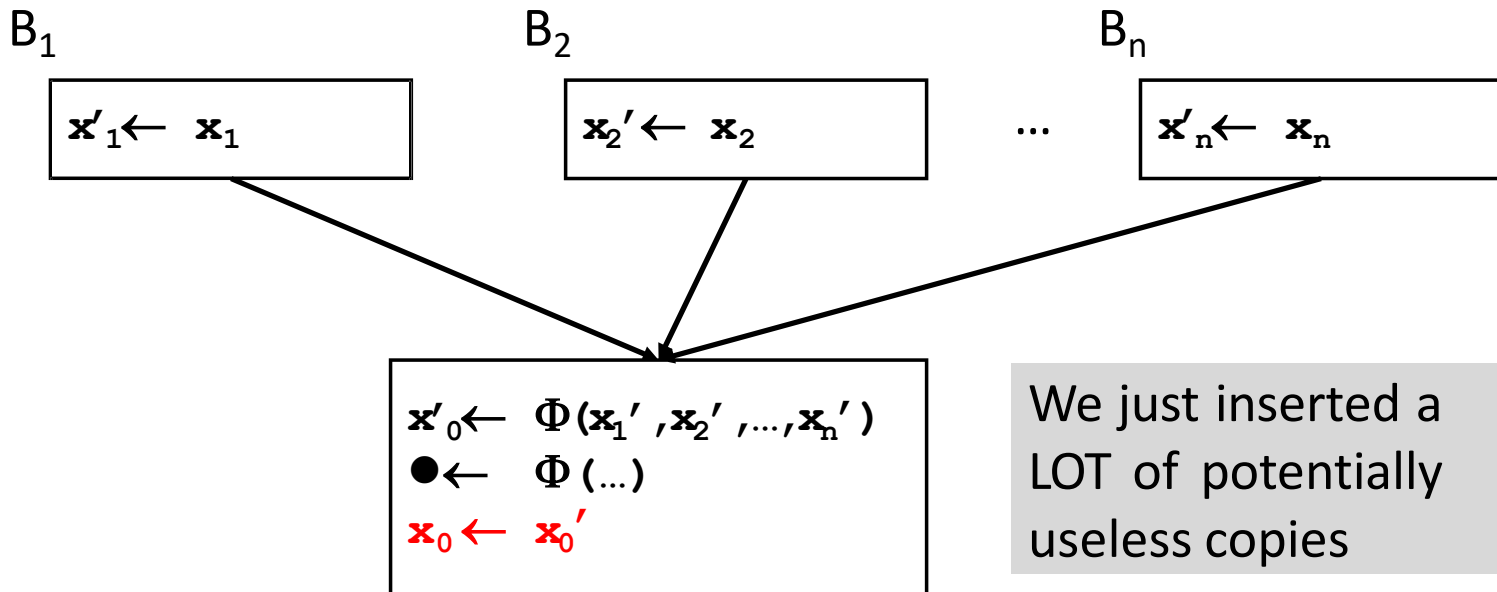
Converting to CSSA

- Goal is to ensure that all Φ related variables do NOT interfere
 - For each argument, insert copy at end of predecessor block and use copy in Φ -function
 - Rename destination
 - Insert copy **AFTER** all Φ -functions in block



Converting to CSSA

- Goal is to ensure that all Φ related variables do NOT interfere
 - For each argument, insert copy at end of predecessor block and use copy in Φ -function
 - Rename destination
 - Insert copy **AFTER** all Φ -functions in block



Register Allocation on CSSA

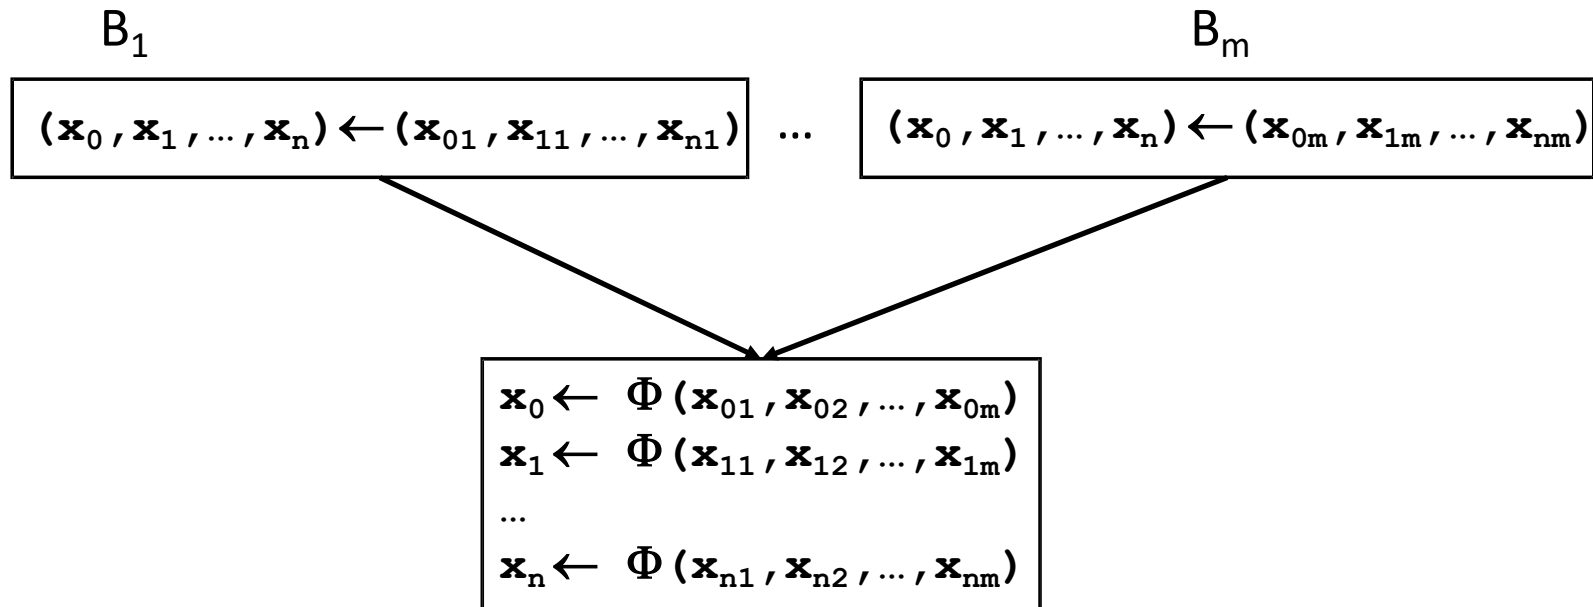
- Build interference graph
- Pre-spill to make it colorable
 - If spill a Φ -related variable, make sure all from same Φ -function use same memory slot!
 - Why do we know this is ok?
 - [Cheat 1: if you spill one, spill them all]
- Color using SEO

Elimination of Φ -functions

- Put parallel copies in predecessor blocks
- Sequentialize the parallel copies

Elimination of Φ -functions

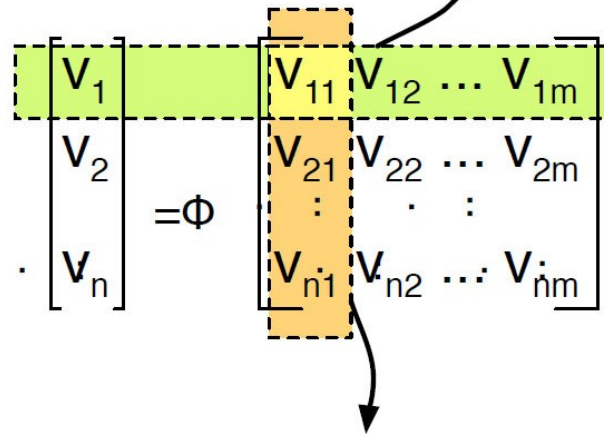
- Put parallel copies in predecessor blocks
- Sequentialize the parallel copies



Elimination of Φ -functions

- Put parallel copies in predecessor blocks
- Sequentialize the parallel copies

$$\Phi\text{-function: } v_1 = \Phi(v_{11}, v_{12}, \dots, v_{1m})$$




parallel copy

$$(v_1, v_2, \dots, v_m) := (v_{11}, v_{12}, \dots, v_{1m})$$

[Pereira&Palsberg 2010]

Parallel Copies

- $(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n) \leftarrow (\mathbf{x}_{01}, \mathbf{x}_{11}, \dots, \mathbf{x}_{n1})$
- Each parallel copy forms a “location transfer graph” [Pereira&Palsberg 2010]
 - edges in graph are the pairwise copies that need to be performed
- In LTG, in-degree is at most 1



Parallel Copies

- $(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n) \leftarrow (\mathbf{x}_{01}, \mathbf{x}_{11}, \dots, \mathbf{x}_{n1})$
- Each parallel copy forms a “location transfer graph” [Pereira&Palsberg 2010]
 - edges in graph are the pairwise copies that need to be performed
- In LTG, in-degree is at most 1
- If we spilled correctly (e.g., all Φ -related variables are spilled to same slot), then also:
 - out-degree of any node is at most 1
 - if node in graph is memory location, then

Spartan Transfer Graphs

- If we spilled correctly (e.g., all Φ -related variables are spilled to same slot), then also:
 - in-degree of any node is at most 1
 - out-degree of any node is at most 1
 - if node in graph is memory location, then
 - in-degree + out-degree is at most 1, or
 - edge on node is self-loop
- These graphs are “Spartan Transfer Graphs” [PP10]

Sequentializing Parallel Copies

- Each connected component forms
 - A Cycle
(Then, all nodes are registers)
 - A Path
(Then 1st may be memory store and/or last node may be memory load)
- Can implement as sequential code:
 - cycles use register swap
 - Paths use moves (mov, ld, st as appropriate)

Parallel Copies

- $(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n) \leftarrow (\mathbf{x}_{01}, \mathbf{x}_{11}, \dots, \mathbf{x}_{n1})$
- Each parallel copy forms a “location transfer graph” [Pereira&Palsberg 2010]
 - edges in graph are the pairwise copies that need to be performed
- If we spilled correctly (e.g., all Φ -related variables are spilled to same slot), then LTG is either cycle or path
 - If cycle, only registers involved
 - If path and memory involved, then
 - 1st copy may be store and last copy may be load


Example LTG to code

$$\begin{bmatrix} r_a \\ r_b \\ m_c \\ r_d \\ r_e \end{bmatrix} = \Phi \begin{bmatrix} \dots & r_b & \dots \\ \dots & m_f & \dots \\ \dots & r_g & \dots \\ \dots & r_e & \dots \\ \dots & r_d & \dots \end{bmatrix}$$


Creates LTG with 3 connected components



```
mov  r_a ← r_b  
ld   r_b ← m_f
```



```
st   m_c ← r_g
```




```
xchg r_d ↔ r_e
```


Putting it all together

- Critical Edge Splitting
- Convert back to Conventional-SSA (CSSA)
- Register Allocation
 - Build interference graph
 - pre-spilling
 - coloring
 - coalescing
- Deconstruct SSA
 - put parallel-copies in predecessors
 - Eliminate parallel copies


Example 1



Copy Folded

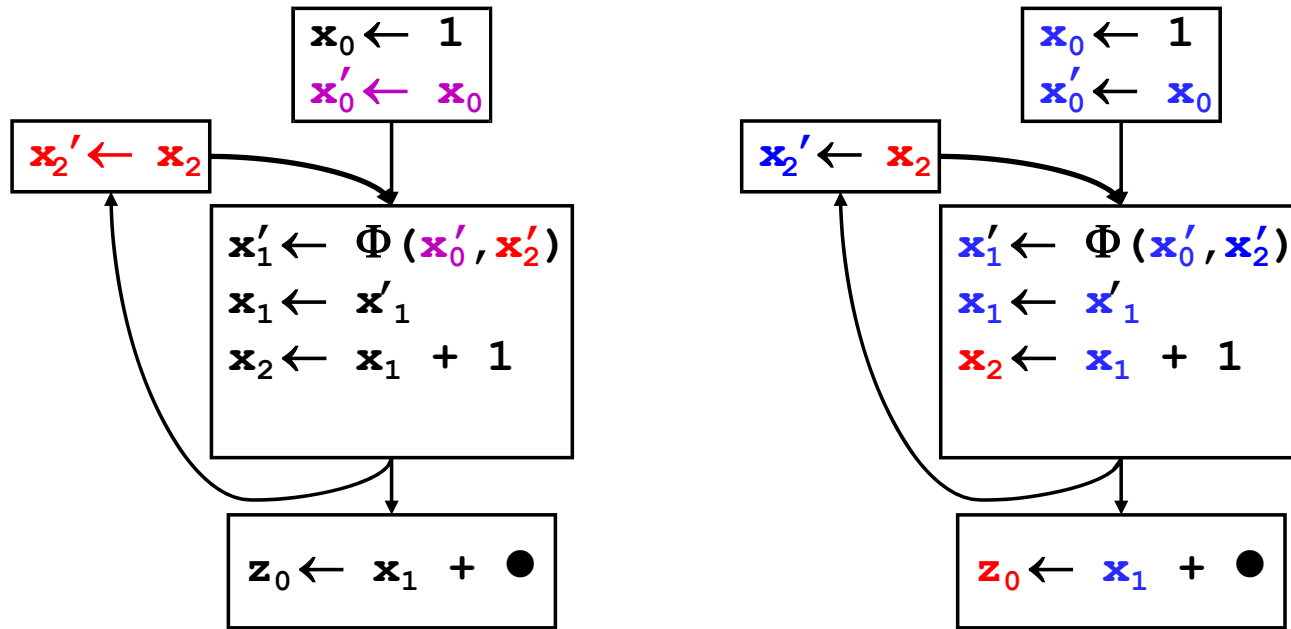


split critical edge



Convert to CSSA

Example 1




Convert to CSSA


register allocation

Done, since $x \leftarrow \Phi(x, x)$ can simply be eliminated.

Well, we can clean up




Convert to CSSA



register allocation

Example 2



Some Fine Tuning

- We added LOTS and LOTS of copies.
- Can reduce added copies when
 - creating CSSA
 - Introducing parallel copies
- Can rely on coalescing, but also ...

Reducing Copies Going to CSSA

- Only need to introduce copies if there is interference!
- As building interference graph mark nodes which are Φ -related. If edge between them, introduce copies and update interference graph.
- Can do even better if also do liveness checking, see [Sreedhar et al, 1999]

Reducing Stores for Spilling

- Every path from LTG that ends in a memory slot will produce a store.
E.g., $a \rightarrow r_1 \rightarrow r_2 \dots r_x \rightarrow m$ will create **st $r_x \rightarrow m$** at the end.
- But, only needs to be done once, e.g., at point of definition.
- So, eliminate store and change register allocator to insert store at definition point
- Similar elimination of loads possible. See [Pereira&Palsberg 2010]

Coalescing

- Coalescing becomes even more important.
- Perform before SSA deconstruction (focus on Φ -related variables)
- (See [Boissinot et.al. 2009])