





But, There are limits to what optimization can do. Indeed, optimization may on rare occasions make a program bigger or slower than it was before.

Moreover, optimization is no substitute for good program design, or especially for intelligent algorithm selection.

Chapter 6 -- Optimization





### 1.1 Basic Blocks

- Some optimization can be done within individual statements, but the most fruitful source of deadwood is in redundancies over sequences of statements.
- For this reason, most machine-independent optimization takes as its starting point the basic block.

Chapter 6 -- Optimization





- This means that there are no branches into or out of the middle of the block
- ♦ and that the entire program can be represented as a set of basic blocks interconnected by branches.
  - + The branches indeed define the boundaries of the blocks.

Chapter 6 -- Optimization



















 In the course of the ensuing optimizations, instructions within a basic block may be modified or deleted.
 Hence it is a good idea to make the gotos refer

- Hence it is a good idea to make the gotos refer to block numbers instead of quadruple numbers, in case the target quad gets deleted.
- You may have noticed that this was done in the breakdown of the previous basic blocks.

Chapter 6 -- Optimization

13

14

At the moment we are interested in the individual blocks; but when we come to global optimization, we will also be concerned with the flow of control and of data between the blocks.

■ The gotos in the code can be pictured as creating paths between the blocks; when we show these paths, we get a flow graph.

Chapter 6 -- Optimization

■ The flow graph for the basic blocks we just created is shown on the next slide.





## 1.2 Local Optimization

- We can do a significant variety of enhancements within a basic block.
- Some of these optimizations can be done globally as well as locally; I am listing them here because then *can* be done locally.
- Details about their implementation must await the development of theory later in the chapter.

Chapter 6 -- Optimization









#### Copy Propagation.

◆ Copy instructions are 3AC instructions of the form

**◆** x := y;

- ♦ When instructions of this sort occur, followed later by a use of x, we can frequently substitute y for x (provided neither has changed between the copy and the use).
- This is another optimization that can be done locally with the sue of a DAG (or globally)

 If x is a temporary, we have now created useless code. We discuss its removal in 6.1.4
 Chapter 6 - Optimization











## 1.3 Loop Optimization

- Most programs spend the bulk of their execution time in loops, and especially in the innermost levels of nested loops.
- There is a well-known saying that 10 percent of the code accounts for 90 percent of the execution time; this 10 percent is usually found in loops.
- Hence loop optimization is of particular importance.



<ul> <li>There are other operations that could be moved that are not reachable by the programmer.</li> <li>for i:= 1 to 30 do</li> <li>for j:= 1 to 30 do</li> <li>for k:= 1 to 30 do</li> <li>for k:= 1 to 30 do</li> <li>To find the address of [i i k] we must compute:</li> </ul>	1
• ro find the address of $[i,j,k]$ we must compute.	
• offset(i, j, k) = $e(k-l_3) + offset(i, j)$ • offset(i, j) = $e(j-l_2)(u_3+l-l_3) + offset(i)$	
• offset(i) = $e(i-l_1)(u_2+1-l_2)(u_3+1-l_3)$	
<ul> <li>where u<sub>i</sub> and l<sub>i</sub> are the upper and lower bounds of each dimension.</li> </ul>	
<ul> <li>Note: offset(i) and offset(i, j) are invariants to the inner loops.</li> </ul>	
Chapter 6 Optimization	27



<ul> <li>T1 := (u3 + 1 - l3);</li> <li>T2 := u2 + 1 - l2);</li> <li>for i := 1 to 30 do</li> <li>begin</li> <li>T3 := e*(l-l1)*T1*T2 {T3 is offset(i)}</li> <li>for(j:= 1 to 30 do</li> <li>begin</li> <li>T4 := T3 + e*(j-l2)*T1 {T4 is offset(i, j)}</li> <li>for k:= 1 to 30 do</li> <li>begin</li> <li>T5 := T4 + e*(k-l3);</li> <li>x[T5] := y[T5] + z[T5]</li> <li>end;</li> </ul>	
<ul> <li>♦ end;</li> <li>and;</li> </ul>	
<ul> <li>end;</li> </ul>	
Chapter 6 Optimization	28



























### 2. DAGs Again

- We saw how to generate DAGs under the control of the parser in Chapter 5.
- If we generate 3AC directly under the parse, then we may have to construct a DAG from scratch during the optimization phase in order to detect and remove common subexpressions.











• Constructing a DAG has to be done with some care, because we must be sure that two instances of an expression really refer to the same thing.

• We can cope with most problems by adding a flag to the record describing each node, indicating whether anything it refers to has been changed. This is sometimes known as a kill bit.

• If our optimizer includes data flow analysis between procedures, then we will know whether any variables are changed by the call and can avoid killing nodes unnecessarily.

Chapter 6 -- Optimization

44

45

## 2.2 Generating 3AC from DAGs

- The DAG, once we have it, tells us what the minimal code must be, and it also shows the dependencies among the various calculations.
- In generating 3AC from the DAG, we work from the leaves upward. Traversing the DAG in order of its dependencies is equivalent to topological sorting.

Chapter 6 -- Optimization

s



In the example given in the text, which used the DAG shown a few slides ago, they cut the size of the block from thirteen 3AC statements to eight, and reduced the number of temporaries from ten to five.
When there are no repeated subexpressions in a basic block, the DAG opens out into a tree.
In this case you can traverse it in preorder and get similar results, but there is a better way.

**3. Data-Flow Analysis**As soon as we expand the scope of optimization beyond the basic block, we must be able to trace how values make their way through a program.
This process is called data-flow analysis.















## 3.2 Data-Flow Equations

- To optimize globally, we have to be able to track the movement of information between basic blocks.
- We must, in fact, consider all the blocks in the procedure because a computed value may be used in some remote block, not necessarily the successor of the block in which it was computed; thus we have to consider the procedure as a whole.

54



- These equations are the *data-flow equations*.
  - ◆ Data-flow analysis considers
    - what happens to variables or expression values within a block, that is, whether they are updated
    - how variables or values reach the beginning of a block

56

- where they go from the end of a block
  - Chapter 6 -- Optimization

◆ Data Flow Analysis takes four different forms, depending on whether we work forward or backward through the flow graph and on whether we may consider any path to or from a block or must consider all paths.

- Each of these forms has its own set of data-flow equations.
  - We need an equation relating the conditions at a block boundary to those in adjacent blocks,

Chapter 6 -- Optimization

 and we need an equation expressing how conditions change within the block.

• Fortunately, the equations for all four forms have the same general appearance, with minor variations, and the methods for solving them run along the same general lines.

- We also need some sort of initial conditions
  - for the beginning of the flow graph
    - (if our analysis is forward)
  - or for the end of the flow graph
    - (if our analysis is backward)
- ♦ We now look at the examples of the four types, develop equations, and then see how they are solved.
  Chapter 6 - Optimization 57



























In order to guarantee that optimization does not change what the program computes, however, we want a conservative solution.

- At issue here is a Type I versus Type II error.
  - We would rather the optimizer miss a possible optimization rather than mess up the program by doing something wrong.

69

### 4. Optimization Techniques Revisited

■ We will now return to the nonlocal optimizations introduced earlier and reconsider them in the light of the information we obtain from data-flow analysis.

Chapter 6 -- Optimization

70

72



### ■ Loop Invariants ◆ A 3AC statement is a loop invariant if its operands + Are constant, or + Are defined outside the loop, or Are defined by some other invariant in the same loop. • To identify loop invariants, you need only go through the body of the loop and apply these tests to every statement. + You can find where the operands are defined by their du-chains. + It may be necessary to go through the loop repeatedly until no more invariant statements can be found.

<ul> <li>Once the invariants have been found, we can consider moving them.</li> <li>Not every invariant can be moved.</li> <li>For Example: <ul> <li>begin</li> <li>3:=5;</li> <li>-:-:-</li> <li>-:-:</li> <li>-:</li> <li>-:</li> </ul> </li> <li>A:=6; <ul> <li>-:</li> <li>-:</li> <li>These are invariants by our example, but clearly they can't be moved; otherwise the value of a won't be changed at the proper point in the loop.</li> <li>So for a invariant 2:</li></ul></li></ul>
<ul> <li>So, for an invariant a:= b+c to be movable, it must be the only definition of a in the loop.</li> </ul>
Chapter 6 Optimization 73





• It is also important to move the statements in
such a way that they appear in the same order in
the preheader as they did in the body of the
loop.
◆ If the loop contains
-

**∗** x:= 5;

**\*** ...

◆ y:- x+1;

Then clearly the assignment to X must still precede the assignment to y after these statements have been moved to the preheader.

Chapter 6 -- Optimization











<ul> <li>Interprocedural data-flow analysis.</li> <li>It is desirable to be able to trace the movement of data between a subprogram and the program that calls it.</li> <li>This is called interprocedural data-flow analysis.</li> </ul>
<ul> <li>If we can do this, then when we are constructing a DAG for a block that includes a function call, we will know which if any nodes in the DAG must be killed.</li> </ul>
<ul> <li>and when we are optimizing a loop, we will know whether it is safe to move a function call outside the loop.</li> </ul>
Chapter 6 Optimization 81







- We will look at some examples of register optimization here, although detailed techniques for selection registers will have to await the discussion of code generation in Chapter 7.
- Examples of instruction use are harder to find, since machine architectures vary so widely.

Chapter 6 -- Optimization



















- known as peephole optimization.
- We will now look at some of the redundancies found and removed by peephole optimization.

Chapter 6 -- Optimization

91





Chapter 6 -- Optimization





■ If we know all these techniques, how many of them are we going to use?

- Presumably the ultimate optimizing compiler is going to se every single one of them;
- But in practice we may settle for much less, particularly because some techniques may be effective only rarely.
- One obvious approach is to restrict our optimizations to those in which the ratio of payoff to cost is highest.

Chapter 6 -- Optimization





- Global optimizations like interprocedural data-flow analysis may be difficult or impossible unless the entire source program is in memory.
- If the amount of memory is limited, we may omit this optimization.
- Another possibility is to write two versions of the compiler and market the more elaborate one as an "optimizing compiler"
  - Even the less elaborate version will usually include some minimal amount of optimization, however.

Chapter 6 -- Optimization

97



Many commercial compilers attempt nothing more than this minimum.