

EECS 583 – Class 2

Control Flow Analysis

LLVM Introduction

University of Michigan

September 8, 2014

Summer is over.

Time to officially
remember what
day of the week it is.

CORNELLIUS



Mid-Autumn Festival

September 8th, 2014

(The 15th day of the 8th month on Traditional Chinese Calendar)




Announcements & Reading Material

- ❖ HW 1 out today, due Friday, Sept 22 (2 wks)
 - » This homework is not hard, but takes lots of time to figure LLVM out, so start soon!!
 - » Part I: get “hello world” application working
 - » Part II: Run profilers (control & memory dep), collect some stats
- ❖ Reading
 - » Today’s class
 - ÿ Ch 9.6 from Dragon book
 - ÿ Or Ch 7.1, 7.3, 7.4 from Muchnick
 - ÿ “Trace Selection for Compiling Large C Applications to Microcode”, Chang and Hwu, MICRO-21, 1988.
 - » Next class
 - ÿ “The Superblock: An Effective Technique for VLIW and Superscalar Compilation”, Hwu et al., Journal of Supercomputing, 1993


From last time: Control Flow Graph (CFG)

- ❖ Defn Control Flow Graph – Directed graph, $G = (V, E)$ where each vertex V is a basic block and there is an edge E , v_1 (BB1) \rightarrow v_2 (BB2) if BB2 can immediately follow BB1 in some execution sequence
 - » A BB has an edge to all blocks it can branch to
 - » Standard representation used by many compilers
 - » Often have 2 pseudo vertices
 - entry node
 - exit node



Weighted CFG



- ❖ Profiling – Run the application on 1 or more sample inputs, record some behavior
 - » Control flow profiling
 - edge profile
 - block profile
 - » Path profiling
 - » Cache profiling
 - » Memory dependence profiling
- ❖ Annotate control flow profile onto a CFG → weighted CFG
- ❖ Optimize more effectively with profile info!!
 - » Optimize for the common case
 - » Make educated guess



Property of CFGs: Dominator (DOM)


- ❖ Defn: Dominator – Given a CFG($V, E, \text{Entry}, \text{Exit}$), a node x dominates a node y , if every path from the Entry block to y contains x
- ❖ 3 properties of dominators
 - » Each BB dominates itself
 - » If x dominates y , and y dominates z , then x dominates z
 - » If x dominates z and y dominates z , then either x dominates y or y dominates x
- ❖ Intuition
 - » Given some BB, which blocks are guaranteed to have executed prior to executing the BB

Dominator Examples




Dominator Analysis

- ❖ Compute $\text{dom}(\text{BB}_i)$ = set of BBs that dominate BB_i
- ❖ Initialization
 - » $\text{Dom}(\text{entry}) = \text{entry}$
 - » $\text{Dom}(\text{everything else}) = \text{all nodes}$
- ❖ Iterative computation
 - » while change, do
 - change = false
 - for each BB (except the entry BB)
 - ♦ $\text{tmp}(\text{BB}) = \text{BB} + \{\text{intersect of Dom of all predecessor BB's}\}$
 - ♦ if ($\text{tmp}(\text{BB}) \neq \text{dom}(\text{BB})$)
 - $\text{dom}(\text{BB}) = \text{tmp}(\text{BB})$
 - change = true




Immediate Dominator

- ❖ Defn: Immediate dominator (idom) – Each node n has a unique immediate dominator m that is the **last dominator** of n on any path from the initial node to n
 - » Closest node that dominates




Dominator Tree

First BB is the root node, each node dominates all of its descendants




BB	DOM	BB	DOM
1	1	5	1,4,5
2	1,2	6	1,4,6
3	1,3	7	1,4,7
4	1,4		



Dom tree

Class Problem



Draw the dominator tree for the following CFG



Post Dominator (PDOM)

- ❖ Reverse of dominator
- ❖ Defn: Post Dominator –
Given a CFG($V, E, \text{Entry}, \text{Exit}$), a node x post dominates a node y , if every path from y to the Exit contains x
- ❖ Intuition
 - » Given some BB, which blocks are guaranteed to have executed after executing the BB
- ❖ $\text{pdom}(\text{BB}_i) = \text{set of BBs that post dominate BB}_i$
- ❖ Initialization
 - » $\text{Pdom}(\text{exit}) = \text{exit}$
 - » $\text{Pdom}(\text{everything else}) = \text{all nodes}$
- ❖ Iterative computation
 - » while change, do
 - $\text{change} = \text{false}$
 - for each BB (except the exit BB)
 - ♦ $\text{tmp}(\text{BB}) = \text{BB} + \{\text{intersect of pdom of all successor BB's}\}$
 - ♦ if ($\text{tmp}(\text{BB}) \neq \text{pdom}(\text{BB})$)
 $\text{pdom}(\text{BB}) = \text{tmp}(\text{BB})$
 $\text{change} = \text{true}$


Post Dominator Examples



Immediate Post Dominator


❖ Defn: Immediate post dominator (ipdom) –
Each node n has a unique immediate post dominator m that is the first post dominator of n on any path from n to the Exit

- » Closest node that post dominates
- » First breadth-first successor that post dominates a node



Why Do We Care About Dominators?

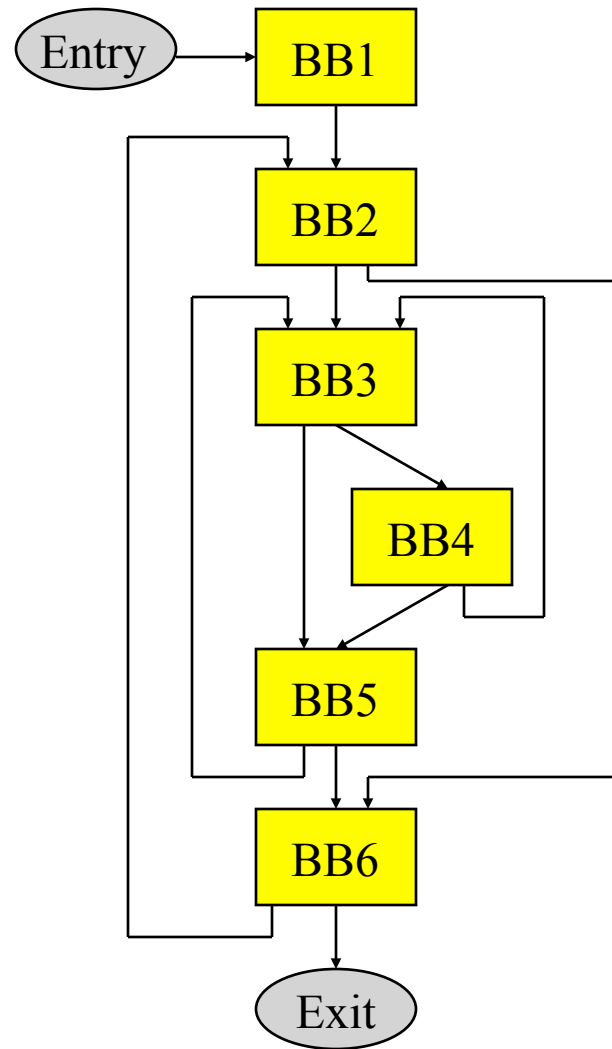
- ❖ Loop detection – next subject
- ❖ Dominator
 - » Guaranteed to execute before
 - » Redundant computation – an op is redundant if it is computed in a dominating BB
 - » Most global optimizations use dominance info
- ❖ Post dominator
 - » Guaranteed to execute after
 - » Make a guess (ie 2 pointers do not point to the same locn)
 - » Check they really do not point to one another in the post dominating BB



Natural Loops

- ❖ Cycle suitable for optimization
 - » Discuss optimizations later
- ❖ 2 properties
 - » Single entry point called the header
 - Header dominates all blocks in the loop
 - » Must be one way to iterate the loop (ie at least 1 path back to the header from within the loop) called a backedge
- ❖ Backedge detection
 - » Edge, $x \rightarrow y$ where the target (y) dominates the source (x)


Backedge Example



Loop Detection

- ❖ Identify all backedges using Dom info
- ❖ Each backedge ($x \rightarrow y$) defines a loop
 - » Loop header is the backedge target (y)
 - » Loop BB – basic blocks that comprise the loop
 - Y All predecessor blocks of x for which control can reach x without going through y are in the loop + y
- ❖ Merge loops with the same header
 - » I.e., a loop with 2 continues
 - » $\text{LoopBackedge} = \text{LoopBackedge1} + \text{LoopBackedge2}$
 - » $\text{LoopBB} = \text{LoopBB1} + \text{LoopBB2}$
- ❖ Important property
 - » Header dominates all LoopBB


Loop Detection Example



Important Parts of a Loop

- ❖ Header, LoopBB
- ❖ Backedges, BackedgeBB
- ❖ Exitedges, ExitBB
 - » For each LoopBB, examine each outgoing edge
 - » If the edge is to a BB not in LoopBB, then its an exit
- ❖ Preheader (Preloop)
 - » New block before the header (falls through to header)
 - » Whenever you invoke the loop, preheader executed
 - » Whenever you iterate the loop, preheader NOT executed
 - » All edges entering header
 - Backedges – no change
 - All others, retarget to preheader
- ❖ Postheader (Postloop) - analogous

Preheaders for each Loop

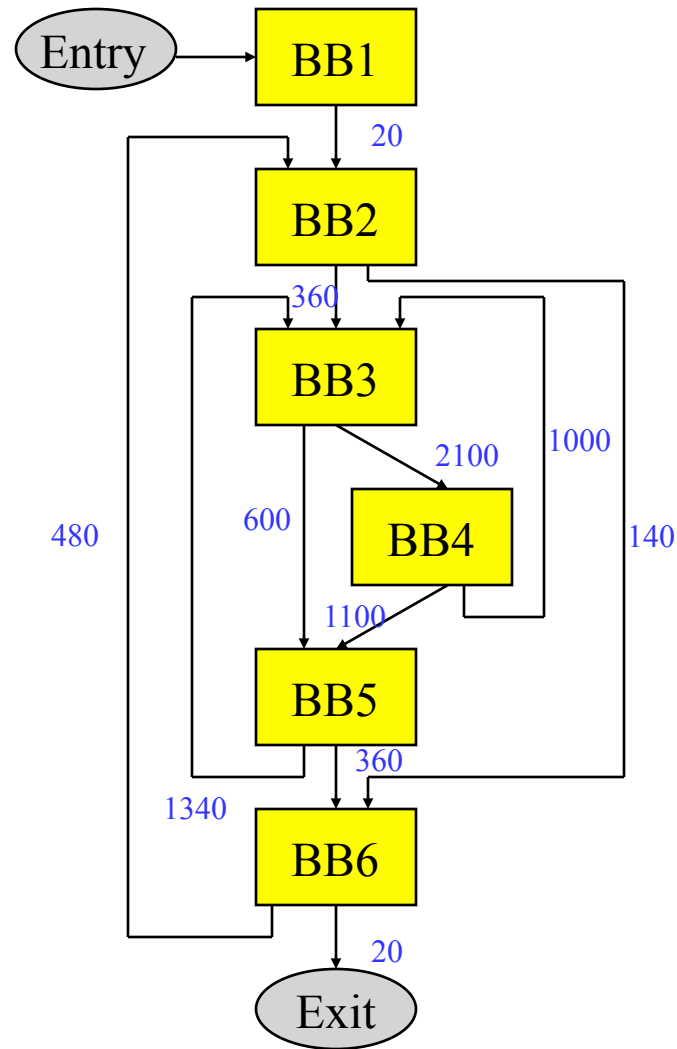


Characteristics of a Loop

- ❖ Nesting (generally within a procedure scope)
 - » Inner loop – Loop with no loops contained within it
 - » Outer loop – Loop contained within no other loops
 - » Nesting depth
 - $\text{depth}(\text{outer loop}) = 1$
 - $\text{depth} = \text{depth}(\text{parent or containing loop}) + 1$
- ❖ Trip count (average trip count)
 - » How many times (on average) does the loop iterate
 - » `for (I=0; I<100; I++)` → trip count = 100
 - » With profile info:
 - $\text{Ave trip count} = \text{weight}(\text{header}) / \text{weight}(\text{preheader})$


Trip Count Calculation Example

Calculate the trip counts for all the loops in the graph



Reducible Flow Graphs

- ❖ A flow graph is reducible if and only if we can partition the edges into 2 disjoint groups often called forward and back edges with the following properties
 - » The forward edges form an acyclic graph in which every node can be reached from the Entry
 - » The back edges consist only of edges whose destinations dominate their sources
- ❖ More simply – Take a CFG, remove all the backedges ($x \rightarrow y$ where y dominates x), you should have a connected, acyclic graph



Non-reducible!

Regions

- ❖ Region: A collection of operations that are treated as a single unit by the compiler
 - » Examples
 - Basic block
 - Procedure
 - Body of a loop
 - » Properties
 - Connected subgraph of operations
 - Control flow is the key parameter that defines regions
 - Hierarchically organized
- ❖ Problem
 - » Basic blocks are too small (3-5 operations)
 - Hard to extract sufficient parallelism
 - » Procedure control flow too complex for many compiler xforms
 - Plus only parts of a procedure are important (90/10 rule)

Regions (2)

❖ Want

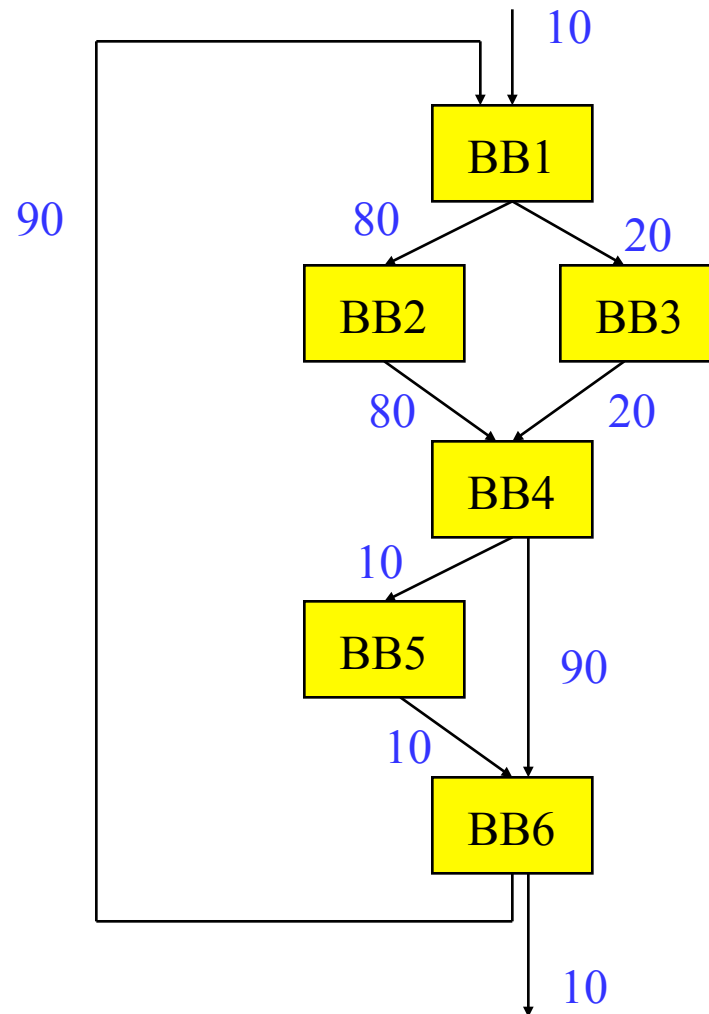
- » Intermediate sized regions with simple control flow
- » Bigger basic blocks would be ideal !!
- » Separate important code from less important
- » Optimize frequently executed code at the expense of the rest

❖ Solution

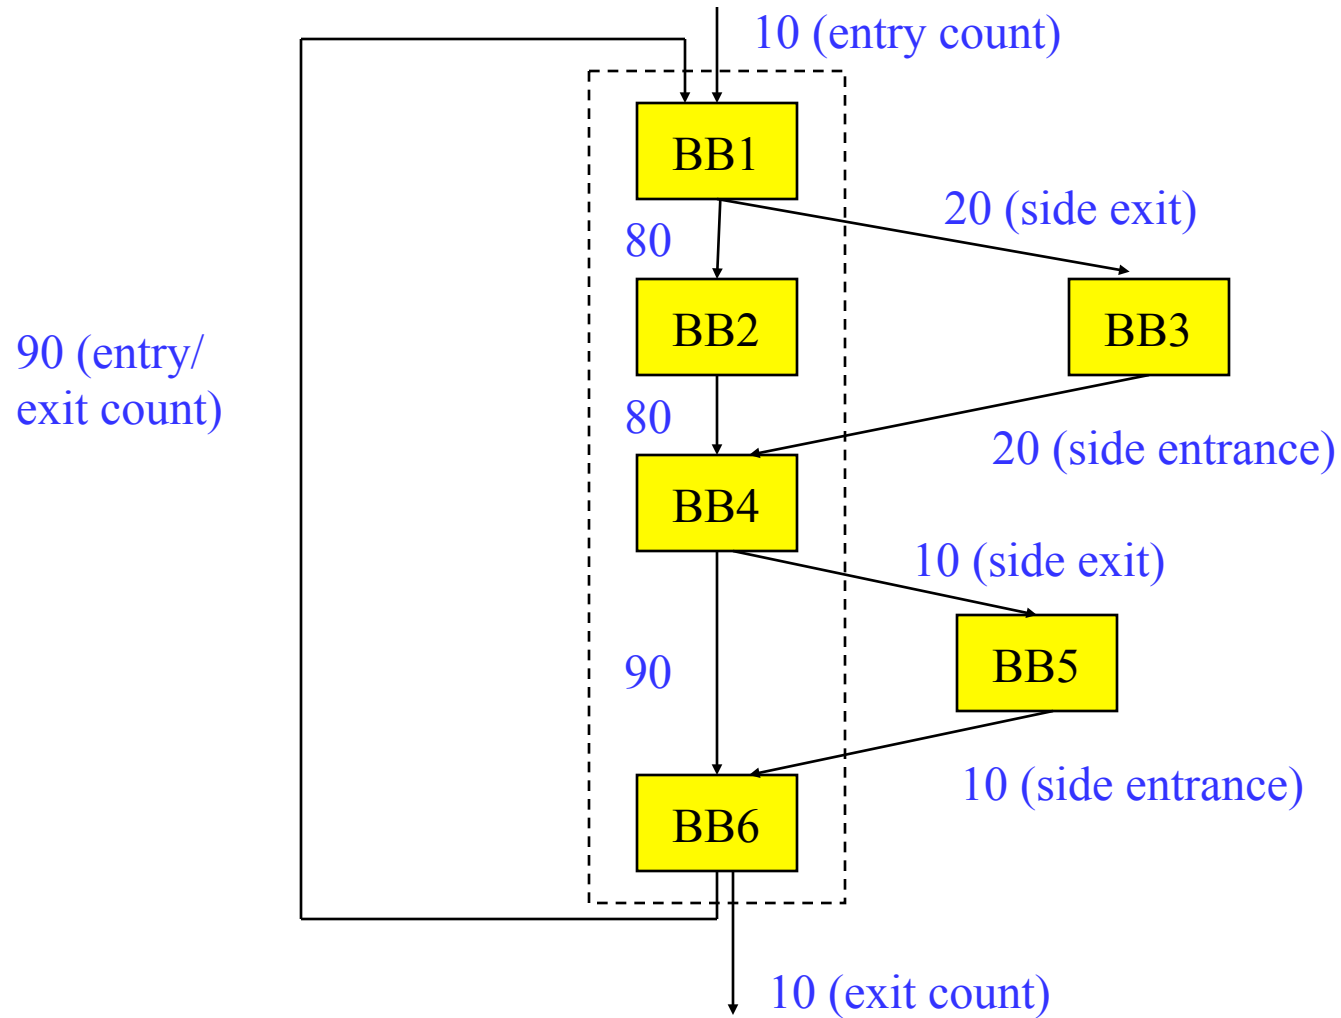
- » Define new region types that consist of multiple BBs
- » Profile information used in the identification
- » Sequential control flow (sorta)
- » Pretend the regions are basic blocks

Region Type 1 - Trace

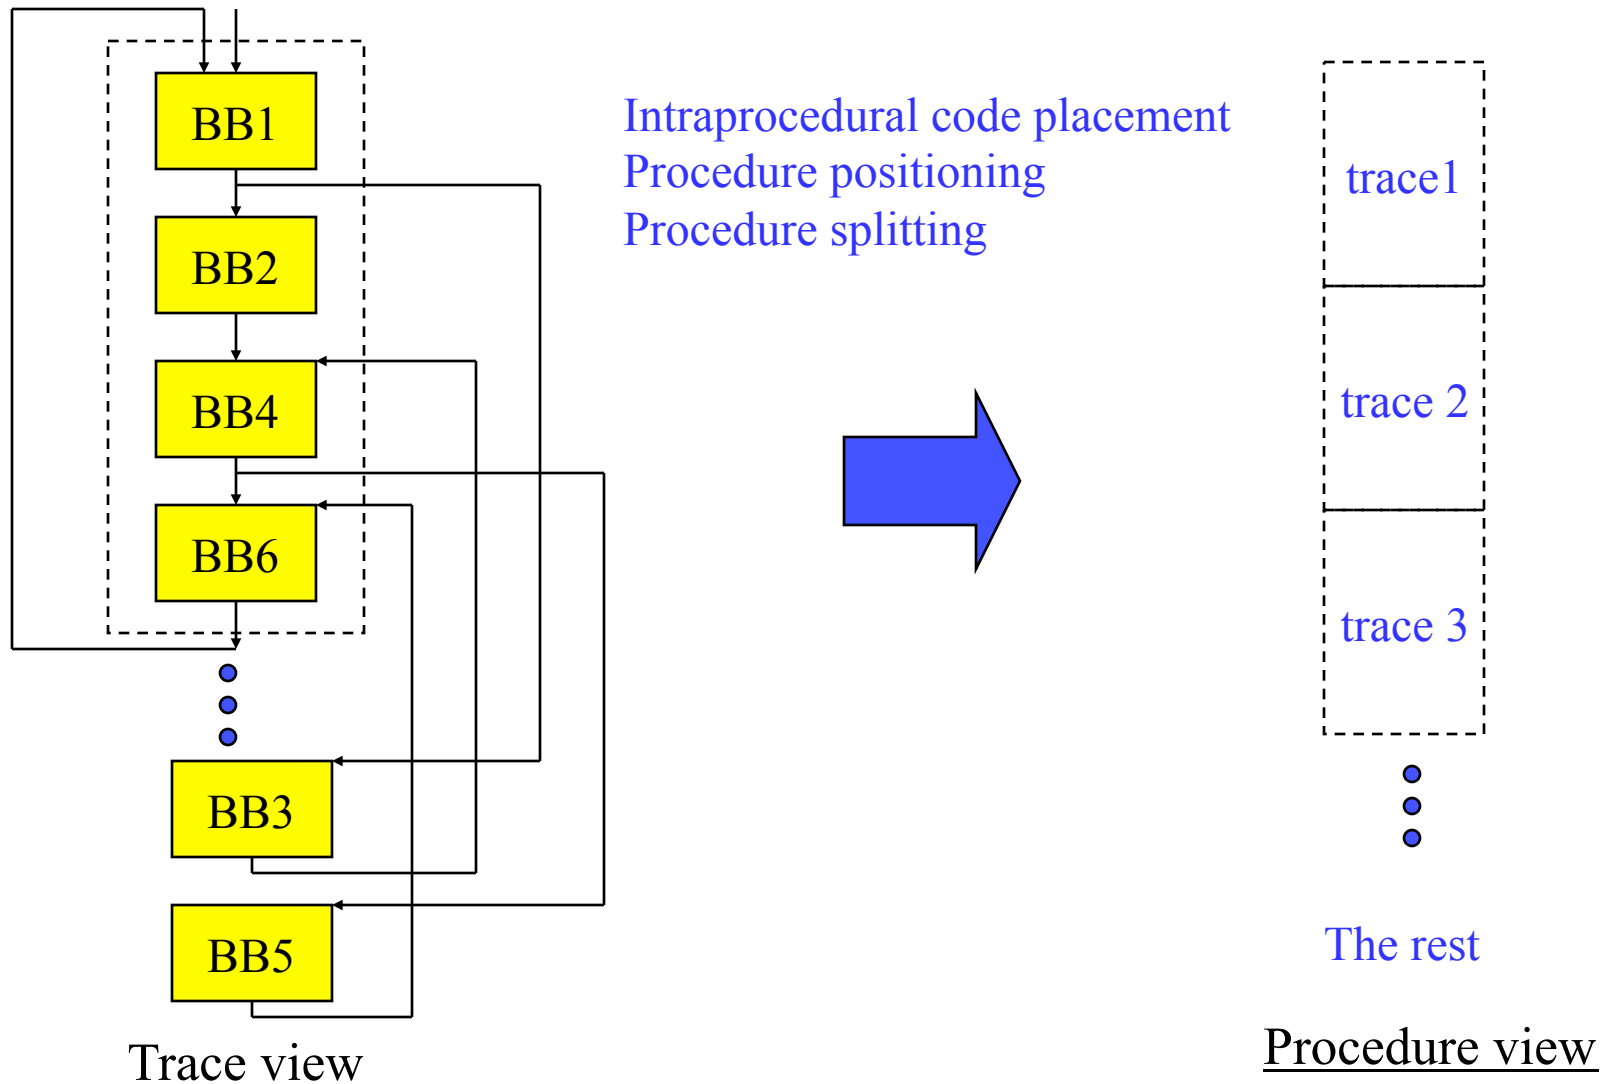
- ❖ Trace - Linear collection of basic blocks that tend to execute in sequence
 - » “Likely control flow path”
 - » Acyclic (outer backedge ok)
- ❖ Side entrance – branch into the middle of a trace
- ❖ Side exit – branch out of the middle of a trace
- ❖ Compilation strategy
 - » Compile assuming path occurs 100% of the time
 - » Patch up side entrances and exits afterwards
- ❖ Motivated by scheduling (i.e., trace scheduling)



Linearizing a Trace




Intelligent Trace Layout for Icache Performance



Issues With Selecting Traces

- ❖ Acyclic
 - » Cannot go past a backedge
- ❖ Trace length
 - » Longer = better ?
 - » Not always !
- ❖ On-trace / off-trace transitions
 - » Maximize on-trace
 - » Minimize off-trace
 - » Compile assuming on-trace is 100% (ie single BB)
 - » Penalty for off-trace
- ❖ Tradeoff (heuristic)
 - » Length
 - » Likelihood remain within the trace



Trace Selection Algorithm

```
i = 0;
mark all BBs unvisited
while (there are unvisited nodes) do
    seed = unvisited BB with largest execution freq
    trace[i] += seed
    mark seed visited
    current = seed
    /* Grow trace forward */
    while (1) do
        next = best_successor_of(current)
        if (next == 0) then break
        trace[i] += next
        mark next visited
        current = next
    endwhile
    /* Grow trace backward analogously */
    i++
endwhile
```

Best Successor/Predecessor

- ❖ Node weight vs edge weight
 - » edge more accurate
- ❖ THRESHOLD
 - » controls off-trace probability
 - » 60-70% found best
- ❖ Notes on this algorithm
 - » BB only allowed in 1 trace
 - » Cumulative probability ignored
 - » Min weight for seed to be chose (ie executed 100 times)

```
best_successor_of(BB)
  e = control flow edge with highest
    probability leaving BB
  if (e is a backedge) then
    return 0
  endif
  if (probability(e) <= THRESHOLD) then
    return 0
  endif
  d = destination of e
  if (d is visited) then
    return 0
  endif
  return d
end procedure
```

Class Problems

Find the traces. Assume a threshold probability of 60%.

