

Understanding the Backward Slices of Performance Degrading Instructions

Craig Zilles and Guri Sohi

University of Wisconsin - Madison

International Symposium on Computer Architecture

June, 2000


Motivation

PROCESSORS ACHIEVE ONLY A FRACTION OF PEAK PERFORMANCE ON MANY PROGRAMS

- *Performance Degrading Events (PDE)*
 - **branch mispredictions**
 - **cache misses**

LARGER CACHES AND PREDICTORS

- *handle easy cases*
- *concentrates PDE's to a fraction of "problem" static instructions*
 - **UNCORRELATED (DATA-DEPENDENT) BRANCHES**
 - **HASH TABLE LOOKUPS AND POINTER CHASING**




Motivation, cont.

PROGRAM BEHAVIOR IS DETERMINISTIC →

build predictors which use the program!

Pre-execution



Pre-executed sub-program feeds prediction for branch fetched by the main thread

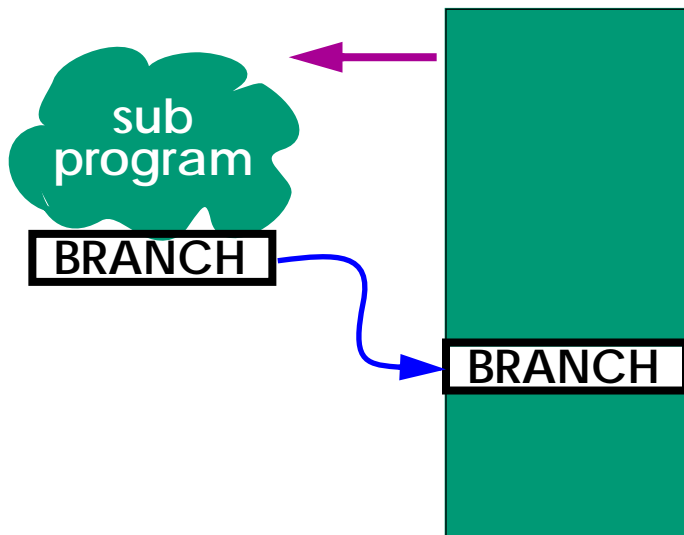
Pre-fetch memory similarly

Only pre-execute instructions which defy normal predictors

AVOID MISPREDICTION

Motivation, cont.

THE EFFECTIVENESS OF PRE-EXECUTION IS DETERMINED BY THE SUB-PROGRAM



- Sub-program must **enable faster execution** of the problem instruction
- Sub-program **size determines "overhead"**

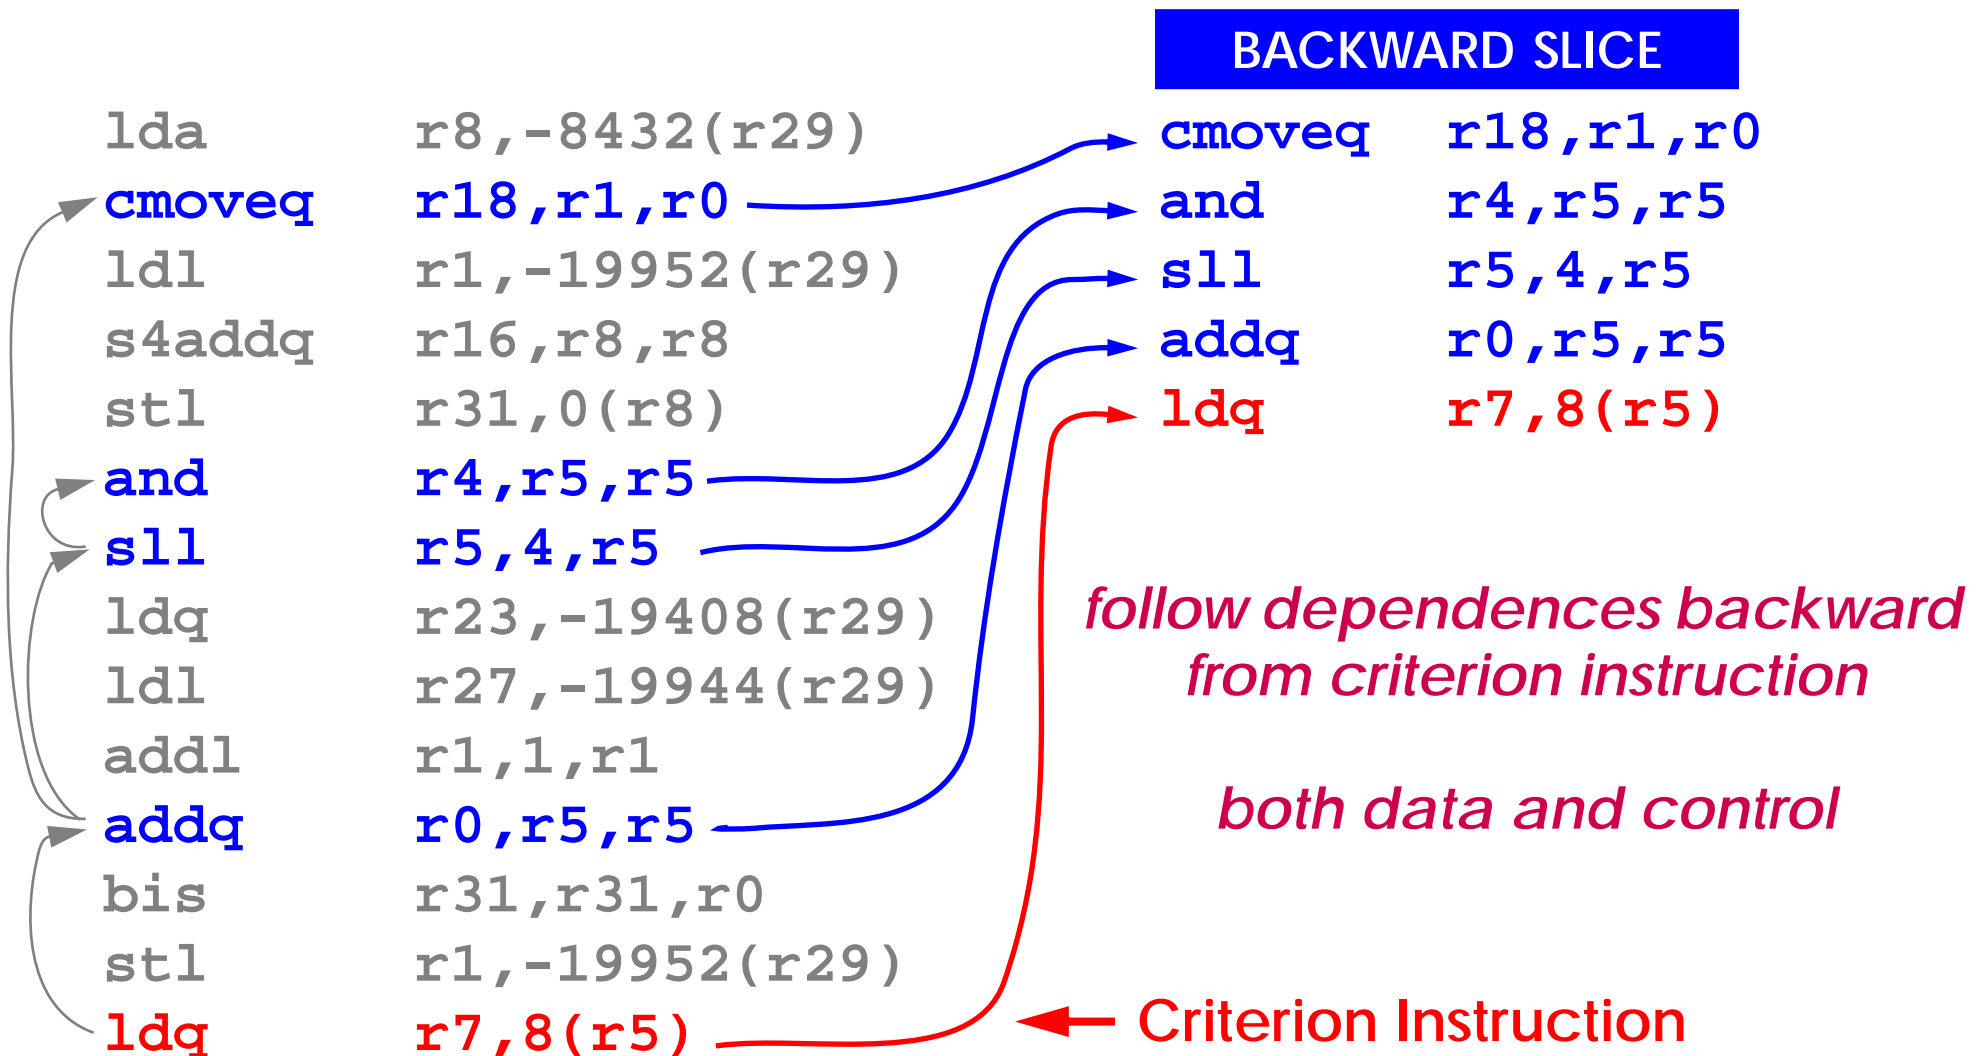
How can we build sub-programs to minimize their size while maintaining accuracy

Overview

- *MOTIVATION*
- PROGRAM SLICING
- EXPERIMENT OVERVIEW
- METHODOLOGY
- CONSERVATIVE SLICES
- EXAMPLE SPECULATIVE OPTIMIZATIONS
- ADDITIONAL OBSERVATIONS
- CONCLUSIONS

Program Slicing

THE SUB-PROGRAM INCLUDES ONLY THE SUBSET OF INSTRUCTIONS WHICH CAN INFLUENCE THE PROBLEM INSTRUCTION.



Experiment Overview

INITIAL CHARACTERIZATION OF SLICES

Categorization of instructions in the slice


- *largest contributors:*
 - **control and memory dependence resolution**
 - **NOT dataflow**

Exploiting well-known phenomena to reduce slice size

- *highly-biased branches*
- *stability of memory dependences*

Methodology

- *SPEC95 integer benchmarks*
 - Alpha architecture, optimized -O4
- *Profiled to identified “problem” static instructions*
 - Frequently caused mispredictions or cache misses
- *Generated ASSEMBLY-LEVEL slices*
 - Looked at the 512 dynamic instructions leading to the criterion
 - Removed NOPS, and SP/GP computation dependences



Conservative Slices

PROGRAMS HAVE **AMBIGUOUS MEMORY DEPENDENCES** AND **COMPLEX CONTROL FLOW** → CONSERVATIVELY CONSTRUCTED SLICES CAN BE LARGE

- *50% of program necessary to compute all store addresses*
- *80% of program necessary to resolve all branches*

SOLUTION: EXPLOIT THE FACT THAT SLICES ONLY PROVIDE HINTS

- construct **speculative slices**
- assume common-case behavior
 - profiling is required to detect the common-case

TWO EXAMPLE OPTIMIZATIONS:


- both targeting ambiguous memory dependences

Speculative Optimizations

A CONSERVATIVE SLICE MUST COMPUTE THE ADDRESS FOR **EVERY STORE WHICH COULD ALIAS WITH A LOAD** IN THE SLICE

THE SET OF MEMORY DEPENDENCES ACTUALLY REALIZED IS A SUB-SET OF THOSE WHICH ARE POSSIBLE

- Profile to identify the store sets
- Only compute store addresses for these stores



Slices built using store sets approximate size of oracle slices


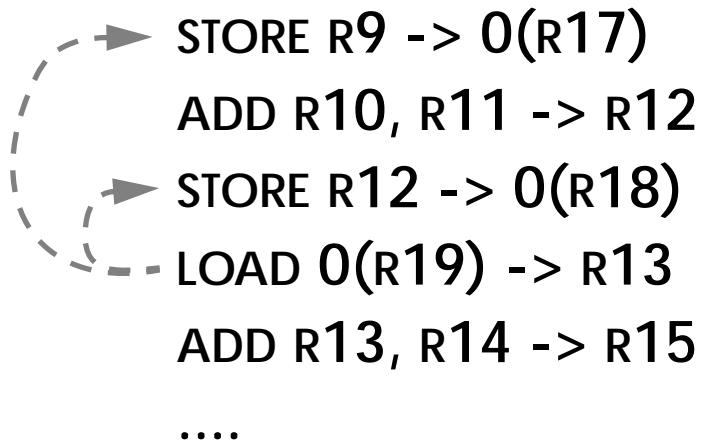
Speculative Optimizations, cont.

WHEN MEMORY DEPENDENCES EXIST, OFTEN THE LOAD COMMUNICATES WITH THE **MOST RECENT** STORE FROM ITS STORE SET.

- much like communication through registers

EXPLOIT TO REDUCE SLICE SIZE:


1. Assume a communication pattern (imprecise transformation)
2. Remove load and store from the slice



Speculative Optimizations, cont.

STATIC LOADS ARE HIGHLY BIASED WITH RESPECT TO THIS BEHAVIOR

- Mis-speculation can be avoided
- Easily profiled to classify the dependences



Removing such load-store pairs can further reduce slice size with little affect on accuracy

Reduced address sub-slice to 1/4 of conservative size

Additional Observations

OFTEN DATA DEPENDENCES ARE CLUSTERED NEAR CRITERION


- possibly influences pre-execution mechanisms

NOT UNCOMMON FOR SLICES TO OVERLAP

- create a single “multi-slice”

SLICES ARE BURSTY

- due to program structure



Additional Observations

FALSE CONTROL DEPENDENCES:

control dependent regions are part of slice, but all paths from the branch contribute to the slice equivalently.

COMMON CASE: CONDITIONAL FUNCTION CALL

```
INT A = ....;
IF (B) {
    FUNCTION();
}
CRITERION → IF (A)
```

```
FUNCTION() {
    SAVE A;
    ....
    RESTORE A;
}
```

OTHER CASE: CODE REPLICATION

Currently refining infrastructure to handle these cases

Summary

PRE-EXECUTION: GENERAL TECHNIQUE FOR HANDLING “PROBLEM” INSTRUCTIONS

- Use the program to predict the program
- Requires small, accurate slices

SPECULATIVE SLICES:

- Exploit common-case behavior to reduce slice size while maintaining accuracy
- Some slices can be reduced to less than 10% of the 512 instructions preceding the criterion while maintaining greater than 95% accuracy
 - Much future work to be done
- Requires sophisticated profile information