

Partial Redundancy Elimination using Lazy Code Motion

Sandeep Dasgupta¹ Tanmay Gangwani²

May 11, 2014

¹Electronic address: sdasgup3@illinois.edu

²Electronic address: gangwan2@illinois.edu

1. Problem Statement & Motivation

Partial Redundancy Elimination (PRE) is a compiler optimization that eliminates expressions that are redundant on some but not necessarily all paths through a program. In this project, we implemented a PRE optimization pass in LLVM and measured results on a variety of applications. We chose PRE because it's a powerful technique that subsumes Common Subexpression Elimination (CSE) and Loop Invariant Code Motion (LICM), and hence has a potential to greatly improve performance.

In the example below, the computation of the expression $(a + b)$ is partially redundant because it is redundant on the path $1 \rightarrow 2 \rightarrow 5$, but not on the path $1 \rightarrow 4 \rightarrow 5$. PRE works by first introducing operations that make the partially redundant expressions fully redundant and then deleting the redundant computations. The computation of $(a + b)$ is added to 4 and then deleted from 5.

```
(1) if (OPAQUE)
(2)   x = a + b;
(3) else
(4)   x = 0;
(5) y=a+b;
```

2. Related Work

Partial Redundancy Elimination

Morel et al. [10] first proposed a bit-vector algorithm for the suppression of partial redundancies. The bi-directionality of the algorithm, however, proved to be computationally challenging. Knoop et al. [9] solved this problem with their Lazy Code Motion (LCM) algorithm. It is composed of uni-directional data flow equations and provides the earliest and latest placement points for operations that should be hoisted. Drechsler et al. [6] present a variant of LCM which they claim to be more useful in practice. Briggs et al. [2] allude to two pre-passes to make PRE more effective - Global Reassociation and Value Numbering.

Value Numbering

Briggs et al. [3] compare and contrast two techniques for value numbering - hash based[4] and partition based[1]. In subsequent work they provide SCC-based Value Numbering [5] which combines the best of the previously mentioned approaches. Cooper et al. [11] show how to incorporate value information in the data flow equations of LCM to eliminate more redundancies.

PRE in LLVM

Since LLVM is a Static Single Assignment (SSA) based representation, algorithms based on identifying expressions which are lexically identical or have the same static value number may fail to capture some redundancies. Kennedy Chow et al. [7] provide a new framework for PRE on a program in SSA form. The present GVN-PRE pass in LLVM appears to be inspired by the work of Thomas et al. [12] which also focuses on SSA.

3. Algorithm Overview

Our algorithm for PRE is a slightly modified version of the iterative bit-vector data flow algorithm by Knoop et al. [9]. It uses four data flow equations to identify for each expression in the program, the optimal evaluation point. The first flow equation calculates down-safe (anticipatable) points for an expression. An expression is said to be down-safe at a point p if computing the expression at p would be useful along any path from p . The second flow equation calculates

up-safe (available) points. An expression is up-safe at a point p if it has been computed on every path from the entry node to p and not killed after the last computation on each path. Using these, the algorithm calculates the *Earliest* property. An expression is said to be *Earliest* at a point p if there doesn't exist an earlier point where the computation of the expression is both down-safe and produces the correct values. Such points are known as *computationally optimal* placement points.

Evaluating the expression at computationally optimal points could negatively impact performance due to increased register pressure. Therefore, the latter half of the LCM algorithm pushes the computation of the expression close to the use of the expression. More specifically, the third flow equation calculates the *Latest* property. An expression is said to be *Latest* at a point p if it is computationally optimal at p , and on every path from p , any later optimal point on the path would be after some use of the expression. Through the fourth and final flow equation, the algorithm determines if it is necessary to allocate a temporary at a point p for the expression. The property is known as *Isolated*. An expression is *Isolated* at a point if it is optimal, and the value of the expression is only used immediately after the point. Therefore, allocation of temporaries at *Isolated* points is avoided.

In summary, the four flow equations provide computationally optimal placement points which require the shortest lifetimes for the temporary variables introduced. In appendix B, we outline all equations.

4. Implementation Details

Value Numbering

Prior research [2] has shown that value numbering can increase opportunities for PRE. LLVM presently has a GVN-PRE pass which exploits this. However, value numbering in GVN-PRE is tightly coupled with the code for removing redundancies, and hence we were not able to use the same for our code. We wrote our own value numbering pass which fed expression value numbers to the PRE stage. It should be noted, however, that we did not implement value numbering from scratch and used an old (now defunct) LLVM pass as a starting point. Most importantly, we augmented the basic value numbering in the following ways -

- Added the notion of leader expression (described below), with associated data structures and functions.
- Functionality to support value-number-based bitvectors rather than expression-name-based bitvectors.
- (Optimization 1) If the expression operator is one of these - AND, OR, CMP::EQ or CMP::NE, and the operands have the same value number, we replace all uses accordingly and then delete the expression.
- (Optimization 2) If all operands of an expression are constants, then we evaluate and propagate constants.
- (Optimization 3) If one operand of an expression is a constant (0 or 1), then we simplify the expression. e.g. $a+0 = a$, $b*1 = b$.
- (Optimization 4) If the incoming expressions to a `Phi` node have the same value number, then the `Phi` node gets that same value number

Reassociation has also been shown to make the code more amenable for PRE. It refers to using associativity, commutativity and distributivity to divide expressions into parts that are constant, loop invariant and variable. We used an already existing LLVM pass (-reassociate) for Global Reassociation. As per our testing, optimizations 2 and 3 (above) are also done by this pass, and hence, we disabled our version for the more robust LLVM version. Optimizations 1 and 4, however, are still our contribution.

Notion Of Leader Expression

The value numbering algorithm computes the RPO solution as outlined in [5]. It goes over the basic blocks in reverse post order and adds new expressions to a hash table based on the already computed value numbers of the operands. We call an expression a ‘leader’ if at the time of computing its value number, the value number doesn’t already exist in the hash table. In other words, out of a potentially large set of expressions that map to a particular value number, the leader expression was the first to be encountered while traversing the function in reverse post order. Leader expressions are vital to our algorithm as they are used to calculate the block local properties of the data flow equations (Appendix A).

Types Of Redundancies

Given two expressions X and Y in the source code, following are the possibilities -

1. X and Y are lexically equivalent, and have the same value numbers
2. X and Y are lexically equivalent, but have different value numbers
3. X and Y are lexically different, but have the same value numbers
4. X and Y are lexically different, and have different value numbers

In the source code, there could be opportunities for redundancy elimination in cases 1, 2 and 3 above. If the source code is converted to an intermediate representation in SSA form then case 2 becomes an impossibility (by guarantees of SSA). Therefore, our algorithm presently handles the cases when X and Y are lexically same/different, but both have the same value number (cases 1 and 3). Driven by this observation, we implement value number based code motion, the details of which are presented below. It should be noted that even though case 2 above is not possible in SSA, the source code redundancies of this type transform into that of type case 4. Figure 1 presents an illustration of the same. **This is not handled in our implementation.**

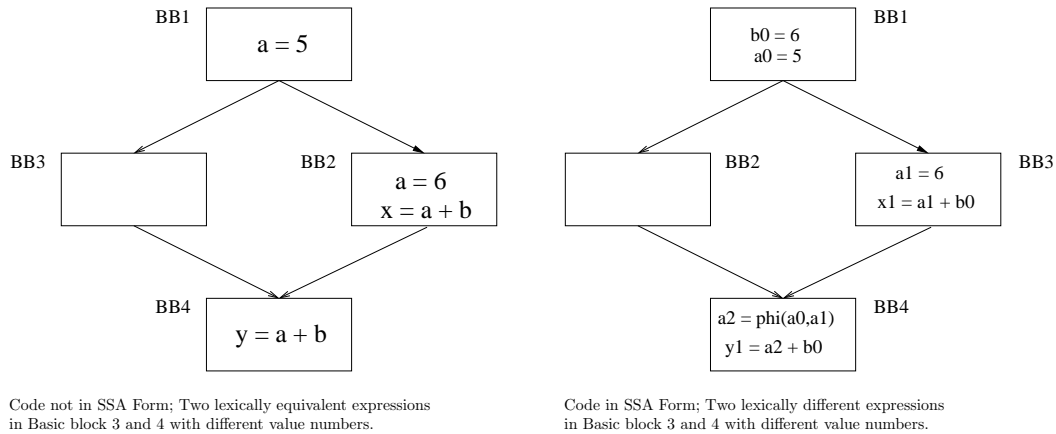


Figure 1

Value-Number driven code motion

We initially implemented the flow equations from the Lazy Code Motion paper [9]. This set included a total of 13 bit vectors for each basic block - 2 for block local properties ANTLOC and TRANSP, and 11 for global properties. These equations, however, could only be applied to single instruction basic blocks. We therefore, derived a new set of equations which are motivated by later work[8] of the same authors. This set of equations apply to maximal basic blocks and entails a total of 19 bit vectors for each basic block in our current implementation - 3 for block local properties ANTLOC, TRANSP, XCOMP and 16 for global properties. Block local properties are defined in appendix A. In appendix C, we include the generalized data flow framework, and show how each PRE equation maps to the framework. We call the algorithm value-number driven because each slot in each of the bit vectors is reserved for a particular value number rather than a particular expression. Also, we make the observation that a large number of expressions in the program only occur once, and are not useful for PRE. Therefore, to further optimize for space and time, we only give bit vector slots to value numbers which have more than one expression linked to them. A downside to this approach is that we could miss opportunities for loop invariant code motion. As a solution, we extend the bit vector to include value numbers which have only a single expression linked to them but only if the expression is inside a loop. Note that we still exclude the cases where the expression is not part of a loop. Figure 2 quantifies the savings we observe using functions from the LLVM multi-source package. In the worst case scenario, the bit-vector width maintained by our algorithm has to be equal to the maximum value-number assigned by the value-numbering pass. However, as the results show, the average ratio of bit-vector width to maximum value-number is 0.18. This reflects a savings of over 80%.

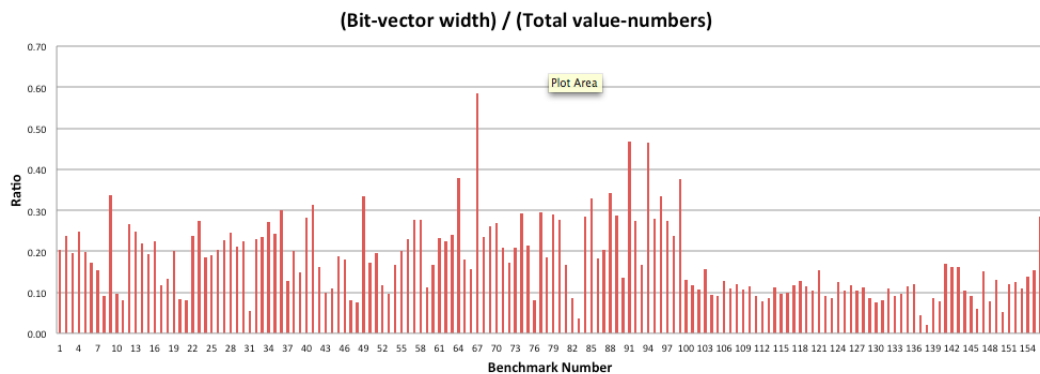


Figure 2

Local CSE

For our data flow equations to work efficiently, a local CSE pass is run on each basic block. Basically, this pass removes the redundancies in straight line basic block code and sanitizes it for the iterative bit vector algorithm. Borrowed from [8], the main idea is to trim the amount of work to be done by the PRE pass. For example, if there are many expressions with the same value number in a basic block, rather than PRE going over all of them, local CSE can weed out the redundancies. We perform this step before calling our data flow framework.

Insert and Replace

To maintain compatibility with SSA, we perform insertion and replacement through memory and re-run the *mem2reg* pass after our PRE pass to convert the newly created load and store instructions to register operations. Following are the major points:

- Assign stack space (*allocas*) at the beginning of the function for all the expressions that need movement.
- At insertion point, compute the expression and save the value to the stack slot assigned to the expression
- At replacement point, load from the correct stack slot, replace all uses of the original expression with the load instruction, and delete the original expression
- *mem2reg* converts stack operations to register operations and introduces the necessary Φ instructions

In appendix E, we have shown in Figure E.1 and Figure E.2, the optimizations performed by our PRE pass.

5. Miscellaneous

Zero-trip Loops

Our algorithm moves the loop invariant computations to the loop pre-header only if placement in the loop pre-header is anticipatable. Such a pre-header is always available for *do-while* loops, but not for *while* and *for* loops. Hence, a modification is required to the structure of *while* and *for* loops which peels off the first iteration of the loop, protected by the loop condition. This alteration provides PRE with a suitable loop pre-header to hoist loop independent computations to. In Figure D.1 (Appendix D) we show the CFG changes. We achieved this effect using an existing LLVM pass *-loop-rotate*.

Critical Edges

A critical edge in a flow graph is an edge from a node with multiple successors to a node with multiple predecessors. Splitting such edges and inserting dummy nodes aids PRE by offering more anticipatable points. We used an existing LLVM pass (*BreakCriticalEdges*) for the same. In many cases, however, the dummy nodes created by this pass do not hold any computation after PRE. We used *-simplifycfg* to clean up the mess created by *BreakCriticalEdges*.

Unresolved Issues

There were a couple of issues on which we would have liked to spend more time. The first is redundancy elimination for expressions which are lexically different in SSA, and have different value numbers. We came up with a few techniques within the bounds of our existing PRE code, but unfortunately, none could be generalized to solve the core problem. The second issue pertains to the insertion step of our algorithm and needs slightly detailed explanation. Suppose that an expression, with value number *vn*, is to be inserted in a basic block. Although our algorithm can handle all cases, for simplicity, assume that the insertion point is the end of the basic block. To insert the expression we scan the list of the expressions in the whole function which have the same value number *vn*. We then clone one of these expressions (called *provider*) and place at the end of the basic block. The trivial case is when the provider is available in the same basic block. If however, the provider comes from another basic block, then we need to ensure that the operands of the provider dominate the basic block we wish to insert the expression in. Not being able to find a suitable provider is the only case where we override the suggestion of the data flow analysis and not do PRE for that expression only. PRE for other expressions proceeds as usual. Our exhaustive testing on multiple suites suggests that this is a very rare occurrence.

Testing

While working on the project, we wrote 25 small test cases to capture the intricate movements of expressions in the partial redundancy elimination algorithm. Most of these contrived test cases, along with our full source code, can be found on our project Github link <https://github.com/sdasgup3/PRE>. For evaluation on real life applications, we chose 3 different suites - LLVM SingleSource, LLVM MultiSource, SPEC2006. For correctness, we checked the output of the binary optimized with our PRE pass with the provided reference output. All benchmarks passed the correctness test. For each suite, we present two sets of performance results. The first set compares the performance of binaries optimized with our version of PRE (henceforth referred to as LCM-PRE) with binaries without PRE optimization (henceforth referred to as BASE). The second set compares the performance of LCM-PRE binaries with binaries optimized with LLVM’s version of PRE (henceforth referred to as GVN-PRE). To remove noise, we run each benchmark thrice and take the average. Also, benchmarks with running time of less than 5 seconds are not accounted for. The next two subsections describe the performance S-curves, following which we summarize in a table, the absolute run-times for three benchmarks from each suite. For a meaningful comparison, we use the same set of optimization knobs for BASE, GVN-PRE and LCM-PRE.

Pass Name	-opt switches
BASE	-mem2reg -loop-rotate -reassociate -mem2reg -simplifycfg
LCM-PRE	-mem2reg -loop-rotate -reassociate -lcm -mem2reg -simplifycfg
GVN-PRE	-mem2reg -loop-rotate -reassociate -gvn -mem2reg -simplifycfg

Table 1: Optimization knobs.

LLVM Single source & Multi source

We ran 45 benchmarks from the SingleSource package. Figure 3(a) shows the S-curve for BASE time over LCM-PRE time. For most of the benchmarks (40/45) we either increase performance (up to 42%) or maintain the same level. 5 benchmarks show slight degradation which is bound by 6.5%. Figure 3(b) shows the S-curve for GVN-PRE time over LCM-PRE time. It is heartening to beat GVN-PRE in a few cases.

Results for the MultiSource benchmarks follow a similar trend. Out of the 45 benchmarks from this package, 41 show improvement (up to 23%) or maintain same performance for BASE time over LCM-PRE time (Figure 4(a)), while degradation for the rest is bound by 5%. GVN-PRE time over LCM-PRE time is shown in 4(b).

Spec2006 Benchmark

We augmented our testing infrastructure to support the SPEC2006 suite. Both SPEC-INT and SPEC-FP were tested. We, however, had to limit our testing to C/C++ benchmarks, and leave out Fortran. Getting SPEC-Fortran benchmarks to run inside LLVM needs extra support. We take our inputs for the SPEC runs from the following source - http://boegel.kejo.be/ELIS/spec_cpu2006/spec_cpu2006_command_lines.html

Out of the 27 runs from SPEC, 19 show improvement (up to 52%) or maintain same performance for BASE time over LCM-PRE time, while degradation for the rest is bound by 10% (Figure 5(a)). Our pass triumphs over GVN-PRE for quite a few cases here as well as shown in 5(b).




Figure 3: Performance evaluation with LLVM SingleSource Benchmark




Figure 4: Performance evaluation with LLVM MultiSource Benchmark




Figure 5: Performance evaluation with SPEC 2006

Benchmark Name	BASE Time (B)	LCM-PRE Time(L)	GVN-PRE Time(G)	B/L	G/L
SingleSource/Benchmarks/Dhrystone/fldry	5.405	4.965	5.263	1.088	1.060
SingleSource/Benchmarks/Misc/ourafft	5.112	4.281	4.286	1.194	1.001
SingleSource/Benchmarks/Misc/lowercase	40.795	28.612	28.628	1.425	1.000
MultiSource/Benchmarks/TSVC/NodeSplitting-flt	7.378	6.706	5.398	1.100	0.804
MultiSource/Benchmarks/TSVC/Expansion-flt	6.339	5.734	5.308	1.105	0.925
MultiSource/Benchmarks/TSVC/Expansion-dbl	7.134	5.787	5.355	1.232	0.925
SPECINT2006/456.hmmmer (ref-input 1)	1547.77	1019.885	993.581	1.517	0.974
SPECINT2006/456.hmmmer (ref-input 2)	715.641	521.318	455.754	1.372	0.874
SPECINT2006/464.h264ref	185.551	168.563	163.363	1.100	0.969

Table 2: Absolute run-times for 3 benchmarks from each suite

Performance Analysis

We analyzed the benchmarks where our pass degrades performance. In this subsection, we summarize our thoughts and findings. To measure the improvements for LCM-PRE over BASE, we switch off all backend optimizations for all runs. More specifically, we use `-O0` while converting the LLVM bitcode to machine code. A major repercussion of using `-O0` is that none of the efficient register allocators (greedy, pbqp) can be used (LLVM restriction). Hence, we were stuck with the fast register allocator which does a very poor job for some of the benchmarks. The performance of LCM-PRE is sensitive to register allocation (because of increased register pressure), and this causes the performance dip over BASE as presented in the S-curves. We substantiate this claim with an example from the SingleSource package (*Benchmarks/Shootout-C++/methcall*) (Table 3). Data from `llc-dump` shows the increased amount of loads and stores to the stack for LCM-PRE. We also gather the dynamic data from Pin, using a simple opcode-mix tool which we wrote. The increased number of stack reads (15%) and writes (6%) at runtime for LCM-PRE confirms our hypothesis. We hold the opinion that using a more powerful register allocator would wipe off most of the performance losses.

Next we explain why we chose to stick with `-O0` rather than using `-O3`. This was done to disable backend optimizations such as `-machine-licm`. We expect major performance gains from the loop invariant code motion done by LCM-PRE, and allowing a backend pass to achieve the same effect on BASE would steal our thunder. This was confirmed experimentally, where using `-O3` in the backend results in LCM-PRE execution time same as BASE for all the benchmarks (no improvement, no degradation).

Dynamic data from Pin tool		
Stats	LCM-PRE	BASE
Total Instructions (in Billion)	72	68
stack-read count (in Billion)	19	17
stack-write count (in Billion)	16	15
Static data from llc tool - fast register allocator		
regalloc-Number of loads added	34	28
regalloc-Number of stores added	35	30

Table 3

Appendix A

Computation of localized sets

For each basic block there are 3 bit vectors dedicated to the block-specific properties, namely **Transp**, **Antloc** and **Xcomp**. As mentioned before, a bit vector is a boolean array of value numbers. Let the leader expression (as defined in the section on value numbering) associated with the value number v be called $L(v)$.

$$\begin{aligned} \text{Transp}(v, B) &= \begin{cases} \textit{false} & \text{iff } \exists x \in \text{operands of } L(v) \text{ such that } \text{Mod}(x, B) = \textit{true} \\ \textit{true} & \text{Otherwise} \end{cases} \\ \text{Antloc}(v, B) &= \text{Eval}(v, B) \cap \text{Transp}(v, B) \\ \text{Xcomp}(v, B) &= \text{Eval}(v, B) \cap \overline{\text{Transp}(v, B)} \end{aligned}$$

where

$$\begin{aligned} \text{Eval}(v, B) &= \{v \mid \text{value number } v \text{ is computed in } B\} \\ \text{Mod}(op, B) &= \text{operand } op \text{ modified in } B \end{aligned}$$

(A.1)

Appendix B

Lazy Code motion Transformations

- Down Safety Analysis (Backward data flow analysis)

$$\begin{aligned} \text{Antin}(b) &= \text{Antloc}(b) \cup (\text{Tranp}(b) \cap \text{Antout}(b)) \\ \text{Antout}(b) &= \text{Xcomp}(b) \cup \begin{cases} \phi & \text{if } b = \text{exit} \\ \bigcap_{s \in \text{succ}(b)} \text{Antin}(s) \end{cases} \end{aligned} \quad (\text{B.1})$$

- Up Safety Analysis (Forward data flow analysis)

$$\begin{aligned} \text{Availin}(b) &= \begin{cases} \phi & \text{if } b = \text{entry} \\ \bigcap_{p \in \text{pred}(b)} (\text{Xcomp}(p) \cup \text{Availout}(p)) \end{cases} \\ \text{Availout}(b) &= \text{Tranp}(b) \cap (\text{Antloc}(b) \cup \text{Availin}(b)) \end{aligned} \quad (\text{B.2})$$

- Earliest-ness (No data flow analysis)

$$\begin{aligned} \text{Earliestin}(b) &= \text{Antin}(b) \cap \bigcap_{p \in \text{pred}(b)} \overline{(\text{Availout}(p) \cup \text{Antout}(p))} \\ \text{Earliestout}(b) &= \text{Antout}(b) \cap \overline{\text{Tranp}(b)} \end{aligned} \quad (\text{B.3})$$

- Delayability (Forward data flow analysis)

$$\begin{aligned} \text{Delayin}(b) &= \text{Earliestin}(b) \cup \begin{cases} \phi & \text{if } b = \text{entry} \\ \bigcap_{p \in \text{pred}(b)} \overline{(\text{Xcomp}(p) \cap \text{Delayout}(p))} \end{cases} \\ \text{Delayout}(b) &= \text{Earliestout}(b) \cup \overline{(\text{Delayin}(b) \cap \text{Antloc}(b))} \end{aligned} \quad (\text{B.4})$$

- Latest-ness (No data flow analysis)

$$\begin{aligned} \text{Latestin}(b) &= \text{Delayin}(b) \cap \text{Antloc}(b) \\ \text{Latestout}(b) &= \text{Delayout}(b) \cap (\text{Xcomp}(b) \cup \bigcup_{s \in \text{succ}(b)} \overline{\text{Delayin}(s)}) \end{aligned} \quad (\text{B.5})$$

- Isolation Analysis (Backward data flow analysis)

$$\begin{aligned}
 \text{Isolatedin}(b) &= \text{Earliestout}(b) \cup \text{Isolatedout}(b) \\
 \text{Isolatedout}(b) &= \begin{cases} U & \text{if } b = \text{exit} \\ \bigcap_{s \in \text{succ}(b)} (\text{Earliestin}(s) \cup \overline{(\text{Antloc}(s) \cap \text{Isolatedin}(s))}) & \end{cases}
 \end{aligned} \tag{B.6}$$

- Insert and Replace points

$$\begin{aligned}
 \text{Insertin}(b) &= \text{Latestin}(b) \cap \overline{\text{Isolatedin}(b)} \\
 \text{Insertout}(b) &= \text{Latestout}(b) \cap \overline{\text{Isolatedout}(b)} \\
 \text{Replacein}(b) &= \text{Antloc}(b) \cap \overline{\text{Latestin}(b) \cap \text{Isolatedin}(b)} \\
 \text{Replaceout}(b) &= \text{Xcomp}(b) \cap \overline{\text{Latestout}(b) \cap \text{Isolatedout}(b)}
 \end{aligned} \tag{B.7}$$

Appendix C

Generalized data flow framework

All the equations in Appendix B can be computed using the generic framework defined below.

C.1 Forward Analysis

$$\begin{aligned} \text{In}(b) &= \alpha(b) \cup \begin{cases} \perp & \text{if } b = \text{entry} \\ \bigwedge_{p \in \text{pred}(b)} \beta(p) & \end{cases} \\ \text{Out}(b) &= \gamma(b) \end{aligned} \quad (\text{C.1})$$

C.2 Backward Analysis

$$\begin{aligned} \text{In}(b) &= \gamma(b) \\ \text{Out}(b) &= \alpha(b) \cup \begin{cases} \perp & \text{if } b = \text{exit} \\ \bigwedge_{s \in \text{succ}(b)} \beta(s) & \end{cases} \end{aligned} \quad (\text{C.2})$$

The following is the function which we call with dataflow equation specific parameters defined subsequently.

`callFramework(Out(b), In(b), $\alpha(b)$, $\beta(b)$, $\gamma(b)$, \bigwedge , \perp , \top , Direction)`

Following is the list of values that we need to plug-in to α , β and γ for the above generic framework to work.

- Down Safety Analysis (Backward data flow analysis)

$$\begin{aligned} \alpha(x) &= \text{Xcomp}(x) \\ \beta(x) &= \text{Antin}(x) \\ \gamma(x) &= \text{Tranp}(x) \cap \text{Antout}(x) \cup \text{Antloc}(x) \\ \bigwedge &= \cap \\ \perp &= \phi \\ \top &= V, \text{ set of all values} \\ \text{Direction} &= \text{Backward} \end{aligned} \quad (\text{C.3})$$

- Up Safety Analysis (Forward data flow analysis)

$$\begin{aligned}
\beta(x) &= \text{Xcomp}(x) \cup \text{Availout}(x) \\
\gamma(x) &= \text{Antloc}(x) \cup \text{Availin}(x) \cap \text{Tranp}(x) \\
\bigwedge &= \cap \\
\perp &= \phi \\
\top &= V, \text{ set of all values} \\
\text{Direction} &= \text{Forward}
\end{aligned} \tag{C.4}$$

- Delayability (Forward data flow analysis)

$$\begin{aligned}
\alpha(x) &= \text{Earliestin}(x) \\
\beta(x) &= \overline{\text{Xcomp}(x)} \cap \overline{\text{Delayout}(x)} \\
\gamma(x) &= \text{Delayin}(x) \cap \overline{\text{Antloc}(x)} \cup \text{Earliestout}(x) \\
\bigwedge &= \cap \\
\perp &= \phi \\
\top &= V, \text{ set of all values} \\
\text{Direction} &= \text{Forward}
\end{aligned} \tag{C.5}$$

- Isolation Analysis (Backward data flow analysis)

$$\begin{aligned}
\beta(x) &= \overline{\text{Antloc}(x)} \cap \text{Isolatedin}(x) \cup \text{Earliestin}(x) \\
\gamma(x) &= \text{Earliestout}(x) \cup \text{Isolatedout}(x) \\
\bigwedge &= \cap \\
\perp &= V, \text{ set of all values} \\
\top &= V, \text{ set of all values} \\
\text{Direction} &= \text{Backward}
\end{aligned} \tag{C.6}$$

Appendix D

Transformations for “Zero-trip Loops”

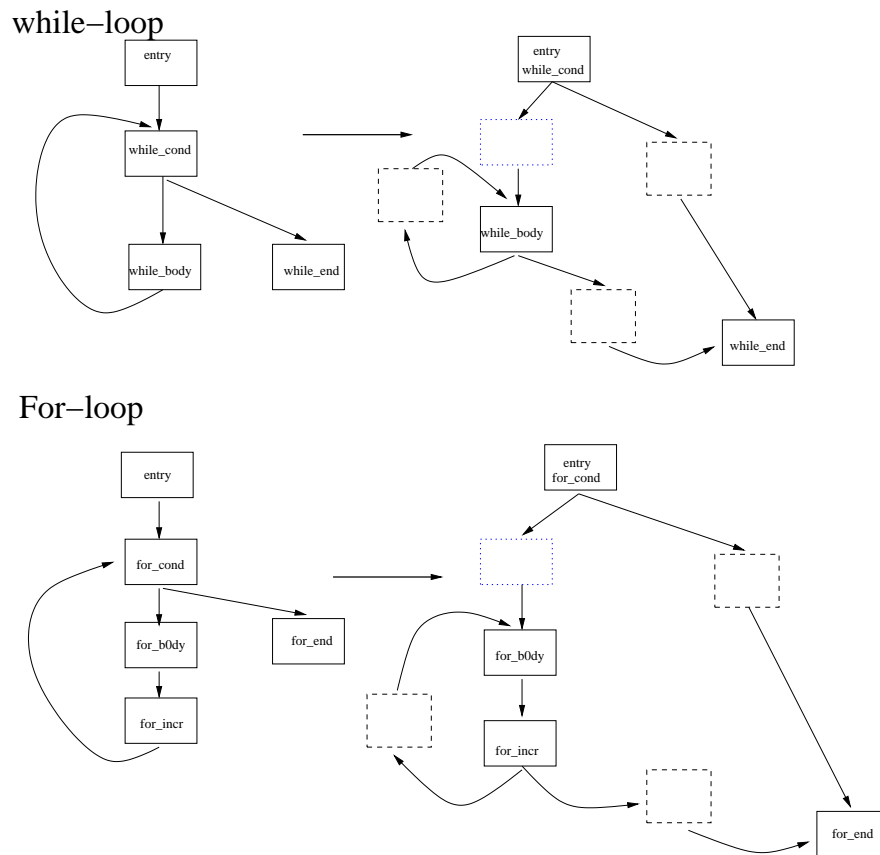


Figure D.1: Loop transformations done by *-loop-rotate*. *do-while* loops remain unaffected. Blue dotted boxes are the ones inserted by loop rotate. PRE can insert the computations in these places.

Appendix E

An Extended Example

Here we show, through an example code, the optimizations performed by our PRE pass. The intention here is to highlight redundancy elimination for expressions $a + b$ & $a < b$. Optimal placements are marked in Figure E.2. Some of the notable observations are:

- Black dotted boxes denote basic blocks inserted because of critical edge splitting
- Blue dotted boxes are the loop pre-headers inserted by -loop-rotate pass. PRE can insert computations here.
- Inserted statements are marked **blue** and replaced ones with **magenta**
- LCSE (Local common subexpression elimination) happened in BB2.
- For the loop BB7, BB9 in Figure E.1, LICM happened wherein the computation of $a + b$ is moved from BB9 (in Figure E.1) to BB8 (in Figure E.2). BB8 is the loop pre-header

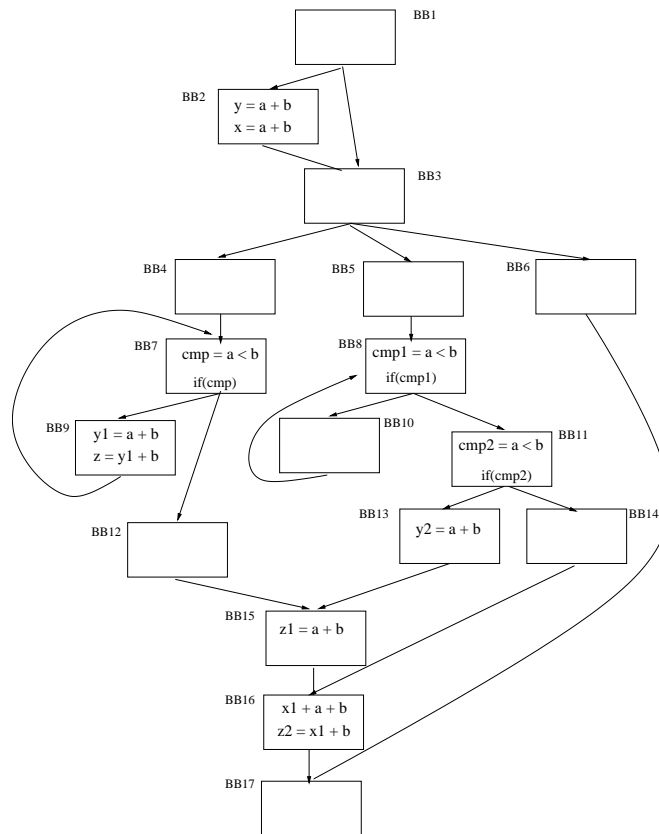


Figure E.1: A motivating example

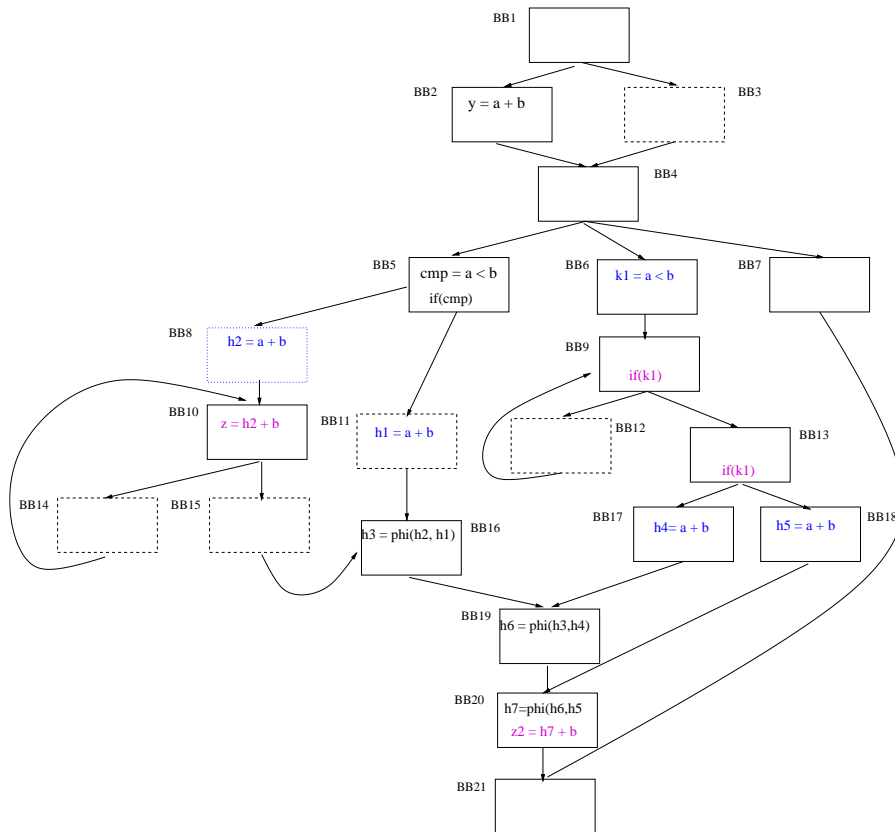


Figure E.2: Lazy code motion transformation on computations $a + b$ & $a < b$.

Bibliography

- [1] B. ALPERN, M. N. WEGMAN, AND F. K. ZADECK, *Detecting equality of variables in programs*, in Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '88, New York, NY, USA, 1988, ACM, pp. 1–11.
- [2] P. BRIGGS AND K. D. COOPER, *Effective partial redundancy elimination*, in Proceedings of the ACM SIGPLAN 1994 Conference on Programming Language Design and Implementation, PLDI '94, New York, NY, USA, 1994, ACM, pp. 159–170.
- [3] P. BRIGGS, K. D. COOPER, AND L. T. SIMPSON, *Value numbering*, *Softw. Pract. Exper.*, 27 (1997), pp. 701–724.
- [4] J. COCKE, *Programming Languages and Their Compilers: Preliminary Notes*, Courant Institute of Mathematical Sciences, New York University, 1969.
- [5] K. D. COOPER AND L. T. SIMPSON, *Scc-based value numbering*, *Software Practice and Experience*, 27 (1995), pp. 701–724.
- [6] K. HEINZ DRECHSLER AND M. P. STADEL, *A variation of knoop, rothing, and steffen's lazy code motion*.
- [7] R. KENNEDY, S. CHAN, S. MING LIU, R. LO, P. TU, AND F. CHOW, *Partial redundancy elimination in ssa form*, *ACM Transactions on Programming Languages and Systems*, 21 (1999), pp. 627–676.
- [8] J. KNOOP, O. RÜTHING, AND S. BERNHARD, *Optimal code motion: Theory and practice*, *ACM Trans. Program. Lang. Syst.*, 16 (1994), pp. 1117–1155.
- [9] J. KNOOP, O. RÜTHING, AND B. STEFFEN, *Lazy code motion*, in Proceedings of the ACM SIGPLAN 1992 Conference on Programming Language Design and Implementation, PLDI '92, New York, NY, USA, 1992, ACM, pp. 224–234.
- [10] E. MOREL AND C. RENVOISE, *Global optimization by suppression of partial redundancies*, *Commun. ACM*, 22 (1979), pp. 96–103.
- [11] L. T. SIMPSON AND K. D. COOPER, *Value-driven code motion*, *IEEE Transactions on Image Processing*, (1995).
- [12] T. VANDRUNEN AND A. L. HOSKING, *Value-based partial redundancy elimination*, in *In CC*, 2004, pp. 167–184.