

CS 243 Final Examination

Winter 2023

March 21, 2023

This is an open-book, open-notes, open-laptop, *closed*-network exam. You have 3 hours to complete the exam. The examination has 8 problems worth 180 points. Please budget your time accordingly. Write your answers in the space provided on the exam. If you use additional scratch paper, please turn that in as well.

Your Name: _____ SUNet ID: _____

The following is a statement of the Stanford University Honor Code:

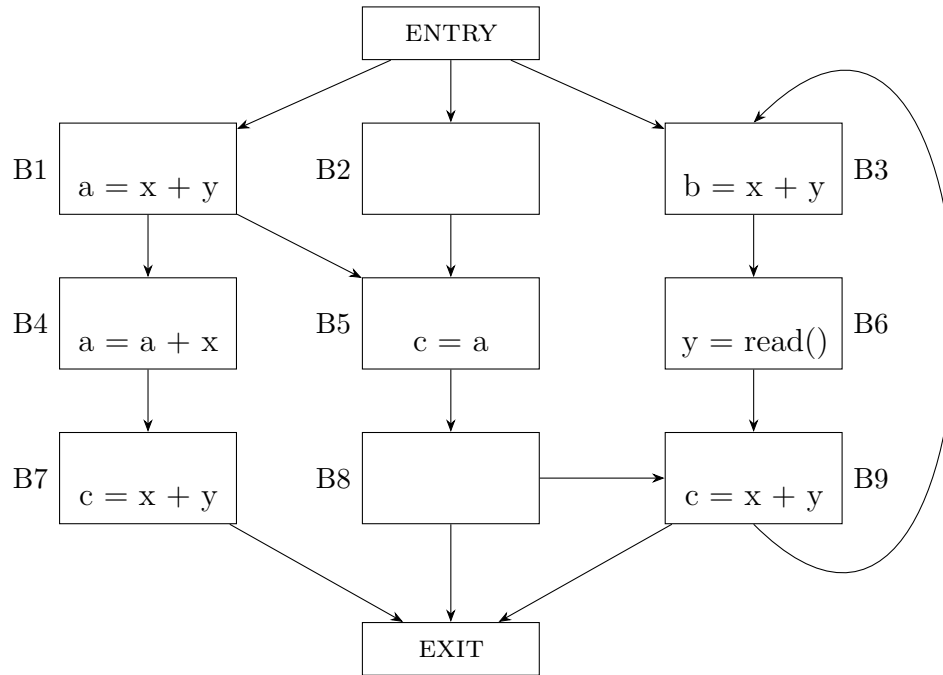
1. The Honor Code is an undertaking of the students, individually and collectively:
 - (a) that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
 - (b) that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.
2. The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.
3. While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

Signature: _____

Problem	#1	#2	#3	#4	#5	#6	#7	#8	Total
Max	21	14	15	20	25	25	30	30	180

Problem 2. Partial Redundancy Elimination [14 points]


Answer the following question based on the program below.



Show the result of running partial redundancy elimination. **Don't apply any optimization other than PRE.** What's the final optimized flow graph? You don't need to show the intermediate steps.

Problem 4. Software Pipelining [20 points]

Consider the following dependence graph for a single iteration of a loop, with resource constraints:



1. What are the bounds on the initiation interval T according to the precedence and resource constraints for this program? [4 points]

2. Is it possible to software-pipeline this loop with the minimum bound found in the previous part? If not, what is the minimum possible initiation interval?

Show the modulo reservation table for the optimal software-pipelined schedule. Also show the code schedule for an iteration in the source loop. [8 points]

3. Can the scheduling algorithm described in class produce the optimal schedule for this loop? If not, show the modulo reservation table and code schedule generated by the algorithm. [6 points]

Problem 5. Pointer Analysis [25 points]

Consider the following simple implementation of a singly-linked list. We use static methods to avoid `this` pointers which may be ambiguous.

```
1 public class Node {
2     public Node n = null;
3
4     // This function is callable as: Node.extend(a, b)
5     public static void extend(Node a, Node b) {
6         if (b == null)
7             b = new Node();    // h3
8         Node c = Node.getTail(a);
9         c.n = b;
10    }
11
12    // This function is callable as: Node.getTail(d)
13    public static Node getTail(Node d) {
14        if (d.n == null)
15            return d;
16        Node e = d.n;
17        return Node.getTail(e);
18    }
19
20    public static void main(String[] args) {
21        Node x = new Node();    // h1
22        Node y = new Node();    // h2
23        Node.extend(x, y);
24    }
25 }
```

Apply flow-insensitive context-insensitive pointer analysis to the program and write down all \mathbf{vP} and \mathbf{hP} tuples derived from the analysis. $\mathbf{h1}$, $\mathbf{h2}$, $\mathbf{h3}$ are the allocation sites, and \mathbf{a} , \mathbf{b} , \mathbf{c} , \mathbf{d} , \mathbf{e} , \mathbf{x} , \mathbf{y} are all the variables.

(This page is intentionally blank.)

Problem 6. Satisfiability Modulo Theories [25 points]

In this problem, we use SMT to analyze the following function. We assume that `data` is a zero-indexed array of length `N`.

```
1  int foo(int data[], int N) { {
2      int v = 0;
3      int i = 0;
4      while (i <= N) {
5          if (data[i] > v) {
6              v = data[i];
7          }
8          i = i + 1;
9      }
10     return v;
11 }
```

1. Notice the while loop on lines 4–9. One method for handling a loop in SMT analysis is to unroll it several times. Show the SSA form of the code when you unroll the loop twice (convert the first two iterations of the `while` loop into `if` statements). [12 points]

2. Let P be the SMT constraints based on the SSA form of the code with the loop unrolled twice. Write a SMT formula that is satisfied if there is an out-of-bound array access. Make sure to take into account path conditions for each array access.

Will an SMT solver return **SAT** or **UNSAT**? If **SAT**, give a satisfying model. If **UNSAT**, briefly explain why. [5 points]

3. Suppose we can safely assume that the `data` array has at least 10 elements. Still unrolling the loop twice, write down the updated SMT formula that is satisfied if there is an out-of-bound array access. Will an SMT solver return **SAT** or **UNSAT**? If **SAT**, give a satisfying model. If **UNSAT**, briefly explain why. [4 points]

4. Again assume that the `data` array has at least 10 elements. What if we unroll the loop 20 times? Will an SMT solver return **SAT** or **UNSAT** on the updated constraint? If **SAT**, give a satisfying model. If **UNSAT**, briefly explain why. [4 points]

3. Does this function have pipelined parallelism? If so, show the transformed code with as many outermost fully permutable loops as possible. (You only need to show the transformed sequential code, there is no need to parallelize it with synchronization.) [9 points]

4. Describe how you would parallelize this function. Discuss the cache performance for the individual processors, as well as the synchronization and communication cost between processors. [9 points]

Problem 8. Optimizing Dynamically Typed Languages II [30 points]

Consider a dynamically-typed language similar to JavaScript and Python. The language has two kinds of instructions of interest to us:

- $x = T(\dots)$, where T is one of **int**, **float**, and **string**. This sets the variable x to a value of type T .
- $x = y + z$, where x, y, z are all variables.

Unlike statically-typed languages like C and Java, the type of a program variable is allowed to change in this language. As an example, the program `x = int(1); x = float(3.14)` will result in variable `x` having an **int** value at first, but afterwards a **float** value.

The addition operator `+` accepts operands of any types with the following rules:

1. The sum of two values of the same type has that same type.
2. The sum of an **int** and a **float** is a **float**.
3. The sum of a **string** and a value of any type is a **string**.

Suppose we are writing an optimizing compiler for this language. Because there are special instructions for adding two integers (**iAdd**), adding two floats (**fAdd**), and adding two strings (**sAdd**), your task is to identify those additions that can be replaced with such special instructions at compile time.

To get full points, your dataflow analysis should identify as many opportunities for optimization as possible.

This problem has four parts: (1) Define your dataflow analysis by filling in the table below. (2) State clearly how you perform the optimization. (3) Is your dataflow analysis monotone? Explain. (4) Is your dataflow analysis distributive? Explain.

(Questions begin on the next page.)

1. Dataflow analysis specification. You may assume that each instruction is in its own basic block. [22 points]

Direction of your analysis (forward/backward)	
Lattice elements and meaning	
Meet operator	
Is there a top element? If yes, what is it?	
Is there a bottom element? If yes, what is it?	
Transfer function	
Boundary condition	
Interior points	

2. How do you optimize the program? [2 points]

3. Is your dataflow analysis monotone? Briefly sketch a proof for it if so, or provide a counterexample if not. [3 points]

4. Is your dataflow analysis distributive? Briefly sketch a proof for it if so, or provide a counterexample if not. [3 points]