

1. (10pt) **Name:** _____

INSTRUCTIONS: Carefully read each question, and write the answer in the space provided. If answers to free response questions are written obscurely, zero credit will be awarded. The correct answer to a free response question with a short answer (i.e., one word or phrase) will never contain any significant words used in the question itself (i.e., “crossword rules”). You are permitted to use **two** 8.5x11 inch sheets of paper (double-sided) containing **hand-written** notes; all other aids (other than your brain) are forbidden. Questions may be brought to the instructor.

This exam is designed to take **80 minutes** (but you may take as much as **150 minutes**, if you want to). There are **800 points** available on the exam, and most questions are worth a number of points that is divisible by ten. These point values are a rough guide to how long I think you should spend on each question.

For **TRUE** or **FALSE** and multiple choice questions, circle your answer.

On free response questions only, you will receive **20%** credit for any question which you leave blank (i.e., do not attempt to answer). Do not waste your time or mine by making up an answer if you do not know. (Note though that most questions offer partial credit, so if you know part of the answer, it is almost always better to write something rather than nothing.)

To get credit for this question, you must:

- Print your name (e.g., “Martin Kellogg”) in the space provided on this page.
- Print your UCID (e.g., “mjk76”) in the space at the top of **each** page of the exam.

Point distribution (blanks for graders only):

Page 1: _____ / 10

Pages 2 and 3: _____ / 320

Pages 4 and 5: _____ / 260

Pages 6 and 7: _____ / 210

Extra Credit: _____ / 0

Total: _____ / 800

I. Multiple Choice and Very Short Answer (120pts). In the following section, either circle your answer (possible answers appear in **bold**) or write a very short (one word or one phrase) answer in the space provided. No partial credit is possible in this section.

2. (20pt) A context-free grammar is the theoretical grounding for which of these? (Select all that apply.)
 - A** typechecking
 - B** lexing
 - C** parsing
 - D** dataflow analysis
3. (20pt) The dataflow analyses that we use for enabling optimizations should be: (Select all that apply.)
 - A** sound
 - B** complete
 - C** conservative
 - D** monotonic
4. (20pt) List scheduling does a _____ of a precedence graph of operations to construct a schedule, one cycle at a time.
5. (20pt) “If B is a subclass of A, then an object of class B can be used wherever an object of class A is expected” is the statement of which of these:
 - A** Liskov substitution principle
 - B** Rice’s theorem
 - C** Church-Turing thesis
 - D** Curry-Howard correspondence
6. (20pt) Which of the following are advantages of using exceptions instead of error return codes in a program? (Select all that apply.) Exceptions...
 - A** ...avoid forgotten error checks.
 - B** ...avoid choosing arbitrary error values.
 - C** ...can carry detailed error information.
 - D** ...enable separation of normal logic and error logic.
7. (10pt) **TRUE** or **FALSE**: regional optimizations are not used in production compilers, because each extended basic block contains exponentially-many paths in the number of constituent basic blocks.
8. (10pt) **TRUE** or **FALSE**: abstract interpretation, dataflow analysis, and pluggable types are equivalently-expressive formalisms. That is, there is nothing that you can express in one of them that you cannot express in the other two.

- | | | | |
|--------------------------|--|------------------------------------|------------------------|
| A. Stop and Copy | B. Boolean Satisfiability | C. Single-Static Assignment | D. Basic Block |
| E. Stack Machine | F. Peephole Optimization | G. Global Optimization | H. Duck Typing |
| I. Mark and Sweep | J. Interprocedural Optimization | K. Activation Record | L. Header File |
| M. Spill | N. Curry-Howard Correspondence | O. Graph Coloring | P. Symbol Table |

II. Matching (200pts). This section contains a collection of terms discussed in class in an “Answer Bank” (choices **A.** through **P.**). Each question in this section describes a situation associated with an answer in the Answer Bank. Write the letter of the term in the Answer Bank that best describes each situation. Each answer in the Answer Bank will be used at most once.

9. (20pt) _____ Maria’s compiler does a whole-program analysis to determine that it will be profitable to inline a function.
10. (20pt) _____ Gottfried writes a program in F*, a language whose type system allows him to encode the program’s full functional specification into its type.
11. (20pt) _____ Bernhard’s programming language has made a design choice that favors expressiveness and flexibility over static safety.
12. (20pt) _____ Leonhard’s program has too many live variables at the same time to fit into the registers on the machine.
13. (20pt) _____ Blaise’s compiler reduces the register allocation problem to a standard NP-complete problem instead of solving it directly.
14. (20pt) _____ Srinivasa needs to typecheck object files that were compiled separately.
15. (20pt) _____ When proving the correctness of an optimization, Sofia relies on the fact that each variable in the IR is referentially transparent.
16. (20pt) _____ René’s compiler replaces an adjacent store and load of the same value with a register move.
17. (20pt) _____ In order to make it easier to show that his compiler is correct, Carl decides to use a simple model of computation.
18. (20pt) _____ Emmy is writing a garbage collector for a language that allows the programmer to store raw pointers.

III. Short answer (470pts). Answer the questions in this section in at most three sentences, unless the question gives other instructions (question-specific instructions have higher priority).

19. (100pt) Rewrite the following program using local optimizations, showing your work, until there are no more local optimizations (from among those we discussed in-class) that you can safely apply. At each step, apply *only one optimization*: do not mix optimizations. The first optimization has been applied for you. After applying each optimization, rewrite *only the lines that were changed by that optimization* and clearly label which optimization you have applied, as in the example.

1 a := 2 * x 2 b := 2 + 2 3 c := a * b 4 d := 6 + b 5 e := c 6 f := e + d 7 g := c + e 8 print g	<u>Constant Folding</u> →	1 2 b := 4 3 4 5 6 7 8 ...
---	------------------------------	---

20. Suppose that we want to add automatic memory management to Cool. We decide to use a **reference counting** garbage collector.

(a) (20pt) Which two constructs' operational semantics rules must we modify to support reference counting? _____ and _____

(b) (50pt) Write the **standard** Cool operational semantics rule for **ONE** of these constructs. Circle the construct that you use for this part in your answer to part a.

(c) (50pt) Write a **modified** operational semantics rule for **the SAME construct you chose in part b** that adds support for reference counting.

(d) (40pt) Describe a direct impact of these modifications on code generation, from the perspective of optimization. For example, you might describe an optimization that is easier (or harder) on the resulting code, or an optimization that now more (or less) profitable.

21. (60pt) Show that an abstract interpretation $I = (A, \sqsubseteq, \alpha, \gamma)$ terminates. If you rely on any standard property of an abstract interpretation or one of its components that we discussed in class, you must give a correct (mathematical) definition of what it means in the context of abstract interpretation.
22. (40pt) Support or refute the following claim: it is always profitable to unroll a loop if you can prove that doing so is safe.
23. (40pt) Support or refute the following claim: a compiler can change the x86-64 calling convention for methods that it generates.

24. (30pt) The goal of a register allocator is *not* just to map the v virtual registers used by a block to the k physical registers that are actually present on the target machine, if the block uses more virtual registers than are actually available (i.e., $v > k$). Why not?
25. (40pt) Support or refute the following claim: it is sound for one pluggable type system to rely on type qualifiers from a *different* pluggable type system as a premise in one of its typing judgments.

IV. Extra Credit. Questions in this section do not count towards the denominator of the exam score.

26. (10pt) In section II (Matching), there is a theme to the names used in the situation descriptions. What is the theme?
27. (10pt) Give the last name of any three of the people who inspired the names used in the situation descriptions in section II (Matching).
28. (10pt) Suppose that we wanted to modify the PA3 and PA4 assignments so that it's no longer necessary to pass `-no-pie -static` to `gcc` when assembling your `.s` files. What would you, the student, need to do differently in PA3 and PA4?
29. (10pt) May I include your final PA4 submission as an anonymous "boss battle" in future iterations of the PA4 leaderboard? (Answer "Yes" or "No". Either of those strings receives full credit for this question; any other string, including the empty string, receives zero. If you are working with a partner, I will only include your submission if you both answer "Yes".)

This page intentionally left blank (you may use it as scratch paper, and the proctor will have more scratch paper at the front if you need more).

