

CSE 501: Implementation of Programming Languages

Goals:

- Understand how languages get implemented efficiently
- Understand what can and can't be done, with how much effort
- Understand state of current research in efficient language implementation

Prerequisites:

- CSE 401 or equivalent
- CSE 505 or equivalent

Readings:

- main text: *Modern Compiler Implementation*, by Appel
- for reference: *Compilers: Principles, Techniques, and Tools*, by Aho, Sethi, and Ullman
- plus important papers from the literature

Grading:

- Midterm: 25%
- Final: 30%
- Homework: 20%
- Project: 25%

Course web page:

<http://www.cs.washington.edu/education/courses/cse501/CurrentQtr>

- office hours, office locations
- course mailing list signup instructions
- on-line copies of all lecture slides, handouts, assignments, etc.
- course project information and instructions

Rough Course Outline

Week	Topic
1/3	Intro; structure of compilers; standard optimizations Standard intermediate representations; control flow, data flow; dependence Dataflow analysis; reaching constants, constant propagation
1/8	Lattice-theoretic data flow analysis framework; integer range analysis Data flow analyzer generators, frameworks
1/15	Holiday Advanced intermediate representations: def/use chains, control dependence tree, SSA form, VDG CSE; loop-invariant code motion
1/22	Inlining Interprocedural analysis; analysis with first-class functions, dynamically dispatched messages
1/29	Procedure specialization, partial evaluation Alias and pointer analysis
2/5	Dependence analysis; automatic parallelization
2/12	Global register allocation Calling conventions; callee-save vs. caller-save Link-time register allocation
2/19	Holiday Instruction scheduling
2/26	Garbage collection
3/5	Implementing functional and object-oriented languages

The Project

Description

For the class project, you will implement a Java bytecode-to-Java bytecode optimizer, as part of a team of ~3 people. This will provide you with concrete experience in engineering the techniques described in the lectures and readings.

We will use the SOOT Java compiler infrastructure from Laurie Hendren's group at McGill. This infrastructure includes routines to read Java .class files, construct a reasonable internal representation of the contents of the .class file, build a simple control flow graph over 3-address code instructions, and write these representations back into a .class file.

The first task is to design and implement a more advanced intermediate representation, such as def/use chains, SSA form, the PDG, or the VDG,

The second task is to design and implement a general framework for doing arbitrary forward or backward dataflow analyses over your intermediate representation. The framework should allow dataflow analyses to be defined in lattice-theoretic terms, and it should automatically handle the process of applying flow functions and meet operators, manage iteration, and report analysis results back to clients. It must handle irreducible control flow graphs. It need not support automatically composing analyses or transformations (like Vortex's framework), but it should support writing a single analysis with little trouble. It should be at least as efficient as a worklist algorithm that (re)analyzes a node only when some information it depends on changes.

The third task is to write several analyses and optimizations. Intraprocedural optimizations at least will benefit from using your framework.

Requirements

All projects should perform the following optimizations:

- peephole optimizations:
 - constant folding (including branch conditions and eliminating unreachable code)
 - arithmetic simplifications: multiplies to shifts
- intraprocedural optimizations:
 - constant propagation
 - common subexpression elimination or pointer analysis
 - loop-invariant code motion or induction variable elimination
 - dead assignment elimination
- interprocedural optimizations:
 - MOD analysis
 - inlining

Written Reports

Each project team should write a report describing the overall organization, the intermediate representation, the interface and algorithms in the dataflow analysis framework, lattice-theoretic specifications of the intraprocedural analyses, and a brief description of the implementation strategy of the other optimizations. This design document should be written incrementally as the compiler is implemented, and will be reviewed as part of the checkpoint dates. The completed report should be no longer than 15 pages.

Time Schedule and “Deliverables”

Several dates are important for completing the project:

1/12: Project Teams Formed - turn in list of members of each team.

2/5: Intermediate Representation - the compiler should build and generate code from an appropriate intermediate representation; no optimizations necessary. A 2-3 page draft of the project report describing the IR should be turned in, too.

2/21: Data Flow Analysis Framework - the compiler should support a general data flow analysis framework, with constant propagation & folding and dead assignment elimination working using the framework. An extended 5-6 page draft of the project report also describing the analysis framework should be turned in, too.

3/5: Completed Optimizer - the compiler should be augmented with all remaining optimizations.

3/9: Final Project Reports Due - the final version of the report should be completed.

No late projects will be accepted. Projects and reports are due at the beginning of class.