

CS415 Compilers Overview of the Course

These slides are based on slides copyrighted by
Keith Cooper, Ken Kennedy & Linda Torczon at Rice
University

- Instructor: Ulrich Kremer (uli@cs.rutgers.edu)
 - My Office Hours via webex/zoom: TBD
 - Lectures 1, 2, and 3: **online**
- Teaching Assistant
 - TA: Jonathan Garcia-Mullen
 - Recitations & office hours will start next week
- Textbooks:
 - Required: Engineering a Compiler by Cooper/Torcsn
 - Recommended: New (or old) Dragon Book
- Web Site:
 - http://www.cs.rutgers.edu/courses/415/classes/spring_2022_kremer/
 - Project descriptions, handouts, homeworks, lecture slides, ...
 - lec01.pdf vs. lec01mod.pdf
- canvas.rutgers.edu and piazza
 - Homework and project questions, grades, homework solutions
- SPN or prerequisites overrides: Please send me an email

- Exams
 - Midterm 25%
 - Final 35%
- Homeworks 10%
- Projects (tentative!)
 - Back End 10%
 - Front End 10%
 - Code generator/optimizer 10%

Cheating and Honor Code.

Notice: *This grading scheme and projects are tentative and subject to change.*

<ul style="list-style-type: none">• Exams<ul style="list-style-type: none">→ Midterm→ Final	<ul style="list-style-type: none">♦ Midterm in class♦ Final is cumulative (scheduled final exam)
<ul style="list-style-type: none">• Homework	<ul style="list-style-type: none">♦ Reinforce concepts, provide practice♦ Number of assignments t.b.d.
<ul style="list-style-type: none">• Projects<ul style="list-style-type: none">→ Back End→ Front End→ Code generator/opt.	<ul style="list-style-type: none">♦ High ratio of thought to programming♦ Single student labs (note academic integrity information)

- Overview § 1
- Instruction Scheduling § 12
- Local Register Allocation § 13
- Scanning § 2
- Parsing § 3
- Context Sensitive Analysis § 4
- Inner Workings of Compiled Code § 6, 7
- Introduction to Optimization § 8
- Code Selection § 11
- More Optimization (*time permitting*)
 - Advanced topics in language design/compilation:
automatic parallelization, GPUs, power & energy, quantum computing

- Lectures in person in SEC 118
- Get an ilab account NOW, if you not have one already.
- I will use slides
 - I will moderate my speed, *you* sometimes need to say "STOP"
- You should read the book
 - Not all material will be covered in class
 - Book complements the lectures
- You are responsible for material from class
 - The midterm and final will cover both, lectures and readings
- CS 415 is not a programming course
 - Projects are graded on functionality, documentation, and project reports more than style. However, things should be reasonable
- Use the resources provided to you
 - See me or the TA in office hours if you have questions
 - **Post questions regarding homework and projects on piazza**
- Email personal issues to me with subject line starting "cs415:"


Science & Engineering Resource Center - Room 118



Classroom Features

System: Digital Classroom

Classroom Map



Campus: Busch

Address: 118 Frelinghuysen Road,
Piscataway, NJ 08854-8019

Seating Chart: [PDF](#)

- Our required text book is available for free online
 - Available through our course website
 - Available on sakai/Resources
- Other recommended textbook is the “Dragon Book”
 - Aho, Lam, Sethi, Ullman: Compilers - Principles, Techniques, and Tools (2nd edition)
 - Older version (Dragon book) also fine

- What is a **compiler**?
 - A program that translates an *executable* program in one language into an *executable* program in another language
 - A good compiler should improve the program, *in some way*
- What is an **interpreter**?
 - A program that reads an *executable* program and produces the results of executing that program
- C is typically compiled, Scheme is typically interpreted
- Java is compiled to bytecode (code for the Java VM)
 - which is then interpreted
 - Or a hybrid strategy is used
 - Just-in-time compilation
 - Dynamic optimization (hot paths)

- Compilers are important system software components
 - They are intimately interconnected with architecture, systems, programming methodology, and language design
- Compilers include many applications of theory to practice
 - Scanning, parsing, static analysis, instruction selection
- Many practical applications have embedded languages
 - Commands, macros, ...
- Many applications have input formats that look like languages,
 - Matlab, Mathematica
- Writing a compiler exposes practical algorithmic & engineering issues
 - Approximating hard problems; efficiency & scalability
 - No free lunch, i.e., there are multi-dimensional tradeoffs

- Compiler construction involves ideas from many different parts of computer science


<i>Artificial intelligence</i>	Greedy algorithms Heuristic search techniques Machine learning
<i>Algorithms</i>	Graph algorithms, union-find, dynamic programming, approximations
<i>Theory</i>	DFAs & PDAs, pattern matching, fixed-point algorithms
<i>Systems</i>	Allocation & naming, synchronization, data locality
<i>Architecture</i>	Pipeline & hierarchy management, instruction set use, parallelism, quantum computing

- Compiler construction poses challenging and interesting problems:
 - Compilers must do a lot but also **run fast**
 - Compilers have primary responsibility for **run-time performance**
 - Compilers are responsible for making it acceptable to use the **full power** of the programming language
 - Computer architects perpetually create new challenges for the compiler by building more **complex machines** (e.g.: multi-core, GPUs, FPGAs, quantum computers, neuromorphic processors)
 - Compilers must/should hide that complexity from the programmer
 - Success requires mastery of complex interactions

It was our belief that if FORTRAN, during its first months, were to translate any reasonable “scientific” source program into an object program only half as fast as its hand-coded counterpart, then acceptance of our system would be in serious danger... I believe that had we failed to produce efficient programs, the widespread use of languages like FORTRAN would have been seriously delayed.

— John Backus


- My own research
 - Compiling for homogeneous parallel architectures
 - Compiler-directed power and energy management
 - Programming models and languages for dynamic networks of mobile devices (SpatialViews, Sarana)
 - Resource-aware security enforcement (establishing trust)
 - Programming models and languages for Autonomous Underwater Vehicles (AUVs)
 - Programming models for program quality/resources tradeoffs
 - Programming models for non-von Neumann machines (e.g. quantum computing, reversible computing, neuromorphic computing)
- Thus, my interests lie in
 - Interplay between compiler, OS, and architecture
 - Static analysis to discern program behavior
 - Run-time performance analysis
 - Physical measurements (power/energy, performance, memory)
 - New ways of thinking about computation



Implications

- Must recognize legal (and illegal) programs
- Must generate correct code
- Must manage storage of all variables (and code)
- Must agree with OS & linker on format for object code


Big step up from assembly language—use higher level notations



Implications

- Use an intermediate representation (IR)
- Front end maps legal source code into IR
- Back end maps IR into target machine code
- Extension: multiple front ends & multiple passes (*better code*)


Typically, front end is $O(n)$ or $O(n \log n)$, while back end is NP-complete



Can we build $n \times m$ compilers with $n+m$ components?


- Must encode all language specific knowledge in each front end
- Must encode all features in a single IR
- Must encode all target specific knowledge in each back end

Limited success in systems with very low-level IRs




Responsibilities

- Recognize legal (& illegal) programs
- Report errors in a useful way
- Produce IR & preliminary storage map
- Shape the code for the back end
- Much of front end construction can be automated



Scanner

- Maps character stream into words—the basic unit of syntax
- Produces pairs — a word & its part of speech
 $x = x + y ;$ becomes $\langle \text{id}, x \rangle = \langle \text{id}, x \rangle + \langle \text{id}, y \rangle ;$
 - $\text{word} \cong \text{lexeme}$, $\text{part of speech} \cong \text{token type}$
 - In casual speech, we call the pair a *token*
- Typical tokens include *number, identifier, +, -, new, while, if*
- Scanner eliminates white space (including comments)
- Speed is important



Parser

- Recognizes context-free syntax & reports errors
- Guides context-sensitive ("semantic") analysis (*type checking*)
- Builds IR for source program

Hand-coded parsers are fairly easy to build

Most books advocate using automatic parser generators

Context-free syntax is specified with a grammar

$$\textit{SheepNoise} \rightarrow \textit{SheepNoise} \ \underline{\textit{baa}} \\ | \ \underline{\textit{baa}}$$

This grammar defines the set of noises that a sheep makes under normal circumstances

It is written in a variant of Backus-Naur Form (BNF)

Formally, a grammar $G = (S, N, T, P)$

- S is the *start symbol*
- N is a set of *non-terminal symbols*
- T is a set of *terminal symbols* or *words*
- P is a set of *productions* or *rewrite rules* $(P : N \rightarrow N \cup T)$

Context-free syntax can be put to better use

1. $goal \rightarrow expr$
2. $expr \rightarrow expr\ op\ term$
3. | $term$
4. $term \rightarrow \underline{number}$
5. | \underline{id}
6. $op \rightarrow +$
7. | $-$

$S = goal$

$T = \{ \underline{number}, \underline{id}, +, - \}$

$N = \{ goal, expr, term, op \}$

$P = \{ 1, 2, 3, 4, 5, 6, 7 \}$

- This grammar defines simple expressions with addition & subtraction over "number" and "id"
- This grammar, like many, falls in a class called "context-free grammars", abbreviated CFG

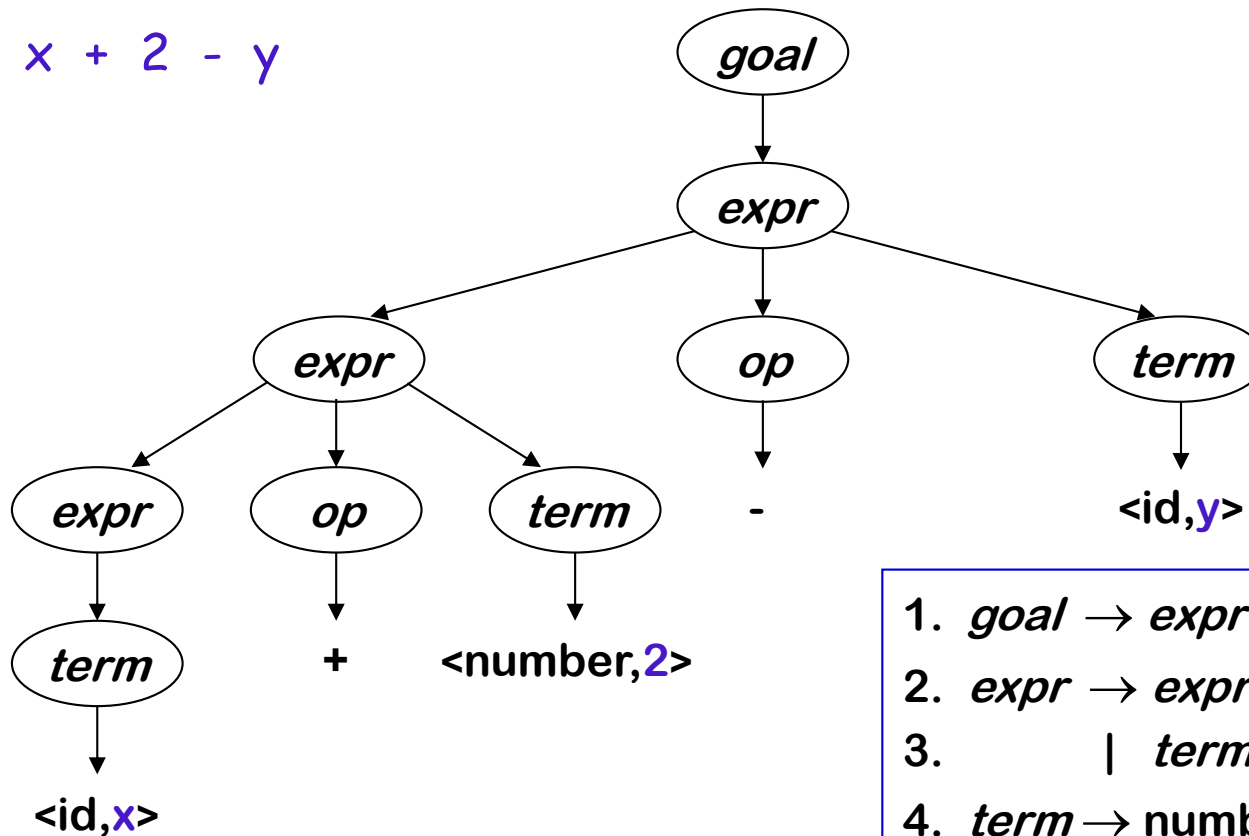
Given a CFG, we can *derive* sentences by repeated substitution

<u>Production</u>	<u>Result</u>	
	<i>goal</i>	\Rightarrow
1	<i>expr</i>	\Rightarrow
2	<i>expr op term</i>	\Rightarrow
5	<i>expr op y</i>	\Rightarrow
7	<i>expr - y</i>	\Rightarrow
2	<i>expr op term - y</i>	\Rightarrow
4	<i>expr op 2 - y</i>	\Rightarrow
6	<i>expr + 2 - y</i>	\Rightarrow
3	<i>term + 2 - y</i>	\Rightarrow
5	<i>x + 2 - y</i>	

To recognize a valid sentence in some CFG, we will need to construct a derivation automatically (forwards or backwards)

A parse can be represented by a tree (*parse tree* or *syntax tree*)


$x + 2 - y$



This contains a lot of unneeded information.

1. $goal \rightarrow expr$
2. $expr \rightarrow expr \ op \ term$
3. | $term$
4. $term \rightarrow \underline{number}$
5. | \underline{id}
6. $op \rightarrow +$
7. | $-$

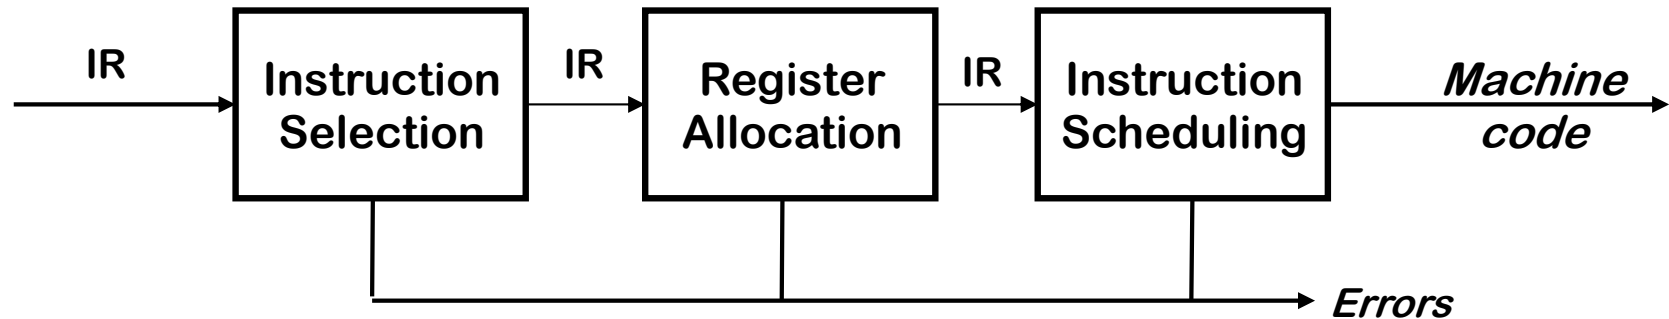
Compilers often use an *abstract syntax tree (AST)*



The AST summarizes grammatical structure, without including detail about the derivation

This is much more concise

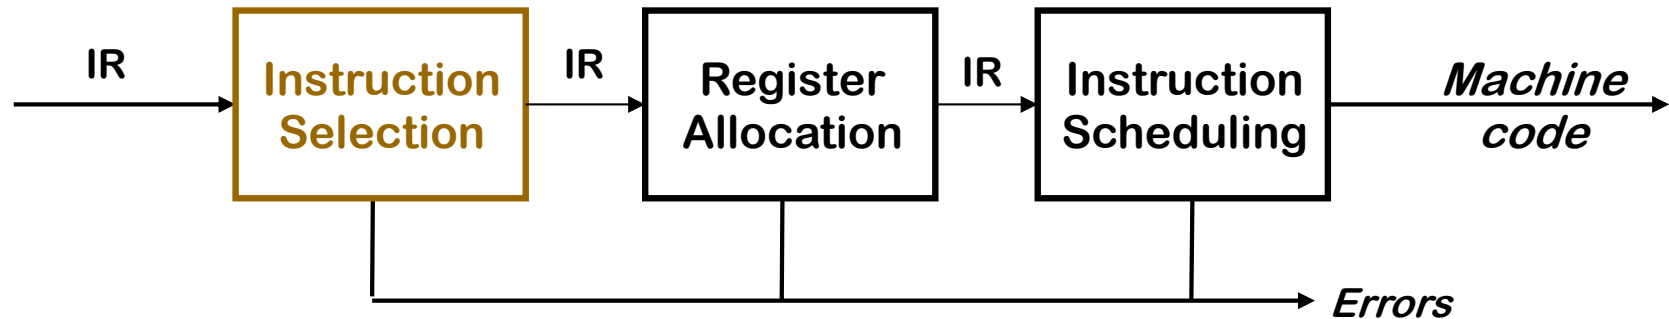
ASTs are one kind of *intermediate representation (IR)*



Responsibilities

- Translate IR into target machine code
- Choose instructions to implement each IR operation
- Decide which value to keep in registers
- Ensure conformance with system interfaces

Automation has been *less* successful in the back end




Instruction Selection

- Produce fast, compact code
- Take advantage of target features such as addressing modes
- Usually viewed as a pattern matching problem
 - *ad hoc* methods, pattern matching, dynamic programming

This was the problem of the future in 1978

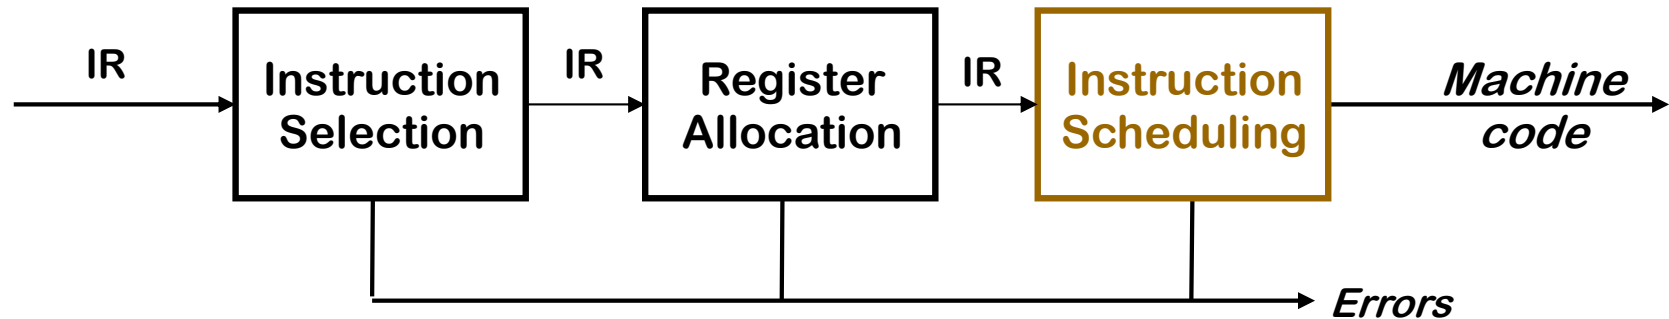
- Spurred by transition from PDP-11 to VAX-11
- Orthogonality of RISC simplified this problem



Register Allocation

- Have each value in a register when it is used
- Manage a limited set of resources
- Select appropriate LOADs & STOREs
- Optimal allocation is NP-Complete (1 or k registers)

Typically, compilers approximate solutions to NP-Complete problems

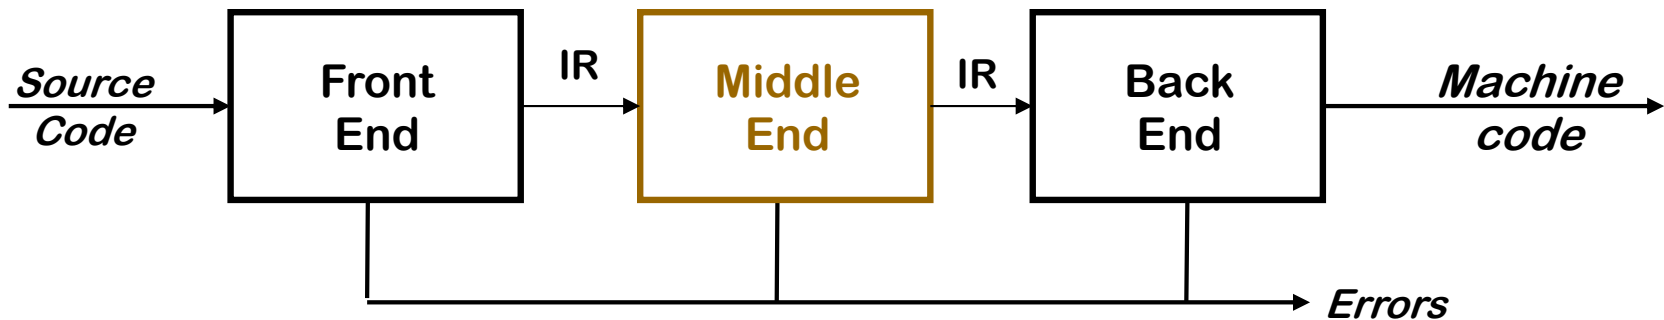


Instruction Scheduling

- Avoid hardware stalls and interlocks
- Use all functional units productively
- Can increase lifetime of variables (changing the allocation)

Optimal scheduling is NP-Complete in nearly all cases

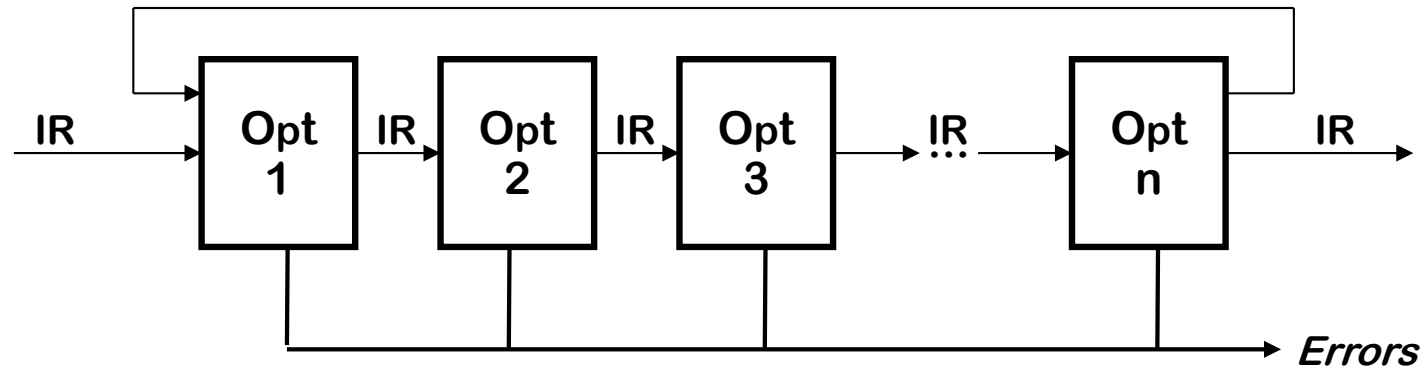
Heuristic techniques are well developed



Code Improvement (or Optimization)

- Analyzes IR and rewrites (or transforms) IR
- Primary goal is to reduce running time of the compiled code
 - May also improve space, power dissipation, energy consumption, ...
- Must preserve "meaning" of the code (may include approximations, i.e., quality of outcomes tradeoffs)
 - Measured by values of named variables or produced output

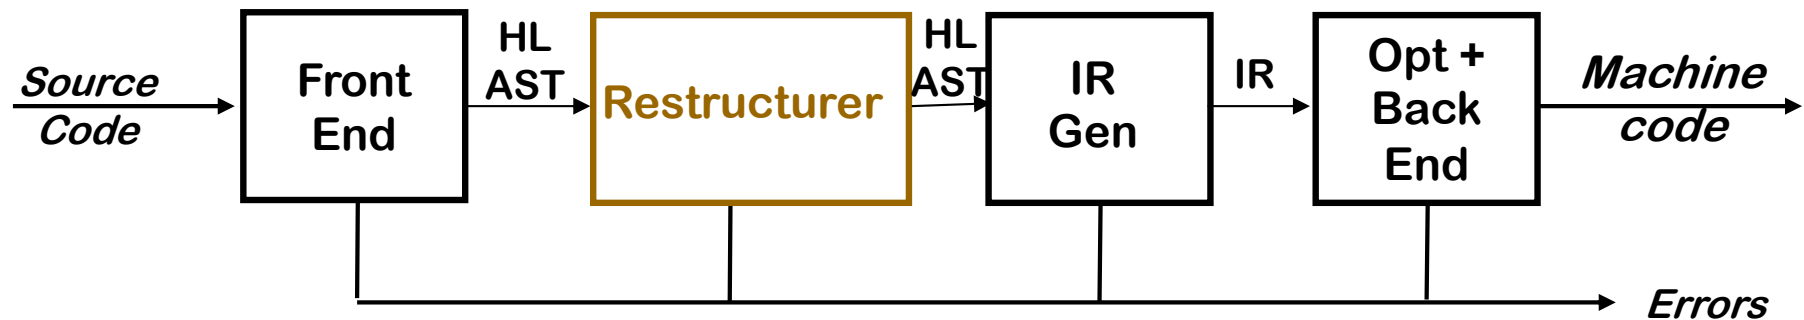
Subject of cs515, and cs516, maybe final weeks of cs415



Modern optimizers are structured as a series of passes

Typical Transformations

- Discover & propagate some constant value
- Move a computation to a less frequently executed place
- Specialize some computation based on context
- Discover a redundant computation & remove it
- Remove useless or unreachable code
- Encode an idiom in some particularly efficient form



Typical **Restructuring** (source-to-source) Transformations:

- Blocking for memory hierarchy and register reuse
- Vectorization
- Parallelization
- All based on dependence
- Also full and partial inlining

- Memory management services
 - Allocate
 - In the heap or in an activation record (*stack frame*)
 - Deallocate
 - Collect garbage
- Run-time type checking
- Error processing (exception handling)
- Interface to the operating system
 - Input and output
- Support of parallelism
 - Parallel thread initiation
 - Communication and synchronization

Instruction Scheduling

Read EaC: Chapters 12.1 - 12.3

No recitations or office hours this week