	Introduction to LLVM
Compiler construction 2011	Register machines
Lecture 5 • Introduction to LLVM. • LLVM language and tools.	Fast but scarce Registers are places for data inside the CPU. + up to 10 limes faster access than to main memory expensive; typically just 32 of them in a 32-bit CPU. Typically, arithmetic operations, conditional jumps etc operate on values stored in registers. Most modern assembly languages use registers, which correspond closely to the machine registers. LLVM (the Low Level Virtual Machine) LLVM is a virtual machine: it has an unbounded number of registers. A later step does register allocation, mapping virtual registers to real machine registers.
Introduction to LLVM	Introduction to LLVM
The LLVM project	The LLVM language
The LLVM project The LLVM infrastructure A collection of (C++) software libraries and tools to help in building compilers, debuggers, program analysers, etc. Tools available on Studat Linux machines. Can also be downloaded to your own computer. Visit 11vm. org. History Starled as academic project at University of Illinois at Urbana-Champaign 2002. Current development mainly at Apple. Growing user base. Related projects • Clang. C/C++ front end; aims to replace gcc.	The LLVM language Characteristic features Three adress-code: two source registers and one destination register: %t2 = add i32 %t0, %t1 One source can be a value: %t5 = add i32 %t3, 7 Instructions are typed: %t8 = fadd double %t6, %t7 store i32 %t5, i32* %tr New register for each result (Static Single Assignment form).

Introduction to LLVM		Introduction to LLVM
Hello world in LLVM		An illegal LLVM program
<pre>@hw = internal constant [13 x i8] c"hello declare i32 @puts(i8+) define i32 @main () { entry: Xt1 = bitcast [13 x i8]* @hw to i6 Xt2 = call i32 @puts(i8* Xt1) ret i32 Xt2 } Comments o String is named @hw, a global constant (global Note escape sequences! o Library function @puts is declared, giving type o @hw is cast to type of argument to puts. Note: Better (type-sale) solution later!</pre>	o world\0A\00" 3* I names start with ©). e signature.	<pre>declare void @printInt(i32 %n) define i32 @main() { entry: %t! = call i32 @sum(i32 100) call void @printInt(i32 %t1) ret i32 0 } define i32 @sum (i32 %n) { entry: %sum = i32 0 %t = i32 0 %t = add i32 %t, 1 %sum = add i32 %t, 1 %sum = add i32 %um, %i %t = icmp eq i32 %i, %n br ii %t, label %end, label %labi } </pre>
Introduction to LLVM		Introduction to LLVM
Corrected program		Optimizing @sum
<pre>define i32 @sum (i32 %n) { entry: %sum = alloca i32 store i32 0, i32+ %sum %i = alloca i32 store i32 0, i32+ %i br label %labl labl: %t1 = load i32+ %i %t2 = add i32+ %sum %t4 = add i32 %t2, %t3 store i32 %t2, i32+ %i store i32 %t2, i32+ %i store i32 %t2, i32+ %i %t5 = icmp eq i32 %t2, %n br i1 %t5, label %end, label %lab) end: ret i32 %t4 }</pre>	Comments • ½ and ½sum are now pointers to memory locations. • Only one assignment to any register. Problem This program has a lot more memory traffic! What can LLVM's optimizer do about that?	<pre>> opt -mem2reg sum.bc > sumreg.bc > llvm-dis sumreg.bc > more sumreg.ll define i32 @sum[i32 %n) { entry: br label %lab1 lab1: %i.0 = phi i32 [0, %entry], [%t2, %lab1] %sum.0 = phi i32 [0, %entry], [%t4, %lab1] %t2 = add i32 %l0, 0, 1 %t4 = add i32 %t2, %sum.0 %t5 = icmp eq i32 %t2, %sum.0 br i1 %t5, label %end, label %lab1 end: ret i32 %t4 }</pre>

Introduction to LLVM	Introduction to LLVM
Φ "functions"	Optimizing the program further
SSA form Only one assignment in the program text to each variable. (But dynamically, this assignment can be executed many times). (But dynamically, this assignment can be executed many times). (Description of the assignment of	<pre>Many optimization passes opt implements many code analysis and improvement methods. To get a default selection, give command line arg -std-compile-opts. Result, selection, give command, selection, give arg -std-compile-opts. Result, s</pre>
Introduction to LLVM	Introduction to LLVM
Personanties is LLUM Optimizing sum further	Analysis of optimized code for @sum
Pendedevisitive Optimizing sum further Result after opt -std-compile-opts define i32 @sum(i32 %n) nouwrind readnone { entry: %tmp3 = add i32 %n, -2 %tmp1 = add i32 %n, -1 %tmp4 = zext i32 %tmp3 to i33	
<pre>Peakative LINW Optimizing sum further Result after opt -std-compile-opts define i32 @sum(i32 %n) nounwind readnone { entry: %tmp3 = add i32 %n, -2 %tmp1 = add i32 %n, -1 %tmp4 = zext i32 %tmp1 to i33 %tmp5 = mul i33 %tmp2, %tmp4 %tmp6 = lahr i33 %tmp6 to i32 %tmp = shl i32 %n, 1 %tmp7 = shl i32 %n, 1 %tmp8 = add i32 %tmp6 to i32 %tmp9 = add i32 %tmp8, -1 r ti 32 %tmp9 }</pre>	Previous loop with execution time O(n) has been optimized to code without loop, running in constant time. Recall 1 + 2 + + n = n(n + 1)/2. Check that optimized code computes this. Why extensions/runcations to and from 33 bits? What happens when n is negative? Opt -std-compile-opts includes many optimization passes. Use -time-passes for an overview. We will discuss some of these algorithms later.

Introduction to LLVM	Introduction to LLVM
printInt and other IO functions	Linking and running the program
<pre>Part of runtime.ll @dnl = internal constant [4 x i8] c"%d\0A\00" declare i32 @printf(i8*,) define void @printInt(i32 %x) { entry: %t0 = getelementptr [4 x i8]* @dnl, i32 0, i32 0 call i32 (i8*,)* @printf(i8* %t0, i32 %x) ret void } We provide this file on the course web site; you just have to make sure that it is available for linking.</pre>	<pre>Linker is llvm-ld > llvm-ld sumopt.bc runtime.bc > ./a.out 5050 > more a.out #!/bin/sh exec lli a.out.bc \${1+"\$@"} So, linking produces two files:</pre>
	CHALMER
Introduction to LLVM	LLVM language and tools
What is in a.out.bc	Types in LLVM
Disassemble it!	
>cat a.out.bc llvm-dis -	An incomplete list
; ModuleID = 'a.out.bc'	Below t and t are types and n an integer literal
target datalayout = "e- most of line skipped	a n bit integers: i n
	a float and double
@dnl = internal constant [4 x 18] c"%d\0A\00"	blobole: label
define i32 @main() {	• Labels. Tabel.
entry:	o The void type: void.
%t0 = getelementptr [4 x i8]* @dnl, i32 0, i32 0	• Functions: $t(t_1, t_2, \ldots, t_n)$.
call i32 (i8*,)* @printf(i8* %t0, i32 5050)	 Pointer types: t*.
ret i32 0	Structures: { t ₁ , t ₂ ,, t _n }.
1	
3	Arrays : [n x t].

LLVM language and tools	LLVM language and tools
Named types and type equality	Identifiers
Named types	
One can give names to types. Examples: %length = type i32 %list = type { i32, %list } * %node = type { %tree, i32, %tree } %tree = type %node * %matrix = type [100 x [100 x double]] Type equality LLVM uses structural equality for types. When disassembling bitcode files that contain several structurally equal types with different names, this may give confusing results.	Local identifiers Registers and named types have local names, starting with %. Global identifiers Functions and global variables have global names, starting with 0. Javalette does not have global variables, but you will need to define global names for string literals, as in 0hw = internal constant [13 x i8] c"hello world\OA\OO" After this definition, 0hw has type [13 x i8] *.
	CHALMERS
LLV/M language and tools	LLVM language and tools
Litterals o Integer and floating-point literals are as expected.	Live large are task Function definitions Simplest form define t gname $(t_1 x_1, t_2 x_2,, t_n x_n)$ { block, block,
Litterals o Integer and floating-point literals are as expected. • true and false are literals of type i1.	LUVE larguage are toos Function definitions Simplest form define I gname (I ₁ x ₁ , I ₂ x ₂ ,, I _n x _n) { block; block;
Littlegaap and toot Constants Literals o Integer and floating-point literals are as expected. o true and false are literals of type i1. o ull is a literal of any pointer type.	LUTH language and tools Function definitions Simplest form define <i>t</i> gname (<i>t</i> ₁ <i>x</i> ₁ , <i>t</i> ₂ <i>x</i> ₂ ,, <i>t</i> _n <i>x_n</i>) { block ₁ block ₂ block _n
Litterals • Integer and floating-point literals are as expected. • true and false are literals of type i1. • null is a literal of any pointer type. Aggregates Constant expressions of structure and array types can be formed; not needed by Javalette.	LUM large approximation Function definitions Simplest form define t gname (t ₁ x ₁ , t ₂ x ₂ ,, t _n x _n) { block, block, blockn, } where gname is a global name (the name of the function), the x _i are local names (the parameters) and the block _i are basic blocks. Basic blocks A basic block is a label followed by a colon and a sequence of LLVM instructions, each on a separate line. The last instruction must be a terminator instruction.

LLVM language and tools	LLVM language and tools
Function declarations	Data layout
Type-checking The LLVM assembler does type-checking. Hence it must know the types of all external functions, i.e. functions used but not defined in the compiled unit. Simple function declaration The basic form is declarer t gname (t_1, t_2, \dots, t_n) For Javalette, this is necessary for IO functions. The compiler would typically insert in each file declare void @printDuble(double) declare void @printDuble(double) declare void @printDuble(double) declare void @printDuble(double)	Layout specifications An LLVM module may specify how data is laid out in memory on the target architecture (endianness, size of pointers, alignment, etc.) We will generate 32 bit code, and your compiler should generate a target layout description as the first line of every generated file. This is one very long line; broken here to fit on the slide. target datalayout = "e-p:32:32:32-11:8:8-18:88- 16:16:16:16-132:32:32-164:32:64-f32:32:32-f4:32:64- v64:64:64-v128:128:128-a0:0:64-f80:32:32-a8:16:32"
declare double GreadDouble()	
CHALMERS	CHALMERS
LLV/M language and tools	LLVM language and tools
LUM transmiss and tools LLVM tools	International sections Input/output functions for Javalette
LILWM tools • The assembler llvm-as. Translates to bitcode (prog.11 to prog.bc). • The disassembler llvm-dis. Translates in the opposite direction. • The disassembler llvm-dis. Translates in the opposite direction. • The interpreter/JIT compiler lli. Executes bitcode file containing a main function. • The linker llvm-ld. Links together several bitcode files and produces a.out. bc and a small script a.out, which calls lli on a.out. bc. • The compiler llc. Translates to native assembler. • The optimizer opt. Optimizes bitcode; many options to decide on which optimizations to run. Use -std-compile-opts to get a default selection. • Drop-in replacement for gcc: clang.	<pre>tutM request set tow Input/output functions for Javalette We provide runtime.ll with a simple implementation of the IO functions (based on functions in C's stdio). You only need to assemble to runtime.bc. Sample part of runtime.ll @dn1 = internal constant [4 x i8] c"%d\0A\00" declare i32 @printf(i8*,) define void @printInt(i32 %x) { entry: %t0 = getelementptr [4 x i8]* @dnl, i32 0, i32 0 call i32 (i8*,)* @printf(i8* %t0, i32 %x) ret void } </pre>

LLVM language and tools	LLVM language and tools
Use of LLVM in your compiler	LLVM instructions
Default mode Your code generator produces assembler file (.11). Then your main program uses system calls to first assemble this with llvm-as and then link together with runtime.bc.	Basic collection Basic Javalette will only need the following instructions: • Terminator instructions: ret and br. • Arithmetic operations:
Other modes More advanced; we do not recommended these for this project. • C++ programmers can use the LLVM libraries to build in-memory representation and then output bitcode file. • Haskell programmers can access C++ libraries via Hackage package LLVM.	 For integers add, sub, mul, sdiv and srem. For doubles fadd, fsub, fmul and fdiv. Memory access: alloca, load, getelementptr and store. Other: iccmp, fcmp and call. Some of the extensions will need more.
CHAIMERS	CHALMERS