



CHALMERS

Compiler construction

Lecture 5: Project extensions

Thomas Sewell

Spring 2020

Chalmers University of Technology — Gothenburg University

Some project extensions:

- Arrays
- Pointers and structures
- Object-oriented languages
- Module system (extension proposal)

Also:

- Stack, Heap & Memory Management
- Some compiler technology

Next time:

- Native (e.g. x86) code

We've covered how to compile the standard JAVALETTE language to LLVM (Part B).

Part C: you must implement at least one extension.
(grading scheme on course syllabus page)

This lecture is an overview of some extensions. More details in the tda283 repository in `project/extensions.md`

Some extensions are more clearly specified than others. You may propose your own extensions, contact us.

Arrays

JAVALETTE restrictions

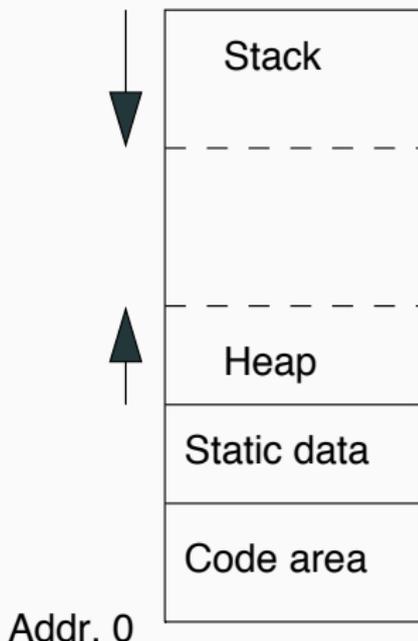
- Only local variables and parameters – no global variables or other non-local declarations
- Only simple data types (`int`, `double`, `boolean`) with fixed-size values
- Only call-by-value parameter passing

All data at runtime can be stored in the activation records on a stack (`alloca`).

More general language features

Global variables, nested procedures, linked structures and pointers, classes, ..., have other requirements: a stack is not sufficient for runtime memory management.

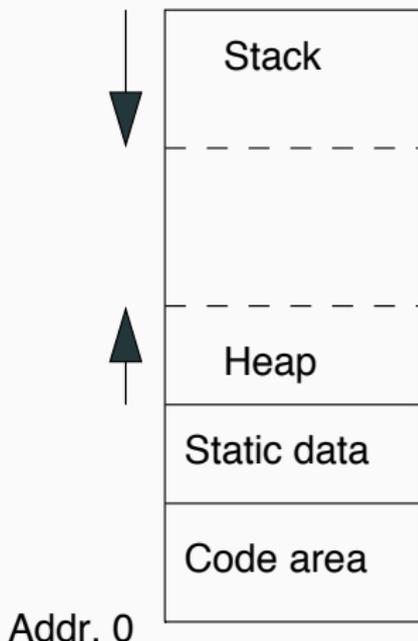
Classic memory organisation



Comments

- Static area for global variables
- Stack and heap grow from opposite ends to avoid predetermined size decisions
- Stack managed by LLVM
- Management of heap much more complicated than stack

Classic memory organisation



Comments

- Static area for global variables
- Stack and heap grow from opposite ends to avoid predetermined size decisions
- Stack managed by LLVM
- Management of heap much more complicated than stack

More comments: this diagram is a bit old, e.g. multi-threading requires multiple smaller stacks, heap can be disjointed.

Java-like arrays

Arrays in the JAVALETTE extensions are similar to Java arrays:

- A variable of type e.g. `int[]` contains a reference to a block of memory on the heap, where array elements are stored.
- Arrays must be explicitly created as in

```
int[] v = new int[20];
```
- Array elements are initialized to `0/0.0/false`
- Arrays have a `length` attribute, with dot-notation
- Arrays can be both function arguments and results

Two extensions (each one credit)



First extension

One-dimensional arrays and 'foreach' statement, as in

```
int sum = 0;
for (int x : v)
    sum = sum + x;
```

The ordinary `for` statement is not required.

Two extensions (each one credit)



First extension

One-dimensional arrays and 'foreach' statement, as in

```
int sum = 0;
for (int x : v)
    sum = sum + x;
```

The ordinary `for` statement is not required.

Second extension

Multidimensional arrays:

- For $n > 1$, an n -dimensional array is a one-dimensional array, each of whose elements is an $n - 1$ -dimensional array

```
int[] [] matrix = new int[10][10];
```

Indexing in two-dimensional array

```
%arr1 = type %struct1*
%arr2 = type %struct2*
%struct1 = type {i32, [0 x i32]}
%struct2 = type {i32, [0 x %arr1]}

define i32 @getElem (%arr2 %m, i32 %i, i32 %j) {
    %p1 = getelementptr %struct2, %arr2 %m,
        i32 0, i32 1, i32 %i
    %p2 = load %arr1, %arr1* %p1
    %p3 = getelementptr %struct1, %arr1 %p2,
        i32 0, i32 1, i32 %j
    %p4 = load i32, i32* %p3
    ret i32 %p4
}
```

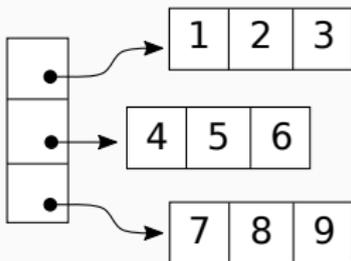
Your generated code may well look different.

First extension

- LLVM type of array hinted at in previous slide
 - The length 0 array means variable size
 - Only useful at the end of a structure
- Use C function `calloc` to allocate 0-initialized memory
- New forms of expression: array indexing, `new` expression, `length`
- Array indexing also allowed in L-values: `arr[1] = x;`
- No need for bounds-checking code
- Example tests in `tester/testsuite/extensions/arrays1`
 - Look out for `arr [1]; arr . length;`

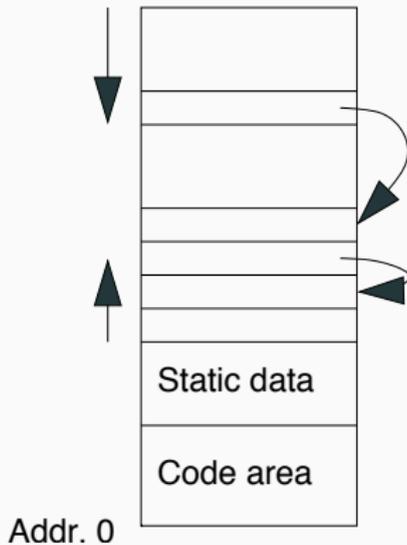
Second extension

- Updating a row is permitted:
`row[1] = 1; matrix[1] = row;`
- Multidimensional arrays are pointer-linked (unlike C).
- `x[1][1][2]` requires multiple `getelementptr` instructions
- `new int[4][4][4]` requires generating code with multiple loops and multiple `calloc`'s



Structures and pointers

- In addition to function definitions, a Javalette file may contain definitions of structures and pointer types
- Structure objects are allocated on the heap, using `new`
- Pointer variable (on stack) may refer to memory structure on the heap



Code examples in JAVALETTE.

```
typedef struct Node *list;

struct Node {
    int elem;
    list next;
}

list cons(int x, list xs) {
    list res;
    res = new Node;
    res->elem = x;
    res->next = xs;
    return res;
}
```

```
int length(list xs) {
    if (xs == (list)null)
        return 0;
    else
        return 1 + length(xs->next);
}

list fromTo(int m, int n) {
    if (m > n)
        return (list)null;
    else
        return cons(m, fromTo(m + 1, n));
}
```

Code examples in JAVALETTE.

```
typedef struct Node *list;

struct Node {
    int elem;
    list next;
}

list cons(int x, list xs) {
    list res;
    res = new Node;
    res->elem = x;
    res->next = xs;
    return res;
}
```

```
int length(list xs) {
    if (xs == (list)null)
        return 0;
    else
        return 1 + length(xs->next);
}

list fromTo(int m, int n) {
    if (m > n)
        return (list)null;
    else
        return cons(m, fromTo(m + 1, n));
}
```

(and more in tester/testsuite/extensions/pointers)

New toplevel definitions

- Structure definitions, exemplified by `Node`
- Pointer type definitions, exemplified by `list`

New expression forms

New statement forms

New toplevel definitions

- Structure definitions, exemplified by `Node`
- Pointer type definitions, exemplified by `list`

New expression forms

- Heap object creation, exemplified by `new Node`
- Pointer dereferencing, exemplified by `xs->next`
- Null pointers, exemplified by `(list)null`

New statement forms

New toplevel definitions

- Structure definitions, exemplified by `Node`
- Pointer type definitions, exemplified by `list`

New expression forms

- Heap object creation, exemplified by `new Node`
- Pointer dereferencing, exemplified by `xs->next`
- Null pointers, exemplified by `(list)null`

New statement forms

- Pointer dereferencing allowed in left hand sides of assignments, as in `xs->elem = 3;`
- Since there is no garbage collection, you should have a `free` statement

Some hints

- Structure and pointer type definitions translate to LLVM type definitions
- Again, use `calloc` for allocating heap memory
 - No loops as `null` now present in language.
- `getelementptr` and `load` will be used for pointer dereferencing
- Info about struct layout may be needed in the state of code generator

Some hints

- Structure and pointer type definitions translate to LLVM type definitions
- Again, use `calloc` for allocating heap memory
 - No loops as `null` now present in language.
- `getelementptr` and `load` will be used for pointer dereferencing
- Info about struct layout may be needed in the state of code generator

From previous lecture: Computing the size of a type

We use the `getelementptr` instruction:

```
%p = getelementptr %T, %T* null, i32 1  
%s = ptrtoint %T* %p to i32
```

Now, `%s` holds the size of `%T`.

Code example (in C)

```
void swap (int *x, int *y) {  
    int tmp = *x;  
    *x = *y;  
    *y = tmp;  
}
```

```
int main () {  
    int a = 1;  
    int b = 3;  
    swap(&a, &b);  
    printf("a=%d\n", a);  
}
```

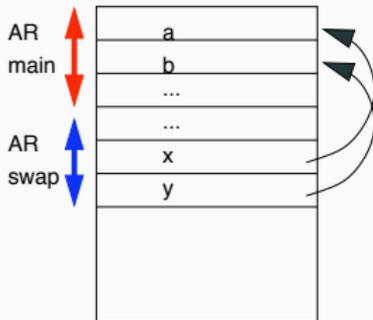
Code example (in C)

```
void swap (int *x, int *y) {  
    int tmp = *x;  
    *x = *y;  
    *y = tmp;  
}
```

```
int main () {  
    int a = 1;  
    int b = 3;  
    swap(&a, &b);  
    printf("a=%d\n", a);  
}
```

Parameter passing by reference

- To make it possible to return results in parameters, one may use pointer parameters
- Actual arguments are addresses
- Problem: makes code optimization much more difficult



Code examples

```
void swap (int *x, int *y) {  
    int tmp = *x;  
    *x = *y;  
    *y = tmp;  
}
```

```
...  
swap (x, x);  
...
```

Code examples

```
void swap (int *x, int *y) {  
    int tmp = *x;  
    *x = *y;  
    *y = tmp;  
}
```

```
...  
swap (x, x);  
...
```

Comments

- With call by reference and pointers, two different variables may refer to the same location; aliasing
- Aliasing complicates code optimization:

```
x := 2  
y := 5  
a := x + 3
```

Here we might want to replace last instr by `a := 5;` but what if `y` is an alias for `x`?

The problem

In contrast to stack memory, there is no simple way to say when heap allocated memory is not needed anymore.

Two main approaches

1. Explicit deallocation

- Programmer deallocates memory (using `free`)
- Potentially most efficient
- Very easy to get wrong (memory leakage or premature returns)

2. Garbage collection

- Programmer does nothing; runtime system reclaims unneeded memory
- Secure but potential runtime penalty
- Acceptable in most situations
- Used in Java, Haskell, C#, ...

General approach

Runtime system keeps list(s) of free heap memory, `malloc` returns a chunk from suitable free list.

Many variations. Some approaches:

1. Reference counting: each chunk keeps a reference count of incoming pointers, when count becomes zero, chunk is returned to free list; problem: cyclic structures
2. When free list is empty, collect in two phases:
 - Mark** Follow pointers from global and local variables, marking reachable chunks
 - Sweep** Traverse heap and return unmarked chunks to free list

Folklore: C is faster than Haskell/Java is faster than Python etc.

Also: `malloc/free` is faster than reference counting is faster than `mark/sweep`.

But ... it's not so simple. Python uses reference counting. Active research on garbage collection shows that `mark/sweep` can be faster than the other styles in many cases.

A more accurate version might be, it's hard to write C programs with complex memory allocation, easier for programs where memory allocation can be reduced to simple rules, and such programs are faster.

Mark/Sweep is standard and essentially necessary for a garbage collector.

It's possible to speed garbage collection up though by switching between fast/local collections of recently allocated objects (the current generation) and more thorough but slower garbage collection operations.

Object-orientation

Class-based languages

We consider only languages where objects are created as instances of classes. A class describes:

- a collection of instance variables; each object gets its own copy of this collection
- a collection of methods to access and update the instance variables

Usually each object contains, in addition to the instance variables, a pointer to a class descriptor. This descriptor contains addresses of the code of methods.

Without inheritance, all this is straightforward; classes are just structures (this design is actually used by e.g. Linux). We propose a little bit more: single inheritance without method override.

```
class Counter {  
    int val;  
  
    void incr () {  
        val++;  
        return;  
    }  
  
    int value () {  
        return val;  
    }  
}
```

```
int main () {  
    Counter c;  
    c = new Counter;  
    c.incr();  
    c.incr();  
    c.incr();  
    int x = c.value();  
    printInt(x);  
    return 0;  
}
```

```
class Point2 {  
    int x;  
    int y;  
  
    void move(int dx, int dy) {  
        x = x + dx;  
        y = y + dy;  
    }  
  
    int getX() { return x; }  
  
    int getY() { return y; }  
}
```

(more in
testsuite/extensions/objects*)

```
class Point3 extends Point2 {  
    int z;  
    void moveZ(int dz) {  
        z = z + dz;  
    }  
    int getZ() { return z; }  
}  
  
int main() {  
    Point2 p;  
    Point3 q = new Point3;  
  
    q.move(2, 4);  
    q.moveZ(7);  
    p = q;  
    ...  
}
```

New toplevel definitions

- Class definitions, consisting of a number of instance variable declarations and a number of method definitions
- Instance variables are only visible within methods of the class
- All methods are public
- All classes have one default constructor, which initializes instance variables to default values (0, false, null)
- A class may extend another one, adding more instance variables and methods, but without overriding
- Classes are types; variables can be declared to be references to objects of a class

New forms of expressions

- Object creation, exemplified by `new Point2`, which allocates a new object on the heap with default values for instance variables
- Method calls, exemplified by `p.move(3,5)`
- Null references, exemplified by `(Point)null`
- Self reference. Within a class, `self` is a pointer that can be used to refer to the current object or call a sibling method.

New scopes and type rules

- Instance variables share syntax with locals/parameters, but have different semantics.
- We have subtyping; if `S extends C`, then `S` is a subtype of `C`; whenever an object of type `C` is expected, we may supply an object of type `S`.

Implementation hints

- Most ideas (and code) from the pointers/structures extension can be reused.
- Method calls will be translated to function call with receiving object (`self` reference) as extra, first parameter.
- Subtyping requires memory compatibility. If `S extends C`, we want the `S` structure to contain a compatible `C` structure.

Example with overriding

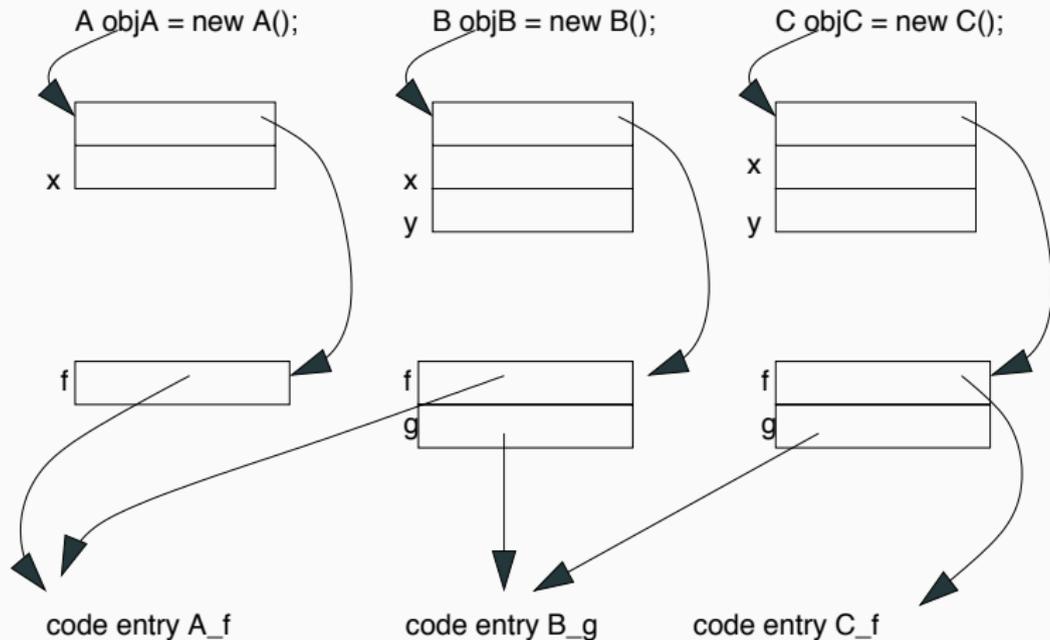
We consider the following classes:

```
class A {  
    int x;  
    void f (int z) {x = z;}  
}
```

```
class B extends A {  
    int y;  
    int g() {return 0;}  
}
```

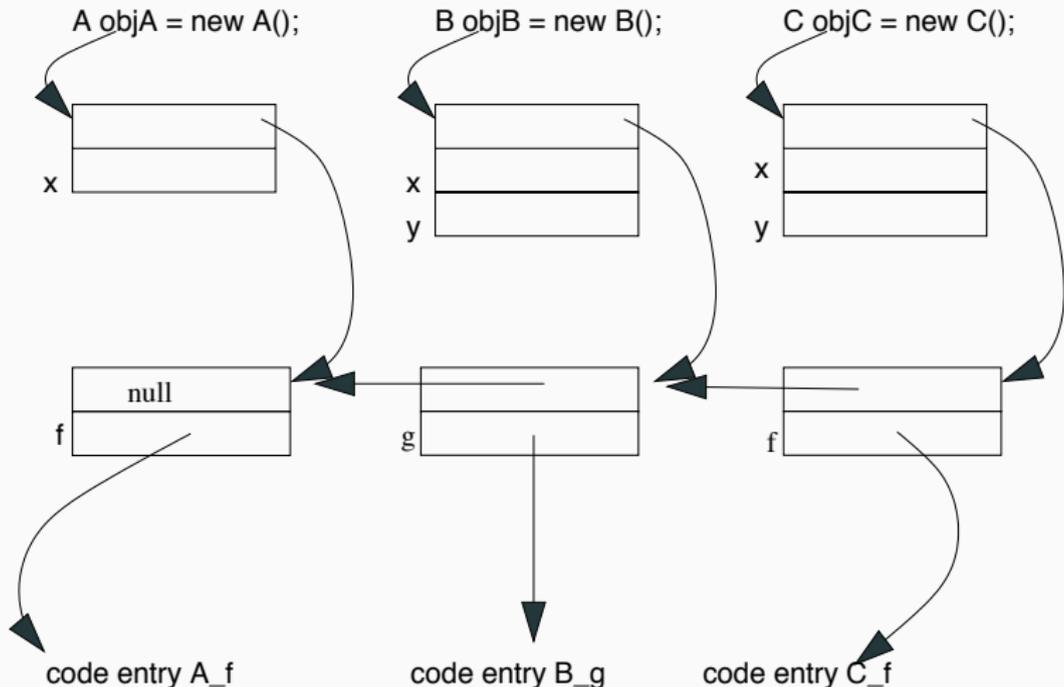
```
class C extends B {  
    void f(int z) {x = z; y = 0;}  
}
```

Objects at runtime



Each object includes instance variables and pointer to class descriptor.

Objects at runtime, variant



Class descriptors linked into list. List searched at runtime.

Code example

```
A obj = objA;
```

```
...
```

```
obj.f(5);
```

```
...
```

```
obj = objC;
```

```
...
```

```
obj.f(3);
```

Code example

```
A obj = objA;  
...  
obj.f(5);  
...  
obj = objC;  
...  
obj.f(3);
```

What code is run?

- The method to execute is determined at runtime by following the link to the class descriptor
- Static type checking guarantees that there is a method with proper signature in the descriptor
- There is an efficiency penalty in dynamic dispatch (so optimization tries to remove it (not required in your implementation))

Modules

Programmer's perspective: Modularity

- Reusability
- Information hiding
- Name space control

Programmer's perspective: Modularity

- Reusability
- Information hiding
- Name space control

Compiler's perspective: Separate compilation

- Smaller compilation units
- Recompilation only of changed units
- Library modules released as binaries

Increasing levels of sophistication

- Inclusion mechanism: concatenate all files before compilation
- Include with header files: headers with type information included for compilation and separate linking
- Import mechanism:
 - Compilation requires interface info from imported files
 - Compilation generates interface and object files
 - Often in OO languages, module = class

Extension proposal

- One module per file
- All modules in same directory (further extension: define search path mechanism)

Observations

- Mainly system for name space control and libraries
- If you want to implement it, you may get credits
- Difficulty: not much support in LLVM

New syntax

If M is a module name, then

- `import M` is a new form of declaration
- `M.f(e1, . . . , en)` is a new form of expression

Unqualified use

A function in an imported module may be used without the module qualification if the name is unique. Name clashes are resolved as follows:

- If two imported modules define a function f , we must use the qualified form
- If the current module and an imported module both define f , the unqualified name refers to the local function

Import

To use functions defined in M , another module must explicitly import M .

Hence, import is not transitive, i.e, if M imports L and L imports K , it does not follow that M imports K .

Dependency

- If M imports L , then M depends on L
- If M imports L and L depends on K , then M depends on K ; dependency is transitive

We assume that dependency is non-cyclic: if M depends on N , then N may not depend on M .

Compiler's tasks

When called by `jl c M.jl`, the compiler must

1. Read the import statements of M to get list of imported modules
2. Recursively, read the import statements of these modules (and report an error if some module not found)
3. Build dependency graph of involved modules
4. Sort modules topologically (and report error if cyclic import found)
5. Go through modules in topological order (M last) and check timestamps to see if recompilation is necessary

Hint: It is OK to require that import statements are in the beginning of the file and with one import per line to avoid need of complete parsing.

Symbol table

You need a symbol table with types of functions from all imported modules. This info is readily available in LLVM files, but needs to be collected (and parsed).

Build the symbol table so that unqualified names will find the correct type signature (i.e., you must check for name clashes).

Note 1 It is a good idea to replace unqualified names by qualified (for code generation)

Note 2 Type declaration for all imported functions must be added to LLVM file

This lecture:

- Language extensions
- Overview of some compiler technology

Next time:

- Backend (e.g. x86) extension