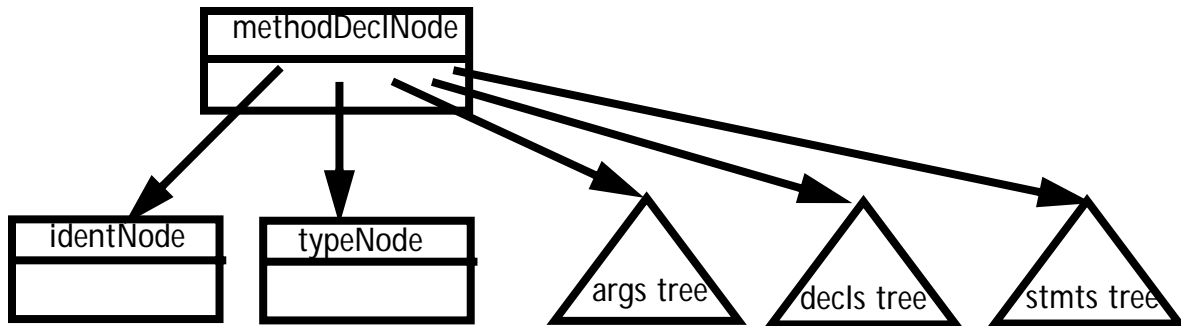


Type Checking Method Declarations

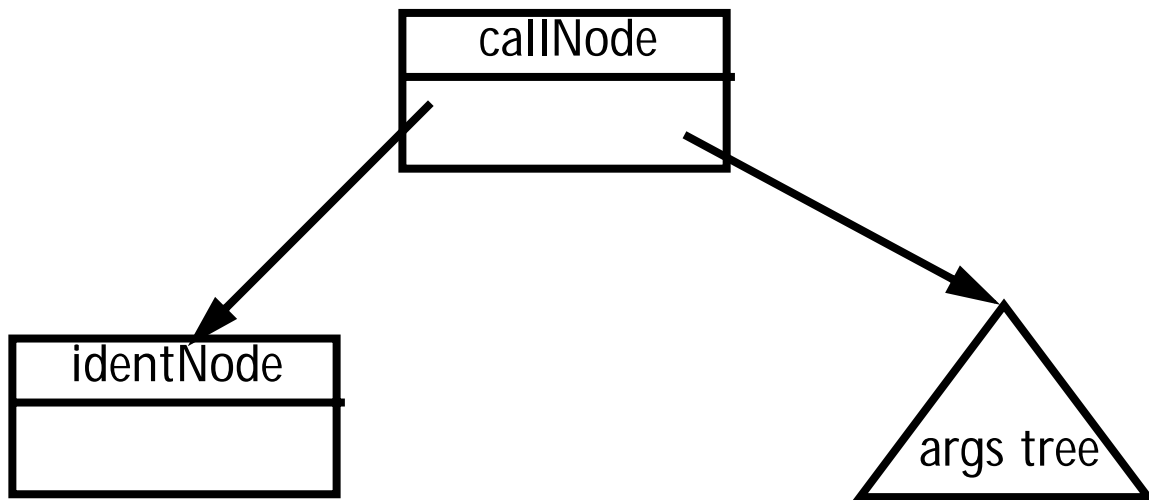


Type checking steps:

1. Create a new symbol table entry `m`, with `type = typeNode.type` and `kind = Method`.
2. Check that `identNode.idname` is not already in the symbol table; if it isn't, enter `m` using `identNode.idname`.
3. Create a new scope in the symbol table.
4. Set `currentMethod = this methodDeclNode`.

5. Type check the args subtree.
6. Build a list of the symbol table nodes corresponding to the args subtree; store it in m.
7. Type check the decls subtree.
8. Type check the stmts subtree.
9. Close the current scope at the top of the symbol table.

Type Checking Method Calls



We consider calls of procedures in a statement. Calls of functions in an expression are very similar.

Type checking steps:

1. Check that `identNode.idname` is declared in the symbol table. Its type should be `Void` and kind should be `Method`.
2. Type check the args subtree.

3. Build a list of the expression nodes found in the args subtree.
4. Get the list of parameter symbols declared for the method (stored in the method's symbol table entry).
5. Check that the arguments list and the parameter symbols list both have the same length.
6. Compare each argument node with its corresponding parameter symbol:
 - (a) Both should have the same type.
 - (b) A Variable, Value, or ScalarParm kind in an argument node matches a ScalarParm parameter. An Array or ArrayParm kind in an argument node matches an ArrayParm parameter.

Virtual Memory & Run-Time Memory Organization

The compiler decides how data and instructions are placed in memory.

It uses an *address space* provided by the hardware and operating system.

This address space is usually *virtual*—the hardware and operating system map instruction-level addresses to “actual” memory addresses.

Virtual memory allows:

- Multiple processes to run in private, protected address spaces.
- Paging can be used to extend address ranges beyond actual memory limits.

Run-Time Data Structures

Static Structures

For static structures, a fixed address is used throughout execution.

This is the oldest and simplest memory organization.

In current compilers, it is used for:

- Program code (often read-only & sharable).
- Data literals (often read-only & sharable).
- Global variables.
- Static variables.

Stack Allocation

Modern programming languages allow recursion, which requires *dynamic allocation*.

Each recursive call allocates a *new copy* of a routine's local variables.

The number of local data allocations required during program execution is not known at compile-time.

To implement recursion, all the data space required for a method is treated as a distinct data area that is called a *frame* or *activation record*.

Local data, within a frame, is accessible only while a subprogram is active.

In mainstream languages like C, C++ and Java, subprograms must return in a stack-like manner—the most recently called subprogram will be the first to return.

A frame is pushed onto a *run-time stack* when a method is called (activated).

When it returns, the frame is popped from the stack, freeing the routine's local data.

As an example, consider the following C subprogram:

```
p(int a) {  
    double b;  
    double c[10];  
    b = c[a] * 2.51;  
}
```

Procedure **p** requires space for the parameter **a** as well as the local variables **b** and **c**.

It also needs space for control information, such as the return address.

The compiler records the space requirements of a method.

The *offset* of each data item relative to the beginning of the frame is stored in the symbol table.

The total amount of space needed, and thus the size of the frame, is also recorded.

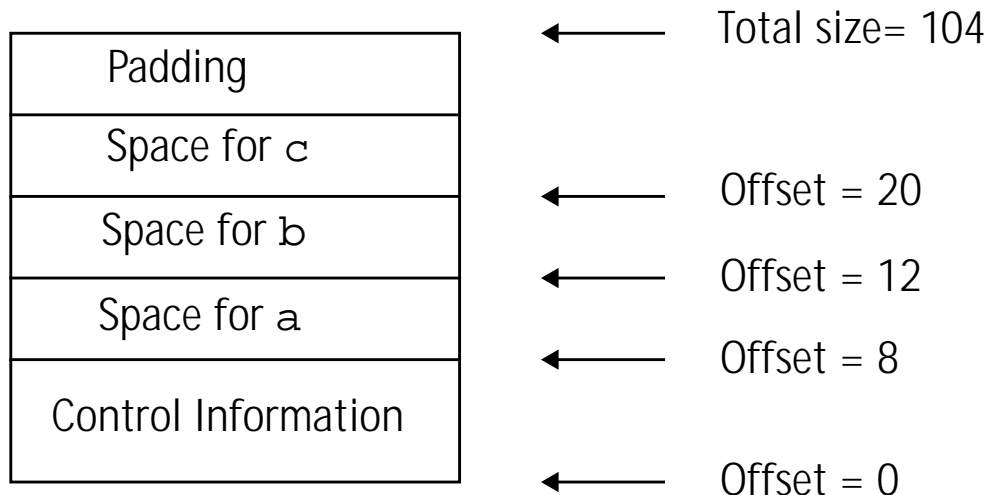
Assume **p**'s control information requires 8 bytes (this size is usually the same for all methods).

Assume parameter **a** requires 4 bytes, local variable **b** requires 8 bytes, and local array **c** requires 80 bytes.

Many machines require that word and doubleword data be *aligned*, so it is common to pad a frame so that its size is a multiple of 4 or 8 bytes.

This guarantees that at all times the top of the stack is properly aligned.

Here is p's frame:



Within **p**, each local data object is addressed by its offset relative to the start of the frame.

This offset is a fixed constant, determined at compile-time.

We normally store the start of the frame in a register, so each piece of data can be addressed as a **(Register, Offset)** pair, which is a standard addressing mode in almost all computer architectures.

For example, if register **R** points to the beginning of **p**'s frame, variable **b** can be addressed as **(R,12)**, with **12** actually being added to the contents of **R** at run-time, as memory addresses are evaluated.

Normally, the literal **2.51** of procedure **p** is not stored in **p**'s frame because the values of local data that are stored in a frame disappear with it at the end of a call.

It is easier and more efficient to allocate literals in a *static area*, often called a *literal pool* or *constant pool*. Java uses a constant pool to store literals, type, method and interface information as well as class and field names.

Accessing Frames at Run-Time

During execution there can be many frames on the stack. When a procedure **A** calls a procedure **B**, a frame for **B**'s local variables is pushed on the stack, covering **A**'s frame. **A**'s frame can't be popped off because **A** will resume execution after **B** returns.

For recursive routines there can be hundreds or even thousands of frames on the stack. All frames but the topmost represent suspended subroutines, waiting for a call to return.

The topmost frame is *active*; it is important to be able to access it directly.

The active frame is at the top of the stack, so the *stack top register* could be used to access it.

The run-time stack may also be used to hold data other than frames.

It is unwise to require that the currently active frame always be at *exactly* the top of the stack.

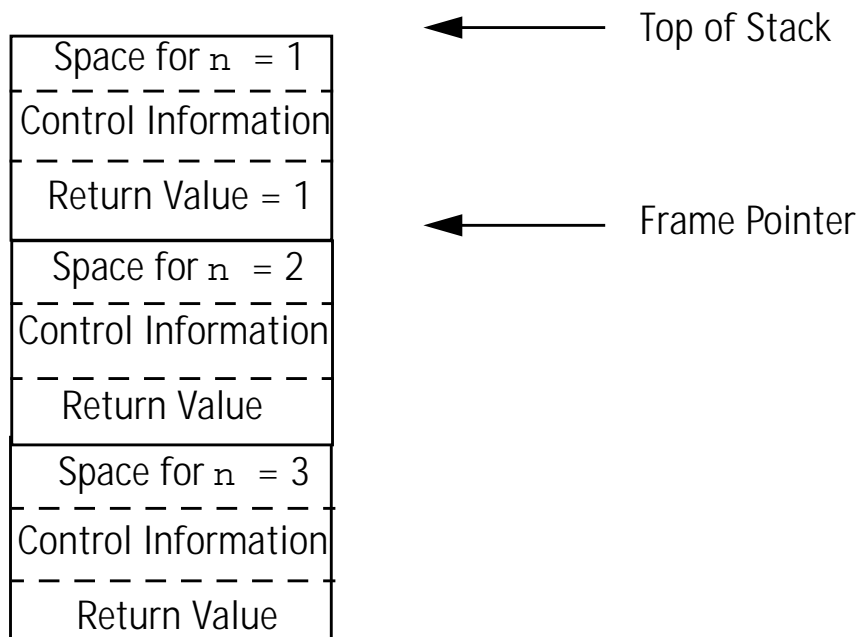
Instead a distinct register, often called the *frame pointer*, is used to access the current frame.

This allows local variables to be accessed directly as $\text{offset} + \text{frame pointer}$, using the indexed addressing mode found on all modern machines.

Consider the following recursive function that computes factorials.

```
int fact(int n) {  
    if (n > 1)  
        return n * fact(n-1);  
    else return 1;  
}
```

The run-time stack corresponding to the call **fact(3)** (when the call of **fact(1)** is about to return) is:



We place a slot for the function's return value at the very beginning of the frame.

Upon return, the return value is conveniently placed on the stack, just beyond the end of the caller's frame. Often compilers return scalar values in specially designated registers, eliminating unnecessary loads and stores. For values too large to fit in a register (arrays or objects), the stack is used.

When a method returns, its frame is popped from the stack and the frame pointer is reset to point to the caller's frame.

In simple cases this is done by adjusting the frame pointer by the size of the current frame.

Dynamic Links

Because the stack may contain more than just frames (e.g., function return values or registers saved across calls), it is common to save the caller's frame pointer as part of the callee's control information.

Each frame points to its caller's frame on the stack. This pointer is called a *dynamic link* because it links a frame to its dynamic (run-time) predecessor.

The run-time stack corresponding to a call of **fact(3)**, with dynamic links included, is:

