

Assignment 4

Structs and Arrays

15-411: Compiler Design

Rob Arnold (rdarnold@andrew) and Eugene Marinelli (emarinell@andrew)

Due: Thursday, October 23, 2008 (1:30 pm)

Reminder: Assignments are individual assignments, not done in pairs. The work must be all your own.

You may hand in a handwritten solution or a printout of a typeset solution at the beginning of lecture on Thursday, October 23. Please read the late policy for written assignments on the course web page. If you decide not to typeset your answers, make sure the text and pictures are legible and clear.

Problem 1

[20 points]

Consider the beloved C struct

```
struct C {
    char cFoo;
    int iBar;
    struct D *next;
    short sBaz;
    double dNaN;
    union {
        unsigned int uiTiberium;
        long long llRockets;
    }
};
```

- Give the total size of the structure and the offset of each member in bytes. Be sure to follow the standard x86-64 alignment convention.
- Is this structure minimally sized? If not, reorder the fields so that it is.
- Unlike C, languages such as Java and C# don't impose a strict ordering requirement on structure members; the runtime is free to reorder the fields as it sees fit. Give a simple algorithm to compute an ordering of fields which gives the smallest structure size.
- How would your algorithm from part c need to change to handle unions? Be sure to handle both anonymous and named unions.

Problem 2

[20 points]

- a. Unlike $L4$, C allows fixed-size arrays in structs. Give an SML struct implementing the following signature, `STRUCT.field_byte_offset` should take a struct and a field identifier and compute the offset of a field in the struct. `sizeof` should take a field type and return its size. Assume that any given struct has been checked for duplicate field names. `Ints` are 4 bytes.

```
signature STRUCT =
sig
  datatype field_type = Int
                        | Array of field_type * int
                        | Struct of field_type list

  type struct = (string * field_type) list

  val sizeof : field_type -> int
  val field_byte_offset : struct -> string -> int
end
```

- b. Suppose we wanted to extend the definition of $L4$ to permit fixed size arrays so they can be directly embedded in structs. Describe in detail how to modify the definition as given in the Lab 4 handout to accomodate this extension. This should include, as you find necessary, extended or modified syntax, static semantics including typing rules, and rules for program execution. You may assume a C-like unsafe model of execution where the results of certain operations, such as accessing an array out of bounds, are undefined.
- c. Discuss what further changes may be required if you want to allow for safe execution.

Problem 3

[20 points]

- a. Although costs for memory are dropping every year, memory consumption is still a problem, especially in large applications. Choosing the right types for variables can play an important role in performance in terms of both memory and speed. C provides many integral data types (`char`, `short`, `int`, `long`, `long long`) and floating point types (`float`, `double`, `long double`) but doesn't give good specifics on their sizes. In contrast, both `C#` and Java provide the same types¹ including absolute sizes so it is easier to pick a type that works in multiple environments. GCC and MSVC have extensions to allow integer structure fields to be sized explicitly in bits (up to 64). Since architectures generally do not provide instructions for writing arbitrary numbers of bits to arbitrary locations, these compilers implement writes to and reads from bitfields using shifting and masking.

¹Except for unsigned numbers in Java

Consider the following C code:

```
extern struct {
    int x:11;
    int y:21;
} s;

int setfield(int a) {
    s.y = a;
    return s.x;
}
```

Implement `setfield` in x86-64 assembly using as few instructions as possible. How would you implement basic arithmetic operations in general?

- b. In C, the `volatile` keyword can be appended to the type of a declaration to indicate that the value being declared may change at any point in the program's execution, for instance, through modification by another thread or by the hardware. The compiler must carefully avoid applying optimizations that assume that a `volatile` variable has a fixed value.

Give a piece of code and an optimization that cannot be applied to the code because a variable involved in the computation is `volatile`.

- c. Considering the way in which writes and reads involving bitfields are implemented by the compiler, describe a race condition that may arise in a multi-threaded program in which multiple threads share access to a struct with `volatile` bitfields, even when each bitfield has a lock which is always used to synchronize read and write access to it. As a compiler-writer, how would you address this problem?