

Xi High-Level ABI Specification

Computer Science 4120
Cornell University

Version of October 19, 2011

Changes

- Oct. 10: Added support for functions returning multiple results.
- Oct. 16: Added more information on calling conventions across platforms.

Introduction

The programs compilers produce are rarely standalone—they have to interface with the operating system’s libraries for things like I/O, memory management, and GUIs. In order to do this, programs must adhere to certain conventions. These conventions are usually specified by platform vendors (e.g. Microsoft, Apple, Intel, or The Linux Foundation) as part of what’s called the platform’s *application binary interface* (ABI).

These specifications are usually extremely helpful to compiler writers, because they remove the pain of having to resolve ambiguities (a process that you may already have found difficult!)

In PA5, Your compiler will interface with the Xi standard runtime that we will provide, but some of the details of that connection already need to be reflected in your PA4 IR. This document specifies the ABI your Xi compilers should follow to properly interface with the library. It assumes a 64-bit system and is meant to be similar to how C function calls work on [x86-64](#) on Windows, Linux and Mac OS X in 64-bit code.

Mangling function names

To better support separate compilation, your implementation must emit symbol names for functions and procedures that include type information exactly as specified below. This allows the linker to catch type errors when modules disagree about the types of functions.

Function and procedure names

The encoding of a procedure or function name is the sequence of the following:

- The string `_I`
- The name of the function, encoded as described below
- The underscore character, `_`
- The encoding of the return type, or `p` if this is a procedure
- Encodings of types of each of the arguments, if there are any

To encode function names all you need to do is replace a single underscore character with two of them (and 2 will get replaced with 4, and so on). Thus, a single underscore character can be known to separate out the function name from the argument information.

Type names are encoded as following:

- `int` is encoded as `i`
- `bool` is encoded as `b`
- An array of type τ is encoded as `a` followed by the encoding of the element type.
- A tuple, from a function returning multiple results, is encoded as `t` followed by the number of arguments returned followed by the encodings of the types returned.

Examples

Declaration	Symbol name
<code>main(args: int[] [])</code>	<code>._Imain_paa<i>i</i></code>
<code>unparseInt(n: int): int[]</code>	<code>._IunparseInt_a<i>ii</i></code>
<code>parseInt(str: int[]): int, bool</code>	<code>._IparseInt_t2i<i>bai</i></code>
<code>eof(): bool</code>	<code>._Ieof_b</code>
<code>gcd(a:int, b:int):int</code>	<code>._Igcd_i<i>iii</i></code>
<code>multiple__underScores()</code>	<code>._Imultiple____underScores_p</code>

Special names

You will need to use the runtime library to allocate heap memory for arrays. To do this, you need to call (in your IR) the function `._I_alloc_i`, passing it a single integer giving the number of bytes to allocate. This function will return the memory address corresponding to the first byte of the allocated memory.

If you detect an array index out of bounds access, you can call the `._I_outOfBounds_p` function, which will print an error message and abort execution.

Note that these names will not conflict with a source-level Xi function or procedure declaration, as identifiers are not permitted to start with `_`.

Nothing else requires special handling; for example, your main function will be expected to be called `._Imain_paai` as it is by above rules.

Memory layout of arrays

All types have sizes that are multiples of 8 and native alignment of 8. Arrays are therefore laid out sequentially, with no need for padding.

To implement the `length` operation on arrays, their size is stored in the “-1”st index of the array, e.g., immediately before cell 0. An array value is a reference to the memory location of cell 0.

Note: the blocks returned by `._I_alloc_i` will always be at least 8-aligned.

Passing arguments and returning results

Presently, two notable conventions for passing and returning function arguments exist: the System V ABI, which is used by Mac OS, Linux, and other Unix-like operating systems, and the independently developed Microsoft calling conventions.

In this course, we require that you support the System V ABI. The course staff will run your submissions on Linux—hence, we recommend that you test your compiler on a Linux machine. If you don’t have access to one, alternative options are to use the CSUGLab or a virtual machine (VMWare, VirtualBox, etc.).

Fortunately, given the right compiler architecture, it is relatively straightforward to support multiple calling conventions. This may be desirable if the members of your group use several different operating systems. This section therefore presents an overview of the System V ABI and important platform-dependent differences that you will need to consider if you choose to support operating systems other than Linux. Since this is only a high-level summary, you may also want to take a look at the actual specifications, which are linked at the bottom of this document.

Registers

In the System V ABI, six registers are designated for passing parameters. In order, these are: `rdi`, `rsi`, `rdx`, `rcx`, `r8`, and `r9`. The register `rax` is used to hold the return value of a function. The registers `rbx`, `rbp`, and `r12-r15` are *callee-saved*, which means that a function that uses these registers must restore their value before returning. All registers not classified as callee-saved are *caller-saved*, which means that their contents are potentially destroyed when calling a function (and their contents must therefore be saved if they are still needed after the function call).

On Windows, four registers are used to pass parameters: `rcx`, `rdx`, `r8`, and `r9`. Again, the register `rax` holds the return value of a function. The registers `rdi`, `rsi`, `rbx`, `rbp` and `r12-r15` are callee-saved.

Stack

When the number of function arguments exceeds six (or four on Windows), the remaining arguments are pushed to the stack right-to-left so that the last argument is the farthest away from the top of the stack. All stack entries are aligned to 8-byte boundaries. Care must be taken to ensure that the stack is aligned to a 16 byte boundary before issuing a function call¹.

On Windows, any function must furthermore allocate a temporary stack region named *shadow space* before issuing a call to another function. Conceptually, this is the stack space that would have been taken up by the first four parameters, had they not been passed using registers. However, this shadow region is always 32 bytes large (even if the function uses less than four parameters).

Passing and returning Xi values

Generally, `int` and `bool` values are passed directly using registers. Arrays have reference semantics, hence they are passed as 64-bit pointers using the same architectural registers. There are two cases when values are not passed using registers: the first one is the aforementioned situation when a function takes too many arguments.

The second case occurs when calling functions that return multiple results. Here, the caller must allocate space for the tuple return value, for instance within its stack frame. It passes a pointer to that memory region as an extra argument, before all the ordinary Xi-language arguments. The callee should save the returned values to that location. For a function returning n values, the memory region must be at least $8n$ bytes long to be able to store all the values.

We recommend that your register allocator computes a static stack frame layout for each function, which accounts for all the memory needed to store caller-save registers, tuple return values, function arguments, as well as unused memory needed to ensure 16-byte alignment and (potentially) the shadow space on Windows targets. Having a static layout means that the stack pointer is not manually modified in the body of a function, which makes it easier to ensure that the stack pointer is 16-byte aligned.

One such layout is the following (with memory addresses decreasing downwards).

Return address
Optional: saved frame pointer
Caller-saved registers
Spilled TEMPs
Optional: 8 bytes for alignment
Scratch space for functions that return tuples
Scratch space for arguments delivered on the stack
Optional: 32-bytes of shadow space

Some of these regions are not always necessary. For instance, a *leaf* function (i.e. a function that never calls other functions) does not need to allocate alignment or shadow space. If your compiler does not use an explicit frame pointer, it is not necessary to store its value directly after the return address. In this case, note that the `rbp` register is caller-saved, hence it still needs to be preserved if a function overwrites its value.

¹Your application may crash when it performs function calls without an aligned stack.

The runtime library comes with few basic Xi example programs and sample assembly output for all three platforms, which may be useful as a reference. The sample assembly code uses the stack layout shown here.

See also

You may find it useful to look at the specifications used in real-life systems, See [AMD64 Application Binary Interface \(v 0.99\)](#). The following document may also be useful: [x86-64 Machine-Level Programming](#). Groups wishing to support Microsoft calling conventions (this is optional) may want to look at the [MSDN documentation](#).