

Sindre Amdal Stephansen

The Challenges of porting Inferno to RISC-V

Master's thesis in Computer Science

Supervisor: Michael Engel

August 2021

Sindre Amdal Stephansen

The Challenges of porting Inferno to RISC-V

Master's thesis in Computer Science

Supervisor: Michael Engel

August 2021

Norwegian University of Science and Technology

Faculty of Information Technology and Electrical Engineering

Department of Computer Science



Norwegian University of
Science and Technology

Contents

1	Introduction	4
2	Background	4
2.1	Plan9	4
2.2	Inferno	5
2.3	RISC-V	5
2.3.1	ISA extensions	5
2.3.2	Privilege levels	6
2.3.3	CSR - Control and Status Register	6
2.3.4	SBI - RISC-V Supervisor Binary Interface	6
2.3.5	Traps	7
3	Implementation	7
3.1	RISC-V compiler	7
3.2	Architecture-specific code structure	8
3.3	Choosing a platform	9
3.4	Platform-specific code structure	9
3.4.1	Header files	9
3.4.2	Functions	10
3.4.3	The configuration file	11
3.4.4	The mkfile	12
3.5	OpenSBI	12
3.6	Address space	12
3.7	Using CSRs	12
3.8	Handling traps	14
3.8.1	Listener interface	14
3.9	Clock and timers	14
3.9.1	Timers in RISC-V	14
3.9.2	Timers in Inferno	15
3.9.3	Implementing the interface	15
3.10	UART	15
3.11	VIRTIO	16
3.11.1	The VIRTIO communication protocol	16
3.11.2	VIRTIO library	17
3.11.3	GPU	18
3.11.4	Input	18
3.11.5	Block device	18
3.12	Graphical output	19
3.13	Initializing the system	20
3.14	Interactive shell	20
3.15	Filesystem	22
3.16	The Just-in-time compiler	24
3.16.1	The Dis instruction set	25
3.16.2	The structure of the virtual machine	25
3.16.3	The structure of the JIT compilers	25
3.16.4	Implementing the JIT for RISC-V	26
3.16.5	Testing the JIT	27

4	Roadblocks	27
4.1	Debugging	27
4.2	Lack of documentation	28
4.3	Floating-point problems	29
4.4	Random crashes	30
4.5	The timer bug	30
5	Conclusion	31
6	Future work	32
7	Acknowledgments	32
8	Bibliography	33
	Appendices	35
A	/os/virtriscv/virtriscv	36
B	/os/virtriscv/mkfile	39
C	/os/virtriscv/main.c	41
D	/os/virtriscv/sdvirtblk.c	44
E	/libinterp/comp-riscv.c	47

Abstract

The RISC-V processor architecture is rapidly rising in popularity, and there will probably be an explosion of smaller RISC-V computers in the coming years, as sensors, in appliances, and more. Because these kinds of computers do not always have the resources to run an operating system like Linux, the Inferno operating system is an alternative, which, with its networked and distributed nature, could be a perfect match for these kinds of systems.

In this thesis I begin to port Inferno to RISC-V, and identify the challenges of both porting and using the operating system.

The first major challenge was to get the system to a stage where it could boot and handle simple input and output. The second challenge was to make the system more usable by implementing drivers. The last challenge was to implement a Just-in-time compiler, to make the system more responsive.

While not fully usable yet, I have made significant progress in porting Inferno. The operating system boots and launches an interactive shell, in which the user can execute commands. It can output to both a serial port and a screen. I have implemented a Just-In-Time compiler, but there are some bugs which cause complicated programs to crash.

This forms the foundation from which a port of Inferno to real hardware can be built.

Sammendrag

Prosesorarkitekturen RISC-V blir stadig mer populær, og det vil sannsynligvis bli en eksplosjon av små RISC-V maskiner de neste årene, som sensorer, i hvitevarer, og mer. Siden disse typene datamaskiner ikke alltid har ressursene til å kjøre operativsystemer som f.eks. Linux er operativsystemet Inferno et alternativ. Infernos distribuerte og nettverksorienterte design kan passe utmerket for disse typene systemer.

I denne oppgaven begynner jeg på å tilpasse Inferno til å kjøre på RISC-V, og identifiserer utfordringene med å tilpasse og bruke operativsystemet.

Den første store utfordringen var å få systemet til et punkt der det kunne starte opp og håndtere enkel kommunikasjon, i form av tekst. Den andre utfordringen var å gjøre systemet brukbart ved å implementere enhetsdrivere. Den siste utfordringen var å implementere en Just-in-Time kompilator, for å gjøre systemet mer responsivt.

Selv om operativsystemet ikke er helt brukbart ennå har jeg gjort store framskritt. Operativsystemet starter opp og viser en kommandolinje der brukeren kan utføre kommandoer. Tekst kan printes både til en seriell port og til en skjerm. Jeg har implementert en Just-in-Time kompilator, men det er noen problemer som får kompliserte programmer til å krasje.

Dette prosjektet danner grunnlaget for å bruke Inferno på RISC-V maskiner.

1 Introduction

The RISC-V platform is gaining ground in research as well as industrial projects. However, system software support so far is mostly focusing on the well-known large open source operating systems (Linux, BSD) or very tiny embedded real-time kernels. The large systems have now grown too big for many applications on restricted hardware platforms, e.g. running on a small FPGA-based system, whereas the traditional real-time operating systems suffer from a lack of useful network integration, memory protection, or orthogonal concepts of files and file systems, which makes their use in networked settings (IoT, Cloud) more challenging.

Thus, the idea of this project is to cover the middle ground by porting the open source Inferno operating system from Bell Labs to RISC-V. This system is already highly portable, but a RISC-V port is missing. Inferno is especially interesting since it is well documented and the low complexity of the system (compared to e.g. Linux) makes it very suitable to work on in the context of a student project.

During this project I started the work to create a port of Inferno to RISC-V, running under QEMU. I managed to get it to compile, print and receive input through UART. It enables and handles traps, schedules processes, starts the virtual machine, and launches an interactive shell which the user can execute commands from. I started to implement a Just-in-Time compiler, but there are bugs which causes crashes with complex programs.

The source code of the project is available at <https://github.com/kalkins/inferno-os/tree/riscv>. The code for the Just-in-Time compiler is at a separate branch, at <https://github.com/kalkins/inferno-os/tree/riscv-jit>.

This report is structured as follows: Section 2 gives an overview over the technologies used in this project. Section 3 goes through the development step by step. Section 4 covers major problems I encountered, and how I dealt with them. Section 5 summarizes the current state of the port. Section 6 discusses what remains in order to have a working port, and how this port could be used in practice.

2 Background

2.1 Plan9

Plan9 from Bell Labs, commonly shortened to Plan9, is a distributed operating system designed to solve some of the problems with UNIX-based workstations [18]. The operating system is designed to be distributed over a network of smaller workstations, giving each the full power of the network. Instead of maintaining UNIX compatibility, Plan9 kept the ideas that worked and redesigned the rest. Plan9 has a new suite of compilers, new libraries, and polished suite of tools.

Plan9 is built around the UNIX concept that everything is a file, and extends it. Most system resources are represented as files in the filesystem, and the files from other computers on the network are seamlessly available in the same filesystem. Because of this each machine can be responsible for a class of services which are made available through files, and any computer on the network can use those services as if they were hosted locally.

Plan9 also incorporated a concept of per-process name spaces, which means that each process has their own view of the filesystem. This is used, for example, by the graphical interface: When a process wants to display something on screen it writes to the `/dev/bitblt` file. The window system process can replace this file in the name space of its subprocesses and therefore intercept all writes. When the subprocess writes to `/dev/bitblt`, believing that it writes to the whole screen, the window system receives the request, translates the coordinates to within the window

given to the subprocess, and writes to the `/dev/bitblt` file in its own name space. This provides a simple way to encapsulate processes and build nested structures.

2.2 Inferno

Inferno is a distributed operating system based on Plan9 which focuses on portability and versatility, intended to be used for phones, TVs, and personal computers [8]. Inferno applications are written in the Limbo language and are compiled to byte-code which runs on a virtual machine, called Dis. The virtual machine was designed to be close to modern processor architectures at the time, and to make Just-in-time compilation fast and easy. The designers of the virtual machine claim that the JIT compiled code is 30-50% slower than native C [26].

2.3 RISC-V

At the time of writing the dominating ISAs, x86 for personal computers and servers, and ARM for embedded systems and phones, are proprietary. For x86 this has resulted in there being only two CPU manufacturers for mid- to high-end systems. For ARM, manufacturers must pay a licensing fee to use the design, which increases manufacturing cost, and there is little room for adapting the design to the rest of the hardware [1]. These designs are also often held back by the requirement of backwards-compatibility.

RISC-V, on the other hand, is a modern, free, open-source instruction-set architecture, which means that anyone can design a processor that fits the specification without paying licensing fees, and the design can be adapted to the hardware. Small embedded devices may implement the ISA in a cheap, straight-forward way, while desktop devices can use advanced techniques to get as much performance as possible.

The full ISA specification can be found in Waterman and Asanović [22, 23].

2.3.1 ISA extensions

One of the most interesting things about RISC-V is that the instruction-set is modular. The specification defines a few base ISAs, and several extensions that may or may not be dependent on other extensions. This allows hardware manufacturers to implement only the functionality that is needed, keeping the hardware simple for embedded devices, while allowing power and functionality to higher-end devices.

The width of registers is defined by XLEN, which is set by the base ISA. Most computational instructions are defined by this value, and therefore automatically use the available width on the platform. Platforms may allow XLEN to be changed at runtime.

The most basic ISA is RV32I, which uses XLEN=32, 32 registers, and defines common instructions like addition, move, load, store and branches. RV64I is another base ISA which builds on RV32I by keeping the instructions but changing XLEN to 64 and defining new instructions for 32-bit values. RV128I similarly changes XLEN to 128 and adds instructions for 64-bit values. RV32E, a base ISA with 16 registers designed for embedded systems, is currently in a draft stage.

The fact that the value of XLEN changes how instructions behave means that a program compiled for RV32I can be loaded on a 64-bit or 128-bit system and use the whole register width automatically. This can cause problems if the developer designs the code around 32-bit registers, especially if the code is designed to overflow. However, if the developer designs the code to work on both 32-bit and 64-bit platforms the full available width can be utilized without having separate versions for each value of XLEN.

Instruction-set extensions can be added to any base ISA in any configuration, as long as their dependencies are also included. The most notable extension is perhaps the **F** extension

which adds 32 registers for single-precision floats, and instructions to handle them. The **D** and **Q** extensions build upon this the same way as RV64I and RV128I, adding support for double and quad precision floats. There is also the **M** extension for multiplication and division, **A** for atomic instructions, and **C** for compressed instructions, which provides shorter variants of common instructions.

2.3.2 Privilege levels

The RISC-V specification defines multiple privilege levels, commonly called modes, in which code can be executed. The current mode determines which privileged instructions are available and how traps are handled. A higher mode can fully control a lower mode, providing security and functionality to operating systems and hypervisors.

Code running in a mode can make it impossible for code in lower modes to know which mode they are running in. When the lower code tries to do an operation that requires a higher privilege level, the upper code can emulate the operation.

There are three modes currently defined:

- Machine mode is the only mandatory mode defined by the specification, and it is the highest possible mode with full access to the platform. However, if only machine mode is available none of the benefits of privilege modes are available.
- User mode is an optional mode, and is always the lowest mode. It is used to run insecure user code, with higher modes granting protection.
- Supervisor mode is an optional mode that can be added between machine mode and user mode. It can be used for running operating systems with the bootloader in machine mode, the OS in supervisor mode, and user applications in user mode.

2.3.3 CSR - Control and Status Register

The RISC-V Zicsr extension defines a separate address space that can contain 4096 Control and Status registers (CSRs), and the instructions to use them [22, Chapter 9]. At the time of writing, over 200 CSRs have been defined.

The CSRs are divided into machine mode, supervisor mode and user mode and can be used to read platform information or enable and handle traps, timers, memory protection and virtual address translation for the different modes. CSRs also include information about XLEN and available extensions, so software can adapt to the platform at runtime.

When referring to a CSR independent of mode I use the format `xstatus`, which refers to `mstatus`, `sstatus`, and `ustatus`.

For the full list and description of CSRs, see Waterman and Asanović [23, Chapter 2, 3.1, 4.1].

2.3.4 SBI - RISC-V Supervisor Binary Interface

The Supervisor Binary Interface (SBI) is a standardized interface between software running in supervisor mode, usually operating systems or unikernels, and software running in higher modes, usually bootloaders or hypervisors. The SBI interface abstracts platform specific functionality, so that programs can be ported to all RISC-V implementations.

The SBI specification currently defines several extensions which the bootloader can offer, like setting timers, sending messages between harts, controlling a hart's state, performance monitoring and resetting the system. Earlier versions of the specification defined functions for reading and writing to a console, but these are now deprecated.

SBI functions are called using a standardized calling convention, which is a hybrid of the RISC-V and Linux calling conventions. Like the standard RISC-V convention, the registers `a0` to `a7` are used for arguments, but like in Linux the `a7` register is used for the ID of the extension that is called. Register `a6` can be used for the ID of the function, if the extension has multiple functions. The call itself is made with an `ECALL` instruction, which causes an exception in the higher modes. The return value is placed into `a1`, with `a0` indicating whether an error occurred.

The specification is still a draft in version 0.3, but it has been implemented by some bootloaders (see section 3.5). It can be found in Dabbelt and Patra [5].

2.3.5 Traps

Traps cause the currently executed code to be stopped, and control is transferred to a trap handler, usually in a higher mode. In RISC-V interrupts are traps that are used as notifications from instructions or devices. Exceptions are errors and environment calls.

Traps are an essential part of operating systems, for system calls, process scheduling, and error handling. In RISC-V traps are layered by mode: First, traps from any mode are sent to machine mode. The program running in machine mode can, before the trap, choose to delegate some or all traps occurring in supervisor or user mode to the program running in supervisor mode. Likewise, user mode traps can be delegated to the program running in user mode from supervisor mode.

Each mode has separate CSRs for enabling and handling traps. There are three classes of interrupts: software interrupts, timer interrupts, and external interrupts. After enabling these, interrupts also have to be enabled globally for the current mode `y` in the `ystatus` CSR. Exceptions can not be disabled, only delegated to a lower mode. When a trap is triggered and sent or delegated to mode `y` the trap handler at the address stored in `ytvect` is called. The specific cause of the trap is stored in the `ycause` CSR.

3 Implementation

3.1 RISC-V compiler

Before I could write any code for the port, I had to find a compiler which was compatible with the Plan9/Inferno compiler architecture, and could compile to RISC-V. Luckily, Richard Miller had already developed and published a RISC-V compiler for Plan9 by the time I started this project. The compiler source can be found in Miller [16].

I began the project by integrating Richard Miller's compiler into Inferno. Plan9 and Inferno have quite similar compiler structure, so this was easy to do. I came across some bugs in the compiler which I fixed as best I could.

After the compiler was integrated I started writing the architecture-specific code necessary for kernel and virtual machine functions, as described in section 3.2.

However, halfway through the project Richard Miller announced that he was working on a new improved compiler, with 64-bit support and more ISA extensions. This new compiler solved all the issues I was having with the old compiler. Richard Miller even added the architecture-specific code for Inferno, which replaced some of my attempts. See section 3.2 for more details. The compiler source can be found in Miller [17]. The compiler was later merged into the Inferno codebase [15].

This new compiler supports both the RV32I and RV64I base extensions, and the **I**, **M**, **A**, **F**, **D**, and **C** ISA extensions. Of these extensions only **C**, for compressed instructions, can be disabled. For the rest, if the platform does not support them the instructions either have to be avoided, or they can cause traps and be emulated in software.

For this project I used the 32-bit compiler because Inferno is a 32-bit operating system.

3.2 Architecture-specific code structure

Some architecture-specific code is necessary for the kernel and Dis virtual machine to run on the hardware. Some of these files were provided by Richard Miller as part of his new compiler, and some I have implemented myself. Most implementations are similar to that of other architectures.

- **Inferno/riscv/include**

These are architecture-specific header files which are used across the kernel.

- **lib9.h**

This file includes other header files.

- **u.h**

This file defines type aliases and floating point configuration constants.

- **ureg.h**

This file defines the **Ureg** struct, which is used to store register values.

- **libinterp**

This folder contains code for the Dis virtual machine. See section 3.16.

- **comp-riscv.c**

This file contains the implementation of the JIT compiler for RISC-V.

- **das-riscv.c**

This file contains the implementation of a RISC-V disassembler, which is used to debug the JIT compiler.

- **libkern**

This folder contains code for kernel libraries. These files were provided by Richard Miller.

- **mkfile-riscv**

Specifies source files for the architecture.

- **frexp-riscv.c**

This file provides functions for double-precision floats.

- **getfcr-riscv.s**

This file provides functions for reading and writing to the floating-point control and status register.

- **memmove-riscv.s**, **memset-riscv.s**, and **strchr-riscv.c**

These files implement the POSIX functions **memmove**, **memset**, and **strchr** for the architecture.

- **vlop-riscv.c** and **vlrt-riscv.c**

These files defines functions for arithmetic operations on integers longer than the platform bit width.

- **utils/libmach**

These files were provided by Richard Miller. The files with an **i** in the name are for RV32I, and those with a **j** in the name are for RV64I.

- **uregi.h** and **uregj.h**

These files define the **Ureg** struct, which is used to store register state.

- `i.c` and `j.c`
These files define the RISC-V registers and address space.
- `idb.c` and `jdb.c`
These files define a RISC-V specific debugger interface.
- `iobj.c` and `jobj.c`
These files provide functions used by the `iar` utility to help it recognize RISC-V object files.

3.3 Choosing a platform

Before any platform-specific code could be written, a platform had to be chosen. There are a few physical RISC-V processors available, but physical hardware can be hard to debug.

Instead, I chose to use QEMU. QEMU is a machine emulator which supports many architectures, and can emulate existing physical platforms [3]. It has support for the RV32I and RV64I base ISAs with all current ISA extensions, and machine, supervisor, and user mode. QEMU also has integrated GDB support, to make debugging easier, and it can emulate many input, storage, networking, and graphical devices.

3.4 Platform-specific code structure

The platform-specific code lives in `os/<platform>`, which is `os/qemuriscv` in this case. This code handles initialization of hardware, and provides functions that the rest of the kernel can use, hiding implementation details behind a common interface.

Because of the standardized nature of RISC-V, most of this code can be used for any RISC-V platform. However, at the moment only the QEMU platform has been added, so the code lives in that folder.

Most of the information about platform porting of Inferno comes from a series of blogposts by LynxLine Labs [10].

3.4.1 Header files

The kernel code often includes specific header files, expecting them to be defined in the platform-specific folder and provided to them through the linker. This means that the code is able to adapt better to the platform, but also that a lot of functions have to be defined before a minimal version of the kernel can be compiled.

Here is a list of the header files that had to be added to compile the kernel, and a description of what they provide:

- `mem.h`
This file defines the memory map of the platform, usually with macros. For more details see section 3.6.
- `dat.h`
This file defines platform-specific data structures like locks, labels, and machine configuration.
- `fns.h`
This file defines most platform-specific functions that other parts of the kernel need.

3.4.2 Functions

The common part of the Inferno kernel declares and uses several functions which are not implemented, which have to be implemented by the platform-specific part of the kernel. Here is an overview over those functions, and what they do:

- `int setlabel(Label*)` and `void gotolabel(Label*)`
These functions handle labels, which contain a program counter and a stack pointer. `setlabel` returns a label with the current program counter and stack pointer values, while `gotolabel` writes the stack pointer to the stack pointer register and jumps to the address in the labels program counter.
- `ulong getcallerpc(void*)`
This function returns the address of the instruction that called the function.
- `int _tas(int*)`
This function does a test-and-set, which is a simple atomic operation: It writes a 1 to the given address, and returns the previous value at that address. This can be used to create locks: When a 0 is returned the lock has been acquired, and a 0 can be written to release the lock. The RISC-V A extension includes an atomic swap instruction which makes this implementation very easy.
- `int splhi(void)`, `int spllo(void)`, `void splx(int)`, `void splxpc(int)`, and `int islo(void)`
These functions enable, disable, or toggle interrupts [14]. `islo` returns non-zero if interrupts are enabled. These can be implemented by setting, clearing, or toggling the supervisor mode interrupts in `sstatus`.
- `void kprocchild(Proc*, void (*)(void*), void*)`
This function configures a kernel process with a stack.
- `int segflush(void*, ulong)`
This function flushes a region to memory and invalidates the region in the instruction cache.
- `void idlehands(void)`
This function is called when process runner has nothing to do. It does not have to do anything.
- `void setpanic(void)`
This function is called to prepare for a panic, if necessary.
- `void dumpstack(void)`
This function dumps debug information about the stack to the user. It is only meant to help with debugging, and can be empty.
- `Timer* addclock0link(void (*)(void), int)`
This function sets a given function to be called after a given delay.
- `void clockcheck(void)`
This function is called to reset the watchdog timer, if necessary.
- `void FPinit(void)`, `void FPsave(void*)`, and `void FPrestore(void*)`
These functions are called from the Dis virtual machine to enable or disable floating point operations.
- `void exit(int)`, `void reboot(void)`, and `void halt(void)`
These functions respectively shut down, reboot, and send the system into an infinite loop.

3.4.3 The configuration file

Each platform must have a configuration file, which describes which parts of the OS should be compiled, global configuration variables, and which files and folders should be included in the filesystem. By convention the file has the same name as the platform, without a file extension, so for this platform the configuration file is `os/virtriscv/virtriscv`. The file uses a specific format, and is parsed by a script before compilation. It is divided into sections, where the section name is at the baseline and the contents are indented, one entry per line.

The configuration file is parsed once to import the mkfile dependencies, once to generate a C file which defines the global variables and links device drivers, and once to generate an assembly file and a header file with the contents of the filesystem.

The following sections are common in the configuration files:

- `dev`
This section defines device drivers source files to include from the `/os/port/` directory with the `dev` prefix.
- `ip`
This section defines C files to include from `/os/ip/`, which contain the network stack.
- `lib`
This section defines libraries to include. These are whole directories at the root level with the `lib` prefix, which are compiled into libraries and linked into the binary.
- `misc`
This section defines C files to include from the platform directory.
- `mod`
This section defines Limbo module definitions to include from `module/`.
- `port`
This section defines C files to include from `os/port/`.
- `code`
This section consists of C code which declares configuration variables, like enabling the JIT compiler.
- `init`
This section has only one entry, which defines the initial program to run in the virtual machine. This program will usually set up the system, then start either the shell or the window manager. The entry is the basename of the Limbo program in the `os/init/` folder. The Limbo program will be added as a mkfile dependency, and compiled when changed.
- `root`
This section defines the filesystem that is included in the binary. Each line specifies a path. If the path ends in a slash it represents a folder which should be present in the filesystem, but does not exist in the local filesystem. If the path does not end with a slash, the file with that path relative to the project root folder is copied into the filesystem, with that path. For example, if the Inferno project root is `/usr/inferno`, and the line `/dis/cd.dis` is in the `root` section, the file `/usr/inferno/dis/cd.dis` will be copied from the local filesystem to `/dis/cd.dis` in the filesystem in the binary.
The exception is the file `/osinit.dis`, which is copied from the location specified in the `init` section.

This section is mostly used to include essential programs and utilities in the binary. The rest of the programs, and other files, should be on a filesystem that is mounted after boot.

The full configuration file for this project is included in appendix A.

3.4.4 The mkfile

The mkfile defines the build process for the platform. It specifies the target architecture, the name of the configuration file, the platform specific header and source files (which are usually not included in the configuration file), and how the resulting binary is compiled and linked. The full mkfile is included in appendix B.

3.5 OpenSBI

OpenSBI is a bootloader developed by the RISC-V foundation which supports SBI and is included by default by QEMU when using the `virt` machine type. It initializes the machine, switches to supervisor mode, and jumps to a specified address, `0x80400000`, where a binary can be placed to be executed. By passing that address to the linker with the `-T0x80400000` flag and exporting to ELF with the `-H5` flag, the resulting ELF can be passed to QEMU and will be loaded correctly and started by OpenSBI.

For this project calls to OpenSBI will only be used to request timers, because RISC-V timers can only be set from M mode, and to shut down the system. The legacy SBI supported console I/O, but this feature is deprecated, and I only used it for debugging other I/O methods.

3.6 Address space

In QEMU RAM starts at address `0x80000000`, with the size being defined by the `-m` command-line parameter. OpenSBI is loaded in at address `0x80000000-0x8001ffff`, and expect the kernel code to be loaded at address `0x80400000`.

There is little documentation about the address space in Inferno, and other implementations are not fully consistent, but it seems like the kernel uses the space below where the kernel code is loaded in, while the user-space uses the space above the kernel. I gave the kernel 8 KiB of stack space from the kernel start at `0x80400000` and downwards. The space from the end of the kernel binary until the end of memory is used for pages for processes. Because the size of RAM can be varied with QEMU, I assume that 128 MiB is available, and the OS will not use more than that. Though Inferno normally does not need that much RAM to run, memory overflow bugs are common during the initial porting process, so it is advantageous to start of with larger RAM sizes. In the future, it might be possible to determine RAM size at runtime and adapt to that.

Supervisor mode does support virtual memory, but I have not used that functionality yet. Because all user processes in Inferno run in a virtual machine, hardware virtual memory is not necessary. However, it would be a useful security measure.

3.7 Using CSRs

Control and Status registers, as mentioned in Section 2.3.3, are used to handle traps, and therefore are vital to an operating system.

The problem with CSRs are that the instructions to operate on them encode both the operation and CSR address. That means that it is impossible to make a generalized function that can do any operation on any CSR at runtime, because the operation and address must be known at compile-time. Instead, separate functions have to be defined for each operation for each CSR.

Some compilers, like GCC, allow these to be implemented in C using inline assembly, but the Plan9 C compiler does not support this, so the functions have to be written in assembly.

Because writing four basically identical functions for a large set of CSR is a boring task, an automatic solution was needed. I created the script `generate-csr.sh` which reads a file `csrregs.h` which contains definitions of CSRs on the form shown in listing 1. The script reads the CSR names and writes function declarations to `csr.h` and implementations to `csr.s`, as shown in listing 2 and 3. These functions return signed numbers because some CSRs have a flag at the MSB position, and the code can check if the CSR value is less than 0 to check the flag regardless of the data width. For example, `xcause` uses the MSB to indicate whether the trap was caused by an interrupt, so the code can simply check whether `xcause` is less than 0 to see if the trap was caused by an error or an interrupt.

When including all currently defined CSRs the resulting code is around 4000 lines (1000 lines of function declarations and 3000 lines of assembly), which compiles to around 3 kilobytes, or 2% of the whole binary. Of course not all CSRs are needed, so the size can be reduced by commenting out sections of `csrregs.h`.

```
#define CSR_ustatus      0x000
#define CSR_wie        0x004
#define CSR_utvec     0x005
```

Listing 1: A snippet from `csrregs.h`

```
long csr_read_ustatus(void);
long csr_write_ustatus(long);
long csr_set_ustatus(long);
long csr_clear_ustatus(long);
```

Listing 2: A snippet from `csr.h`

```
TEXT csr_read_ustatus(SB), $-4
CSRRS CSR(CSR_ustatus), R0, R8
RET

TEXT csr_write_ustatus(SB), $-4
CSRWR CSR(CSR_ustatus), R8, R8
RET

TEXT csr_set_ustatus(SB), $-4
CSRRS CSR(CSR_ustatus), R8, R8
RET

TEXT csr_clear_ustatus(SB), $-4
CSRRC CSR(CSR_ustatus), R8, R8
RET
```

Listing 3: A snippet from `csr.s`

3.8 Handling traps

OpenSBI delegates most traps to supervisor mode, and starts the kernel in supervisor mode. Enabling interrupts when in supervisor mode requires writing the trap handler address to `stvec`, setting the `sie` bit in `sstatus`, and setting the bits corresponding to the desired traps in the `sie` CSR. The trap handler must save all registers, and restore them before returning, to avoid corrupting the state of the interrupted code. Another trap can occur while a trap is being handled, so the registers and stack have to be treated carefully.

QEMU has a separate layer for hardware interrupts through a platform level interrupts controller (PLIC) [4]. This controller is mapped at `0x0c000000`, and handles UART and disk interrupts. Interrupts for those devices are marked as external in the `xcause` CSR, and the PLIC has to be queried to get the exact cause. For UART interrupts are triggered when the input buffer starts being filled or the output buffer is empty.

3.8.1 Listener interface

I implemented a trap listener interface based on the one in the pc port, which allows various parts of the operating system to add and remove trap listeners at any time. There can be multiple listeners for each trap. PLIC interrupts are separated into a separate bus, selected with the `tbf` argument (name kept for consistency with the pc port).

The interface consists of the following functions:

- `void intrenable(long irq, void (*f)(Ureg*, void*), void* a, int tbf, char *name)`

This function enables a trap listener with the given interrupt request number (`irq`) and bus (`tbf`). If it is the only listener for a maskable interrupt, the interrupt is unmasked.

- `int intrdisable(int irq, void (*f)(Ureg *, void *), void *a, int tbf, char *name)`

This function disables a trap listener with the given interrupt request number (`irq`) and bus (`tbf`). If it is the only listener for a maskable interrupt, the interrupt is masked.

3.9 Clock and timers

3.9.1 Timers in RISC-V

The RISC-V specification defines a standard way of reading wall-clock time and setting timers. The platform should implement a machine mode accessible memory-mapped register, `mtime`, which ticks up at a fixed rate, though the rate might be different for each platform. There should also be a memory-mapped `mtimecmp` register. A timer interrupt should happen when the value of `mtime` is greater than the value of `mtimecmp` [23, Chapter 3.1.10].

These registers are not accessible from supervisor or user mode. Instead, the current time can be read from the `time` and `timeh` CSRs. These can be implemented to point to `mtime`, or the request can be intercepted and handled in machine mode.

The RISC-V specification does not define a way for lower privilege levels to set timers, instead leaving it up to the machine mode software to define a method for this. The SBI specification defines a method for this with the `void sbi_set_timer(uint64_t stime_value)` in the Timer extension, which allows software in supervisor mode to request a timer interrupt at a given time [5].

3.9.2 Timers in Inferno

Inferno implements most of the functionality for multiplexed timers on a single native timer. The following functions are left to be implemented in the platform-specific code:

- `uvlong fastticks(uvlong *hz)`
This function returns the current value of the real-time clock, and writes the period of the clock to `hz`.
- `void timerset(uvlong next)`
This function sets a timer interrupt to trigger when the real-time clock reaches the value of `next`.
- `void clockcheck(void)`
It is unclear what this function does. It is only called when busy-waiting for locks, and in some platform-specific drivers. All platforms implement it as an empty function. Some implementation comments mention that the function is used to reset watchdog timers.
- `void delay(int milliseconds)` and `void microdelay(int microsecond)`
These functions busy-waits for a given number of milli- or microseconds.

3.9.3 Implementing the interface

There does not seem to be a standardized way to find the clock period, but through testing I found that the period in QEMU is 10 000 000 Hz. I later verified this in the QEMU source code [19, include/hw/intc/sifive_clint.h, line 57].

In addition to the required functions I implemented the following functions:

- `void clockinit(void)`
This function enables timer interrupts and calls `timerset` to set a timer infinitely far in the future. It is called during setup before any timers are set.
- `void clockintr(Ureg *ureg, void*)`
This is the handler for timer interrupts. It sets a new timer infinitely far in the future, and calls `timerintr` function in Inferno.

During testing, I discovered that setting a timer to -1 through SBI immediately caused a timer interrupt, even though the SBI documentation specifies that this is a method to set a timer infinitely far in the future. Through testing, I discovered that setting timers higher than 2^{61} sometimes immediately triggers the timer interrupt. As a temporary workaround, I used the value 2^{60} as infinity, as that is over 3000 years in the future. See section 4.5 for more details.

3.10 UART

The serial port is an essential way for an operating system to communicate with the outside. In QEMU all output from the operating system through the serial port is printed to the screen, and anything the user types in the terminal QEMU is running in is sent through the serial port to the operating system.

QEMU emulates the 16550a UART to handle serial port communication. The UART has eight byte-wide registers which are mapped to the memory addresses `0x10000000-0x10000007`. The pc port of Inferno includes a driver for the 8250 UART line, which supports the 16550a. The

driver needed a little configuration for finding the UART port and setting up interrupts, but it mostly worked right out-of-the-box.

The UART port is set up by calling `i8250console` during system initialization, which sets up the FIFO, and configures the UART to use 9600 baud, 8 data bits, 1 stop bit, and no parity. For use with QEMU the configuration has little impact, as the whole system is emulated.

3.11 VIRTIO

While UART is useful for basic input and output, other devices are necessary for a fully usable system. QEMU can emulate a large variety of such devices, which gave me the choice of which drivers I wanted to implement. While the codebase for Inferno includes drivers for several devices I chose not to use them because they are old, possibly unstable, and might have compatibility problems with QEMU.

Instead, I chose to use VIRTIO [21] devices, because VIRTIO uses the same communication protocol for all types of devices, which eases driver development. VIRTIO also performs better than other drivers on QEMU, because it reduces the layers of abstraction between the host and guest systems. The disadvantage of VIRTIO is that it is not implemented on real hardware, and the drivers are therefore only usable for testing or running virtualized systems.

QEMU supports VIRTIO over PCI or memory-mapped IO (MMIO). While the pc port includes a PCI driver which could be adapted for the RISC-V port, I instead decided to use MMIO because of the simplicity of using such an interface. This requires that the VIRTIO addresses and irq numbers are configured at compile-time.

3.11.1 The VIRTIO communication protocol

For VIRTIO over MMIO all VIRTIO devices have a predefined address region. For QEMU these addresses are `0x10001000`, `0x10002000`, up to `0x10008000`, which gives a maximum of 8 VIRTIO devices. The memory region for each device starts with a set of *device registers*, which are used to negotiate device features, setup interrupts, and give the device pointer to the communication queues. After the registers there is a device-specific *configuration space*, which usually contains information about the device.

Data is sent between the driver and the device using *Virtqueues*. Each type of device has a different number of virtqueues for different purposes. This implementation uses *Split Virtqueues*, which separates the buffers the device should read from, and the buffers it should write to. A split virtqueue consists of three ring buffers: the *Descriptor table* containing pointers to buffers and metadata, the *Available Ring* with indexes of descriptors the device should handle, and the *Used Ring* with indexes of descriptors which the device has handled. Often the driver needs to send data and get a response, for which it allocates one descriptor and buffer pair for the device to read, and another pair for the device to write the response to, and sends them together in a descriptor chain. The driver is responsible for allocating the virtqueue and all buffers. The driver usually deallocates the associated buffer after a response.

The VIRTIO specification often describes messages as a single structure. However, such structures do not have to be sent by a single descriptor, but can be split up. The device will look at the size of each buffer associated with a descriptor, and reassemble the structure from there. This allows a structure to contain both read-only and write-only fields, as they can be split into descriptors that specify if the device can read or write. Structures can also contain arrays of undefined length, usually for large data transfers. These arrays do not need to be contiguous with the rest of the structure as long as they are referred to by a separate descriptor [21, Chapter 2.6.4].

3.11.2 VIRTIO library

Because VIRTIO uses a common communication protocol for all devices, I implemented a library which handles this communication, to make each driver simpler. The library provides flexible interrupt handling by letting response handlers be set per message, in addition to setting a default response handler for each VIRTIO queue. It uses the platform-specific header files and the interrupt listener functionality described in section 3.8.

The library has the following interface:

- `void virtio_init(void);`
This function is called during system initialization, and it checks that VIRTIO is available, and collects an internal list of the available devices.
- `virtio_dev *virtio_get_device(int type);`
This function returns the first unused device of the given type, which corresponds to the Device ID in the VIRTIO specification.
- `int virtio_setup(virtio_dev *dev, char *name, virtq_dev_specific_init virtq_init, le64 features);`
This function resets, configures, and initialized a VIRTIO device. `virtq_init` is called at right time in the negotiation process to allocate the queues needed for the device. `features` is the feature flags the driver supports. Only the features which both the device and driver supports are enabled.
- `void virtio_disable(virtio_dev *dev);`
This function resets a VIRTIO device.
- `void virtio_enable_interrupt(virtio_dev *dev, virtio_config_change_handler config_change_handler);`
This function enables interrupts for a VIRTIO device. `config_change_handler` is the listener for device configuration changes.
- `void virtio_disable_interrupt(virtio_dev *dev);`
This function disables interrupts for a VIRTIO device.
- `int virtq_alloc(virtio_dev *dev, uint queueIdx, ulong size);`
This function allocates a VIRTIO queue with a given index and size for a VIRTIO device.
- `int virtq_add_desc_chain(virtq *queue, virtq_intr_handler handler, void *handler_data, uint num, ...);`
This function adds a descriptor chain to the given VIRTIO queue. `handler` is the response handler for the chain. `handler_data` is a value that will be passed to the handler. `num` is the number of descriptors in the chain. For each descriptor there should be three sequential arguments, the address, the size, and a flag indicating whether the descriptor is writable by the device.
- `void virtq_free_chain(virtq *queue, virtq_desc *head);`
This function will free a previously allocated chain, as long as each descriptor was allocated separately.
- `void virtq_make_available(virtq *queue);`
This function will make all current descriptor chains in a queue visible to the device.

- `void virtq_notify(virtio_dev *dev, int queuenum, int notify_response, int avail_idx);`
This function will send a notification to the device of the descriptors in the available ring up to index `avail_idx`.
- `virtq_used_elem *virtq_get_next_used(virtq *queue);`
This function returns the next element in the used ring.

3.11.3 GPU

The VIRTIO GPU device uses one or multiple framebuffers to transmit display data from the driver to the device. The device has a copy of the framebuffer, called a resource, in its own memory. To set up a framebuffer the driver has to request the device to create a resource, then allocate the framebuffer and request that the framebuffer is connected to the resource, then request the device to use the resource for a given scanout (screen). When the screen should be updated the driver must send a message that a region of the framebuffer is invalidated, then request that the device flushes the region of the resource to the screen [21, Chapter 5.7].

At first, I implemented the driver to invalidate and flush the framebuffer for every write. However, this resulted in a lot of small updates, which were visibly slow. Instead, I implemented an update queue, and a timer which drains the queue and flushes each region. Updates which are close together are merged, to reduce the number of messages sent to the device.

3.11.4 Input

The VIRTIO input device represents all kinds of input devices, like keyboards, mice, joysticks etc. Unlike other kinds of devices the input device does not need to be polled, but writes to the next available descriptor whenever an event occurs. The driver allocates all descriptors during initialization, but does not deallocate them after use because they will simply be overwritten the next time the device gets to that index in the descriptor table.

Each input event consists of a type, a code, and a value, conforming to the evdev interface used by the Linux kernel [21, Chapter 5.8]. The evdev interface is described in Torvalds [20, Version 5.12.10, Documentation/input/event-codes.rst]. A full list of the key codes is available in Torvalds [20, Version 5.12.10, include/uapi/linux/input-event-codes.h].

Keyboard drivers in Inferno only interact with the rest of the operating system by adding the typed characters to the keyboard queue `kbdq`. This means that each driver has to keep track of modifier keys, and has to define the keymap. The pc port includes a keyboard driver which supports evdev events, so I used that driver with small modifications to work with VIRTIO. This driver uses the standard US keymap.

I started to implement a mouse driver, which uses the same device type as the keyboard but sends different events and key codes. However, because the window manager is not available (see section 3.14) there is limited use for it, and it is harder to test.

3.11.5 Block device

The VIRTIO block device represents a hard drive, which is usually backed by a file in the host file system. The device is fairly straight-forward to use, the metadata like block size and capacity is given in the device configuration space. Read and write requests are sent on the same format, containing a sector number to start from and an array of data to read from or write to, depending on the operation. The device responds by writing a status code to the end of the request structure.

In Inferno storage device drivers are represented by a `SDifc` structure which contains the name and function pointers to the standard storage device functions, or `nil` if the function is not defined for that device.

I implemented the following storage device functions for this driver:

- `SDev* pnp(void)`
This function discovers, sets up, and returns a linked list of all storage devices on this interface.
- `SDev* id(SDev*)`
This function gives each storage device in the given linked list a unique name. I used the naming scheme "virtblkX", where X is an incrementing number.
- `int enable(SDev*)`
This function enables interrupts from the given device. Returns 1 if successful, otherwise returns 0.
- `int disable(SDev*)`
This function disables interrupts from the given device. Returns 1 if successful, otherwise returns 0.
- `int verify(SDunit*)`
This function performs the equivalent of an SCSI inquiry command. Returns 1 if successful, otherwise returns 0.
- `int online(SDunit*)`
This function retrieves the storage device block size and storage capacity. Return 1 if successful, otherwise returns 0.
- `long bio(SDunit* unit, int lun, int write, void* data, long nb, long bno)`
This function performs a read or write request to or from the buffer `data`, starting at block `bno` until block `bno+nb`. Returns the number of bytes read or written. Because the function can not return the number of bytes until the operation is finished, the function is blocking. The function name probably means "buffered I/O", as it is linked to the Limbo library `Bufio` [11].

It is worth mentioning that there is an alternative to `bio` in the `int rio(SDreq*)` function. I decided not to implement this yet because `SDreq` seems to be based on SCSI, and `bio` seemed much easier to implement. From reading other implementations of `rio` I am not sure how it is supposed to work, but the name might mean "raw I/O".

The full driver implementation is included in appendix D.

3.12 Graphical output

In addition to the GPU driver there has to be an interface between Inferno and the driver which implements screen functions used in Inferno. For this I used the `screen.h` and `screen.c` files from Richard Miller's port of Plan9 to Raspberry Pi, modified for Inferno by *Lab 18, we have a screen!* [13]. This interface is designed for a framebuffer, so it was easily adapted to the VIRTIO GPU driver.

With these files in place a border is drawn around the screen when Inferno starts. All printed text, except that printed only to UART with `iprint`, is displayed on the screen. User input is printed as the user types in it. When the text reaches the bottom the window is scrolled down, to keep the most recent text in view.

3.13 Initializing the system

When QEMU starts it first gives control to OpenSBI, running in machine mode. OpenSBI sets up the machine, then calls the kernel in supervisor mode at address 0x80400000. The function at that address is called `_start()`, which is shown in listing 4.

The Plan9 assembler usually inserts a function prologue which allocates `x+4` bytes stack space automatically, based on the `$x` parameter, and stores the link register. An epilogue is inserted to load the link register and reset the stack. However, because the stack pointer is not initialized yet the first function has to be declared with `$(-4)`, which prevents the assembler from inserting a prologue and epilogue.

The `_start()` function sets up the registers for the rest of the kernel. It sets the stack pointer, register `R2`, to a predefined address from `mem.h`. It also uses a pseudoinstruction to set the static base, which is the address at the start of the kernel, to register `R3` for relative addressing. After the registers are initialized it calls `main()`, which continues the initialization from C code.

```
#include "mem.h"

TEXT _start(SB), $(-4)
    /* set static base */
    MOVW $setSB(SB), R3

    /* set stack pointer */
    MOVW $(MACHADDR+MACHSIZE-4), R2

    /* call main */
    JAL R1, main(SB)
```

Listing 4: The `_start()` function

The `main()` function first initializes the memory for the kernel [12]. First the bss section, used for static variables and located after the kernel binary, is cleared. Then the memory pool, located after the bss section until the end of memory, has to be defined for the kernel to know which portions it can use. Currently, the size of the memory is not checked at runtime, so the emulator has to be started with at least the same amount of memory as the kernel expects, which is currently 128 MiB.

After the memory is initialized, traps are enabled and timers are initialized. Then the print queue and device drivers, like UART, input, and GPU, are initialized. Then the screen is initialized, and the OS information is printed.

Finally, user processes and the VM are initialized, and the Dis binary `/osinit.dis` is executed. The `main()` function is shown in listing 5. The full `main.c` file is included in appendix C.

3.14 Interactive shell

Starting the interactive shell is the baseline for a usable Inferno installation. For the shell to be available, the Dis file for the shell itself and all programs which should be available from the shell must be included in the `root` section of the platform configuration file. The shell is started from the Dis init file, by loading the shell module and spawning a shell instance in a new thread. If shell commands should be executed during initialization, they can be executed directly using the shell module. The init code necessary to start the shell is included in listing 6. A screenshot of the system after starting the shell and running the `ls` command is shown in figure 1.


```

void
main() {
    // Clear bss
    memset(edata, 0, end-edata);
    memset(m, 0, sizeof(Mach));

    // Initialize the memory pool
    confinit();
    xinit();
    poolinit();
    poolsizeinit();

    // Enable traps and timers
    trapinit();
    clockinit();

    // Set up UART and the print queue
    printinit();
    i8250console();

    // Set up VIRTIO drivers
    virtio_init();
    input_init();

    // Initialize the screen
    screeninit();

    print("\nRISC-V QEMU\n");
    print("Inferno OS %s Vita Nuova\n\n", VERSION);

    // Start processes
    procinit();
    links();
    chandevreset();

    eve = strdup("inferno");

    userinit();
    schedinit();
}

```

Listing 5: The main() function

The next step up from the interactive shell is to start the window manager. However, the `wm` requires so many smaller programs to be included that it is unsuited to be compiled in the binary, and should be provided using a harddrive. However, as will be discussed in section 3.15, this is not possible yet.

```

implement Init;

include "sys.m";
    sys: Sys;
    print: import sys;
include "sh.m";
    sh: Sh;
include "draw.m";
    draw: Draw;
    Context: import draw;

Bootpreadlen: con 128;

Init: module
{
    init: fn();
};

init()
{
    sys = load Sys Sys->PATH;
    sh = load Sh Sh->PATH;

    sys->bind("#i", "/dev", sys->MREPL); # draw device
    sys->bind("#c", "/dev", sys->MAFTER); # console device
    sys->bind("#S", "/dev", sys->MAFTER); # storage devices

    spawn sh->init(nil, "sh" :: "-i" :: nil);
}

```

Listing 6: The Limbo code to start the interactive shell

3.15 Filesystem

As mentioned in 3.4.3, a simple filesystem is included in the binary to provide the programs needed to initialize the system. However, this filesystem is read-only, and while it is possible to cram in all the available Limbo programs, the binary quickly becomes unreasonably large. Instead, a separate filesystem should be mounted to provide the rest of the Limbo programs, and user-modifiable files. This could be done over the network to another computer using the 9P protocol, but for this project I used a harddrive utilizing the VIRTIO block device driver, as described in 3.11.5.

The harddrive QEMU presents to the driver is backed by a file in the host filesystem, where I allocated a partition usable for Inferno. First I tried using the `kfs` filesystem native to Plan9 and

```
Console
RISC-V QEMU
Inferno OS Fourth Edition (20151010) Vita Nuova

Starting init0()
Initial Dis: "/osinit.dis"
init: starting shell
; ls
boot
chan
dev
dis
env
fd
fonts
icons
lib
locale
man
module
n
net
osinit.dis
prog
services
tmp
usr
;
```

Figure 1: The system after starting the shell and running a command.

Inferno, but I had trouble finding Linux tools for it on the host side. It is possible to run Inferno hosted under Linux to format the partition, however that was a very cumbersome process. In addition, when mounting the filesystem in Inferno running on RISC-V, the kfs driver constantly had to check the filesystem, and froze when trying to mount it.

Instead, I used the FAT filesystem, which Linux fully supports. Inferno has a driver for FAT32, however it is uncertain how well all the features and extensions of the filesystem is supported.

After creating and partitioning the harddrive file, I mounted it on the host system and copied over the entire `/dis/` folder with all the compiled Limbo programs. I then added the command line parameters shown in listing 7 to QEMU to use the file as the harddrive, accessible as a VIRTIO block device. The drive is detected by the driver at boot, and is available in Inferno under `/dev/virtblk00/`. However, the partitions are not detected or represented by files automatically. To do that, the `fdisk` tool has to read the partition table and write the configuration to the disk control file. Because the partition type is FAT, the partition file will automatically be `/dev/virtblk00/dos`. Then the partition file must be mounted using the `dossrv` tool. Finally, the `dis` folder on the harddrive has to be bound to `/dis`, so all the program files are where they are expected. The full commands to achieve this is listed in listing 8. To reduce the number of manual commands during system setup, these commands are executed in the init file, before the shell is started.

Unfortunately the filesystem is read very slowly, because the block device driver is asked to read small sequential blocks. In addition, the system freezes halfway through reading files from the filesystem, like when using the kfs filesystem. However, it is unclear if this bug is in the block device driver, the filesystem driver, or some other program.

```
-drive if=none,format=raw,file=hdd.img,id=hdd -device
↪ virtio-blk-device,scsi=off,drive=hdd
```

Listing 7: The flags passed to QEMU to set up the hard drive with the VIRTIO block device driver.

```
disk/fdisk -p /dev/virtblk00/data > /dev/virtblk00/ctl
dossrv -f /dev/virtblk00/dos -m /n/local
bind /n/local/dis /dis
```

Listing 8: The Inferno shell commands to set up and mount the filesystem on a harddrive

3.16 The Just-in-time compiler

A Just-in-time (JIT) compiler dynamically translates one set of instructions to instructions native to the processor it is running on, at runtime [2]. This approach sacrifices some time and memory to compile the program, but the result will run faster than when using an interpreter. How fast the JIT compiles, and how fast the resulting code runs, depends on the similarity between the instruction sets, and which optimizations the JIT performs.

The Inferno OS includes a framework for JIT compilers which compile Dis programs to native instructions. As all user space programs are Dis programs in Inferno, a JIT compiler is essential to get a responsive system.

3.16.1 The Dis instruction set

The Dis virtual machine uses an instruction set modeled after CISC-processors, providing three-operand memory-to-memory instructions. The authors compare this approach to that of the Java stack-based virtual machine, and notes that the memory-to-memory approach is closer to common processors and makes the JIT compiler more efficient on non stack-based processors [26].

The instructions are organized into modules, which are loaded and compiled to native code separately. Each module has a data segment, and each function gets allocated a frame for local variables. The instructions can access values in the module data or function frame, or indirectly access values whose addresses are stored in one of those locations.

The instruction set has instructions for various datatypes, including 8-bit unsigned integers, 32 and 64-bit signed integers, 64-bit double precision floating-point, UTF-8 encoded strings, pointers, memory, and memory containing pointers.

The virtual machine uses reference-counted garbage collection [26]. As a result of this, pointers have to be handled using special instructions, to ensure that the garbage collector tracks every instance of the pointer. This includes instructions which allocate memory, so that task is moved from the programmer to the virtual machine [7].

The virtual machine has a few registers, to store the program counter, module data pointer, function frame pointer etc., but the registers are not directly accessible through the instruction set.

3.16.2 The structure of the virtual machine

The Dis virtual machine defines a C struct for the virtual registers, which is used by the interpreter and the compiled instructions. When moving between the interpreter and compiled code, or from the interpreter to an instruction handler, the normal C calling convention is disregarded, and the virtual registers are used instead. The handler for each instruction is separate from the rest of the interpreter, so the compiled code can call a handler in isolation if there is an instruction that is too complex or too infrequent to implement in the JIT compiler.

When reaching the entry point of each module the virtual machine will check whether the module has been compiled yet, execute it if it has, or try to compile it if it has not. If the compilation fails it uses the interpreter as a fallback. It seems to be possible to set the `MUSTCOMPILE` or `DONTCOMPILE` flags in the Dis binary to either force the module to be compiled, or be handled by the interpreter [6]. However, it does not seem like these flags are used by the Limbo compiler.

The Inferno JIT compilers are very simplistic compilers. They use a mixed code approach [2], but the decision to use native or interpreted code is done per instruction based on complexity, not based on how frequently a section is executed. After the first compilation, Inferno does not call the JIT compiler again for the same module, so no further optimizations are possible.

3.16.3 The structure of the JIT compilers

The only public function of a JIT compiler is the `compile` function, which is called when a new module should be compiled. However, the existing JIT compilers seem to follow the same basic structure.

The compilation is done in two passes. In the first pass each compiled Dis instruction is overwritten by the next one so that the total size and offsets are known for the second pass.

The JIT compiler will try to optimize the compiled instructions based on the information in the instruction, like in the size of the datatype, or by calculating based on the immediate value.

There are many Dis instructions which are complex or not supported by the native instruction-set. These are often delegated to the interpreter by loading the operators into the virtual registers

and calling the handler function for that instruction.

The JIT compilers often add macros, which are basically functions, to reduce code duplication of sections which are general and is not optimized at compile time. These macros are placed after each module.

3.16.4 Implementing the JIT for RISC-V

The best way to implement a new JIT compiler would be to start by delegating all instructions to the interpreter, then implementing one instruction at a time. However, I did not understand the way the Dis JIT compilers usually worked when I started this, so I did not see that possibility. Instead, I decided to look through the code of another JIT compiler line by line, and copy or translate each line as I understood what it did. This means that my JIT compiler implements roughly the same instructions as the one I based the code on. I mostly based my code on the ARM JIT compiler, because I am most familiar with ARM assembly. However, while ARM and RISC-V are RISC architectures, their instruction sets are quite different, so translating was sometimes hard. I sometimes used the MIPS JIT compiler as a second reference because the instruction set is much closer to RISC-V, but the structure and naming convention made the code hard to read.

The JIT compiler has to be careful how it allocates registers. RISC-V usually has 32 registers, however the Plan9 compiler only uses the first 16 to be more compatible with compressed instructions and the planned RV32E instruction set, which only has 16 registers. I decided to use the same restriction for the JIT compiler. Three registers have to be permanently reserved for the current frame pointer, module pointer, and pointer to the virtual registers, five registers are used for storing values for a single Dis instruction, one register is used for constructing 32-bit numbers from immediate values, and one register is used to store the address when loading double indirect operands. Finally, one register is used to store `H`, the Dis value for invalid pointers. Keeping `H` in a register simplifies comparing values to `H`, since it otherwise would have to be loaded into a register each time.

The JIT compiler starts by compiling a module preamble, which sets up the fixed registers, then jumps to the first compiled instruction. Then the first pass is compiled, each instruction overwriting the last, storing the compiled size of each Dis instruction. Then the buffer for the second pass is allocated based on the sizes, and the second pass then writes into the allocated buffer. Finally, initializers and destructors for each datatype is compiled.

The current implementation of the JIT compiler assumes that the `M` and `D` RISC-V extensions are supported by the processor. Arithmetic operations on 64-bit integers are emulated using 32-bit integers instead. Arithmetic operations for 32- and 64-bit integers and 64-bit floating-point numbers is implemented, and have undergone some simple test cases.

The conversion between 64-bit integers and 64-bit floats was initially delegated to the interpreter, however this uncovered an issue with the C compiler. When casting a float to an integer in C, the compiler will round the value by adding 0.5 or -0.5, depending on the sign, and then convert it in software, rounding down to the closest integer. However, the C compiler seems to expect that some registers hold float constants, like 0.5, but because I did not know about this these registers have not been set up, and the registers default to `NaN`. For the JIT compiler I stepped around these problems by handling more of the conversion in assembly. For conversions from floating-point to 64-bit integer I handle the rounding in assembly, then call the `_d2v` function, which uses bit manipulation to handle the rest of the conversion. For conversions from 64-bit integers to floating-point I translated the algorithm in the `_v2d` function to assembly, then optimized it to eliminate branches and reduce the number of instructions. The result is that these operations have been implemented using short and efficient assembly code, and the problem

with the C compiler has been circumvented.

When writing the code I added comments to explain the logic and exactly what was happening in the generated instructions, both to make it easier for me to come back to, and for future readers looking to understand the JIT compilers.

The full implementation of the RISC-V JIT compiler is included in appendix E.

3.16.5 Testing the JIT

The JIT compiler can be enabled by setting the `cflag` global variable in the configuration file higher than 0. The virtual machine will then try to compile all Dis modules before executing them. Of course, in a 2.5k line JIT compiler implemented in one go there was bound to be bugs. The first test run crashed with an illegal access exception, so I started working on ways to ease the debugging process, to more easily fix this and future bugs.

The most important debugging tool for the JIT compiler is the disassembler. While not required for the JIT compiler to work, it is common to implement a disassembler a separate file. For RISC-V I implemented the disassembler in the file `/libinterp/das-riscv.c`. The disassembler simply takes a pointer to the start of the compiled instructions, and the number of instructions, and prints out the address and assembly for each instruction. I used the standard RISC-V assembly syntax instead of the Plan9 assembly syntax, because it closely resembles how the instructions are laid down in the JIT compiler code. The JIT compiler calls the disassembler for each instruction after pass 2, prefaced with information about the associated Dis instruction. This makes it easy to follow the flow of the program, and check that the fields of the Dis instruction were used correctly.

Sometimes it can be useful to isolate the compiled code for a Dis instruction, to verify that it is correct despite other Dis instruction implementations which have bugs. In such cases it is useful to make all other Dis instructions use the interpreter handler during testing.

One easy way to check the logic of the compiled program is to insert an illegal instruction. This causes an exception, and the exception handler prints out the contents of the registers. For this I created the macro `CRASH()` which inserts a 32-bit 0, which is an illegal instruction in RISC-V. I have used this to check values loaded from memory, and to check which path of a branch was taken.

Sometimes the compiled code calls to C code, which crashes on illegal pointers. The stack trace will often track back to the calling compiled code, but not further because the compiled code does not use the stack. In these cases the called C code can be modified to print the contents of the virtual registers, which contains the PC of the associated Dis instruction, and arguments to the called function.

There have been many bugs in the JIT compiler, including illegal memory accesses, incorrect jump and branch offsets, mistranslations from ARM assembly, etc. Currently, the JIT compiler can correctly compile a simple print and if-statements. However, when trying to start the full shell the program crashes because of an index error. Further work is required to fix all the bugs and make the JIT compiler fully usable.

4 Roadblocks

4.1 Debugging

Debugging is an important part of any software development process, and even more so with something as complex as an operating system. However, operating systems are harder to debug because they lack the inherent framework that applications running inside operating systems

have. Especially in the early stages of development it is hard to get any debugging information out from the system.

Normal print-debugging can not be used until the system has established a communication channel with the outside, normally through UART, or SBI if the bootloader supports it. It is possible to print to the screen once that is set up, but the screen updates slower, and might not update at all during crashes. In addition, because the graphical pipeline is more complicated than the serial pipeline, trying to print to the screen can trigger bugs and even crash or freeze the system. Therefore, all debug printing is done through UART, while all other prints go both through UART and to the screen.

One alternative is GDB. Normally GDB is a very useful tool, and it is supported by QEMU, so it can debug from the very first instruction. However, most of the functionality of GDB is not available when running Inferno because of incompatibility with the Plan9 compilers. Plan9 and Inferno have their own symbol table format, which GDB does not support, and even if it could, the symbol table does not include enough information about types and line numbers to be usable. The result is that GDB can be used to debug, but it can not link the instructions in the binary to the source code. It can only be used to show a disassembly at the current location, show the value of registers, and set breakpoints at addresses. Passing the `-a` flag to the linker causes it to print the Plan9 assembly for the whole binary, including addresses, which can be used to figure out where to set breakpoints. However, because Plan9 assembly is different from the style used in GDB, and it includes many pseudoinstruction, figuring out which part is currently executing and finding bugs is very hard, and requires a lot of cross-referencing.

Inferno includes a debugger, called Acid [24, 9], which can debug the kernel [25], and which can be used from the host system through a serial port connection. To do this with QEMU, the serial port has to be exposed as a virtual `tty` using the `-serial pty` flag to QEMU. Then Acid has to be executed on the host system, with the command `acid -R <pty path> <OS binary path>`. However, the OS does not seem to respond to the messages sent by Acid, instead interpreting them as normal user input. This might be because the UART driver or some other step in the pipeline does not correctly identify the Acid control sequence.

The current debugging situation is not ideal, and it means that a lot of time is required to debug even the smallest bugs. The best solution would be to make the linker output contain enough debugging information for GDB to use, but that would require major alterations to the code and structure of the compiler and linker. With UART working, debugging by printing has become a simpler alternative for most situations, at least to find where the bug is.

When debugging the JIT compiler, the debugging methods I use for C does not work as well. The method I have used the most is to insert an illegal instruction as a breakpoint, which causes the trap handler to print information about the registers and the stack. While less flexible than other methods, it is very easy and fast to add and remove the illegal instructions, and to recompile and test. I sometimes use GDB, but that requires setting a breakpoint after each module is compiled, then looking through the disassembled instructions generated by the JIT and find the right address to break on. This makes it a lot more cumbersome to do rapid incremental testing. Inserting debug print statements into the compiled instruction stream is technically possible, but requires a lot of extra code and setup to make efficiently usable in assembly. However, because the compiled code often calls the interpreter, print statements can be inserted into the C code of the interpreter.

4.2 Lack of documentation

While Inferno (or rather Plan9) does have good documentation in general, the documentation relevant to porting, such as documentation of the kernel functions, compilers, and utilities, are at

best a minimal description of behavior, not a full guide.

For example, the documentation for the assembler explains addressing modes, function definitions, and calling conventions, but does not list the available instructions. Because Plan9/Inferno uses an assembly syntax that is meant to be similar for all platforms, but therefore is very different from more common assembly dialects, it is very hard to be sure what the assembly code actually does. I have had to read through the lexer and parser code on several occasions to figure out what an instruction does or how it should be used.

In addition, there is basically no official documentation for porting Inferno to new architectures or platforms. That means that I usually have to read through the code for other platforms, and develop incrementally, figuring out what I need to implement based on the errors I receive. This can be quite difficult because the compiler has vague error messages, and sometimes does not refer to the error location.

The result of this is that this project was very slow and time-consuming compared to how much code I had to write.

4.3 Floating-point problems

As mentioned in section 3.16.4, I came across some problems with the Plan9 C-compiler when performing floating-point arithmetic, because it assumes that certain registers hold common double-precision constants. While I discovered this while working on the JIT compiler, this problem affects all double-precision floating-point operations that uses the values the compiler assume are in registers.

Looking through the source code, it seems like this approach is fairly common, though it is not present in the x86 and ARM compilers. The registers the RISC-V compiler expects are listed in table 1. Floating-point constants that are not in this list are either constructed from other constants, or put into the data segment by the linker, and loaded when needed.

To solve this I added a function that is called during system initialization, which loads these values into the registers. With this change, double-precision floating point operations work as expected in C and when using the interpreter.

Because double-precision values can not be used for single-precision operations, and vice versa, separate constants are needed for single-precision operations. Because double-precision is the primary Dis floating-point type, and single-precision is only included for compatibility [7], single-precision values are not given permanent registers, and are always loaded from memory.

However, there seems to be a problem with the single-precision constants. All single-precision constants I have tested have had the hexadecimal representation `0x000001f4`. After digging around in the compiler code I discovered that when the linker writes the constant to the data segment, it is read from memory as if it was the highest four bytes of a double-precision value. However, because the value is stored in a four byte single-precision value, the read overflows. See listing 9. This bug is not present in the other compilers, and once found it was easy to fix.

Table 1: The floating-point constant register the compiler expects.

FP register	Value
28	0.0
29	0.5
30	1.0
31	2.0

```

fl = ieeedttof(p->to.ieee);
cast = (char*)&fl;
for(; i<c; i++) {
    buf.dbuf[l] = cast[fnuxi8[i+4]]; // Original version
    buf.dbuf[l] = cast[fnuxi4[i]];  // Fixed version
    l++;
}

```

Listing 9: The bugged code from `utils/il/asm.c` that writes single-precision floating-point constants to the data section.

4.4 Random crashes

The OS seems to have several bugs that trigger at random, which are hard to track down and fix. Because these issues usually appear right after booting, and less than half the time, I have not made them a priority to fix, and instead just try to start the OS again. However, they are serious problems that have to be fixed before Inferno on RISC-V can be used for any real-world purpose.

Here is a list of the issues that I know of:

- There is a problem with the memory management functions that crashes the system. The likelihood of it triggering seems to scale with the number of allocations that are requested. From my testing, the problem seems to be a buffer overflow which overwrites the heap metadata, causing a crash when the memory management functions walks the heap. The source of this overflow will be hard to determine.
- Once in a while, right after starting, QEMU reports an illegal memory access and crashes, before the OS got far enough to print. It might even be a problem with OpenSBI. This should be easy to debug with GDB, however it rarely happens, so one has to set up GDB and run QEMU again and again until it happens. Print-debugging would have to be done through SBI as the crash happens before the UART driver is set up.
- Very rarely the system seems to freeze at boot, without printing anything and without any error messages from QEMU. This might be because of a loop that does not end properly, but it's hard to determine because of the same difficulties with debugging as the issue above.

4.5 The timer bug

As mentioned in section 3.9.3, setting timer interrupts infinitely far in the future did not work as expected. While this does not affect programs which sets new timers continuously, like an operating system scheduling processes, it can cause problems for programs which only occasionally set timers.

I decided to test this by writing a small program that only set the timer, to reduce the number of possible errors. I discovered that this problem exists for both 32-bit and 64-bit RISC-V, and both when running in **S** mode and setting timers through OpenSBI, and when running in **M** mode and setting the timers directly. This means that the problem lies in QEMU. After looking through the QEMU source code I found the function `sifive_clint_write_timecmp` in the file `/hw/intc/sifive_clint.c`, which handles setting timers for all RISC-V platforms in QEMU. The function is listed in listing 10.

The first issue is on line 63, which converts the number of ticks until the next timer to nanoseconds. For the virtual RISC-V target, `NANOSECONDS_PER_SECOND/timebase_freq` is 100,

which is then multiplied with `diff`, causing an overflow. This is not really a problem if the timer is set to -1, because the result is still close to -1, but if only the most significant bit is set it will overflow to zero, causing an immediate interrupt.

The second problem is that `timer_mod`, which is the general function for setting timers for all QEMU platforms, takes a signed 64-bit integer as its second argument. A bit further down the chain, in `timer_mod_ns_locked`, this value is set to 0 if it was below 0.

These two problems combined mean that if the number of ticks until the next timer, multiplied with 100, has the most significant bit set, a timer interrupt is triggered immediately.

This bug has been reported to the QEMU developers.

```

38  /*
39  * Called when timecmp is written to update the QEMU timer or immediately
40  * trigger timer interrupt if mtimecmp <= current timer value.
41  */
42  static void sifive_clint_write_timecmp(RISCVCPU *cpu, uint64_t value,
43                                         uint32_t timebase_freq)
44  {
45      uint64_t next;
46      uint64_t diff;
47
48      uint64_t rtc_r = cpu_riscv_read_rtc(timebase_freq);
49
50      cpu->env.timecmp = value;
51      if (cpu->env.timecmp <= rtc_r) {
52          /* if we're setting an MTIMECMP value in the "past",
53             immediately raise the timer interrupt */
54          riscv_cpu_update_mip(cpu, MIP_MTIP, BOOL_TO_MASK(1));
55          return;
56      }
57
58      /* otherwise, set up the future timer interrupt */
59      riscv_cpu_update_mip(cpu, MIP_MTIP, BOOL_TO_MASK(0));
60      diff = cpu->env.timecmp - rtc_r;
61      /* back to ns (note args switched in muldiv64) */
62      next = qemu_clock_get_ns(QEMU_CLOCK_VIRTUAL) +
63             muldiv64(diff, NANOSECONDS_PER_SECOND, timebase_freq);
64      timer_mod(cpu->env.timer, next);
65  }

```

Listing 10: The implementation of the timer handling for RISC-V in QEMU.

5 Conclusion

Through this project I have investigated and tried to solve the challenges of porting the Inferno operating system to RISC-V, including setting up the boot process, programming trap handling, some important drivers, and a Just-In-Time compiler. I have not investigated the challenges of actual hardware platforms, but I believe that porting Inferno to such devices, with the necessary capabilities, should be easy after the groundwork I have done here.

At the end of this project I have created a port of Inferno to RISC-V which can print and receive input, output to a screen, run user processes in an interpreter, and mount a harddrive. In addition, I have started the work on a JIT compiler, which can compile arithmetic and function call instructions correctly.

6 Future work

There is still a lot of work to do for Inferno to be fully usable on RISC-V. The crashes mentioned in section 4.4 have to be fixed, the block-device driver has to be improved to increase the performance, and prevent the freeze mentioned in section 3.15. More drivers have to be implemented, especially for real hardware. In addition, the JIT compiler has to be fully tested and fixed to increase the performance of the system. Multi-core support can also be added to increase performance.

After these issues are fixed, and the system has been ported to real hardware, it can be tested on a network of embedded devices. This will show whether Inferno on RISC-V is practical and competitive for the IoT market.

7 Acknowledgments

I would like to thank my supervisor, Michael Engel, for his help and advice through this project.

I would also like to thank Richard Miller, who wrote the Plan9 RISC-V compiler, without which this project would not have gotten as far.

8 Bibliography

- [1] Krste Asanović and David A Patterson. “Instruction sets should be free: The case for RISC-V”. In: *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2014-146* (2014).
- [2] John Aycock. “A brief history of just-in-time”. In: *ACM Computing Surveys (CSUR)* 35.2 (2003), pp. 97–113.
- [3] Fabrice Bellard. “QEMU, a fast and portable dynamic translator.” In: *USENIX Annual Technical Conference, FREENIX Track*. Vol. 41. 2005, p. 46.
- [4] Palmer Dabbelt, Drew Barbier, and Abner Chang, eds. *RISC-V Platform-Level Interrupt Controller Specification*. RISC-V Foundation. Mar. 2020. URL: <https://github.com/riscv/riscv-plic-spec/blob/master/riscv-plic.adoc> (visited on 12/08/2020).
- [5] Palmer Dabbelt and Atish Patra, eds. *RISC-V Supervisor Binary Interface Specification*. RISC-V Foundation. Sept. 2020. URL: <https://github.com/riscv/riscv-sbi-doc/blob/master/riscv-sbi.adoc> (visited on 12/08/2020).
- [6] *Dis object file*. Vita Nuova Limited. URL: <http://www.vitanuova.com/inferno/man/6/dis.html> (visited on 07/18/2021).
- [7] *Dis Virtual Machine Specification*. Lucent Technologies Inc, Vita Nuova Limited. Sept. 2000. URL: <http://www.vitanuova.com/inferno/papers/asm.pdf> (visited on 06/05/2021).
- [8] Sean Dorward et al. *The Inferno Operating System*. Lucent Technologies, Bell Labs. 1997. URL: <http://www.vitanuova.com/inferno/papers/bltj.pdf> (visited on 12/08/2020).
- [9] Tad Hunt. *Acid Reference Manual*. Vita Nuova, Lucent Technologies. 2000. URL: <http://www.vitanuova.com/inferno/papers/acid.pdf> (visited on 08/01/2021).
- [10] *Inferno OS*. 2014. URL: <http://lynxline.com/category/inferno/> (visited on 12/08/2020).
- [11] Brian W. Kernighan. *A Descent into Limbo*. Vita Nuova, Bell Labs. 2005. URL: http://doc.cat-v.org/inferno/4th_edition/limbo_language/descent (visited on 08/01/2021).
- [12] *Lab 10, Bss, memory pools, malloc*. 2013. URL: <http://lynxline.com/lab-10-bss-menpools-malloc/> (visited on 12/08/2020).
- [13] *Lab 18, we have a screen!* 2013. URL: <http://lynxline.com/lab-18-we-have-a-screen/> (visited on 06/15/2021).
- [14] *Lab 9, coding assembler part*. 2013. URL: <http://lynxline.com/lab-9-coding-assembler-part/> (visited on 12/08/2020).
- [15] Richard Miller. Nov. 9, 2020. URL: <https://bitbucket.org/inferno-os/inferno-os/pull-requests/8> (visited on 07/30/2021).
- [16] Richard Miller. Mar. 3, 2021. URL: <https://9p.io/sources/contrib/miller/riscv-old.tar> (visited on 07/30/2021).
- [17] Richard Miller. May 28, 2021. URL: <https://9p.io/sources/contrib/miller/riscv.tar> (visited on 07/30/2021).
- [18] Rob Pike et al. *Plan 9 from Bell Labs*. Bell Labs. URL: <http://9p.io/sys/doc/9.html> (visited on 12/08/2020).
- [19] *QEMU*. URL: <https://gitlab.com/qemu-project/qemu> (visited on 07/30/2021).
- [20] Linus Torvalds, ed. *Linux Kernel*. URL: kernel.org (visited on 06/05/2021).

- [21] *Virtual I/O Device (VIRTIO) Version 1.1*. Oasis Open. Apr. 2019. URL: <https://docs.oasis-open.org/virtio/virtio/v1.1/virtio-v1.1.pdf> (visited on 06/05/2021).
- [22] *The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Document Version 2019121*. Tech. rep. RISC-V Foundation, Dec. 2019.
- [23] *The RISC-V Instruction Set Manual, Volume II: Privileged Architecture, Document Version 20190608-Priv-MSU-Ratified*. Tech. rep. RISC-V Foundation, June 2019.
- [24] Phil Winterbottom. *Acid: A Debugger Built From A Language*. 1995. URL: http://doc.cat-v.org/plan_9/2nd_edition/papers/acid/ (visited on 12/08/2020).
- [25] Phil Winterbottom. *Native Kernel Debugging with Acid*. URL: http://doc.cat-v.org/inferno/4th_edition/kernel_debugging/ (visited on 06/05/2021).
- [26] Phil Winterbottom and Rob Pike. *The design of the Inferno virtual machine*. Lucent Technologies, Bell Labs. URL: <http://www.vitanuova.com/inferno/papers/hotchips.pdf> (visited on 06/05/2021).

Appendices

A /os/virtriscv/virtriscv

```
1 dev
2     root
3     cons
4     env
5     mnt
6     pipe
7     prog
8     srv
9     dup
10    uart
11    sd
12
13    pointer
14    draw    screen
15    pointer
16
17    ip      bootp ip ipv6 ipaux iproute arp netlog ptclbsum iprouter plan9 nullmedium pktmedium netaux
18
19 ip
20     tcp
21     udp
22     ipifc
23     icmp
24     icmp6
25     ipmux
26
27 lib
28     interp
29     math
30     draw
31     memlayer
32     memdraw
33     tk
34     sec
35     kern
36
37 misc
38     uarti8250
39     sdvirtblk
40
41 mod
42     sys
43     draw
44     tk
45     math
46
47 port
48     alarm
49     alloc
50     allocb
51     chan
52     dev
53     dial
54     dis
55     discall
56     exception
57     exportfs
58     inferno
59     latin1
60     nocache
61     nodynld
62     parse
63     pgrp
64     print
65     proc
66     qio
67     qlock
68     random
69     sysfile
70     taslock
71     xalloc
72
73 code
74     int kernel_pool_pcmt = 10;
```



```

75     int main_pool_pcmt = 40;
76     int heap_pool_pcmt = 20;
77     int image_pool_pcmt = 40;
78     int cflag=0;
79     int swcursor=1;
80     int consoleprint=1;
81
82     init
83         virtriscvinit
84
85     root
86         /chan /
87         /dev /
88         /dis
89         /lib /
90         /env /
91         /fd /
92         /net /
93         /prog /
94         /n /
95         /n/local /
96         /n/dos /
97         /tmp /
98         /dis/lib
99         /dis/disk
100        /osinit.dis
101        /dis/sh.dis
102        /dis/tiny/sh.dis
103        /dis/ls.dis
104        /dis/mc.dis
105        /dis/lc
106        /dis/ps.dis
107        /dis/ns.dis
108        /dis/cat.dis
109        /dis/bind.dis
110        /dis/mount.dis
111        /dis/mmtgen.dis
112        /dis/listen.dis
113        /dis/export.dis
114        /dis/unmount.dis
115        /dis/sleep.dis
116        /dis/pwd.dis
117        /dis/echo.dis
118        /dis/cd.dis
119        /dis/netstat.dis
120        /dis/styxlisten.dis
121        /dis/time.dis
122        /dis/lib/arg.dis
123        /dis/lib/auth.dis
124        /dis/lib/lock.dis
125        /dis/lib/rand.dis
126        /dis/lib/random.dis
127        /dis/lib/dial.dis
128        /dis/lib/bufio.dis
129        /dis/lib/timers.dis
130        /dis/lib/string.dis
131        /dis/lib/filepat.dis
132        /dis/lib/readdir.dis
133        /dis/lib/workdir.dis
134        /dis/lib/daytime.dis
135        /dis/lib/nametree.dis
136        /dis/lib/styxservers.dis
137     # disk support
138        /usr /
139        /usr/inferno /
140        /dis/dd.dis
141        /dis/fs.dis
142        /dis/dosrv.dis
143        /dis/lib/fslib.dis
144        /dis/lib/fsproto.dis
145        /dis/lib/fsfilter.dis
146        /dis/zeros.dis
147        /dis/disk
148        /dis/disk/calc.tab.dis
149        /dis/disk/fdisk.dis
150        /dis/disk/format.dis

```

```
151         /dis/disk/ftl.dis
152         /dis/disk/kfs.dis
153         /dis/disk/kfscmd.dis
154         /dis/disk/mbr.dis
155         /dis/disk/mkext.dis
156         /dis/disk/mkfs.dis
157         /dis/disk/pedit.dis
158         /dis/disk/prep.dis
159         /dis/lib/disks.dis
160         /dis/lib/styx.dis
161     # misc
162         /dis/math/sieve.dis
163     # structure
164         /boot /
165         /man /
166         /fonts /
167         /icons /
168         /module /
169         /locale /
170         /services /
```

B /os/virtriscv/mkfile

```
1 #                                     --makefile--
2 <../../mkconfig
3
4 #Configurable parameters
5
6 CONF=virtriscv                       #default configuration
7 CONFLIST=virtriscv
8
9 SYSTARG=$OSTARG
10 OBJTYPE=riscv
11 INSTALLDIR=$ROOT/Inferno/$OBJTYPE/bin #path of directory where kernel is installed
12
13 LOADADDR=0x80400000
14
15 <$ROOT/mkfiles/mkfile-$$SYSTARG-$$OBJTYPE #set vars based on target system
16
17 <| $SHELLNAME ../port/mkdevlist $CONF
18
19 HFILES=\
20     mem.h\
21     dat.h\
22     fns.h\
23     io.h\
24
25 OBJ=\
26     load.$0\
27     clock.$0\
28     portclock.$0\
29     mul64fract.$0\
30     tod.$0\
31     plic.$0\
32     sbi.$0\
33     inb.$0\
34     dump.$0\
35     csr.$0\
36     trap.$0\
37     intr.$0\
38     virtio.$0\
39     input.$0\
40     kbd.$0\
41     mouse.$0\
42     gpu.$0\
43     archvirtriscv.$0\
44     main.$0\
45     $RISCVOBJ\
46     $IP\
47     $DEVS\
48     $ETHERS\
49     $LINKS\
50     $PORT\
51     $MISC\
52     $OTHERS\
53     $CONF.root.$0\
54
55 LIBNAMES=${LIBS:%=lib%.a}
56 LIBDIRS=$LIBS
57
58 CFLAGS=-wFV -I$ROOT/Inferno/$OBJTYPE/include -I$ROOT/include -I$ROOT/libinterp
59 KERNDATE=${$NDATE}
60
61 default:V: i$CONF
62
63 i$CONF: $OBJ $CONF.c $CONF.root.h $LIBNAMES
64     $CC $CFLAGS -DKERNDATE=$KERNDATE $CONF.c
65     $LD -l -o $target -H5 -T$LOADADDR $OBJ $CONF.$0 $LIBFILES
66
67 install:V: i$CONF
68     cp i$CONF $INSTALLDIR/i$CONF
69
70 <../port/portmkfile
71
72 trap.$0: csr.h
73
74 main.$0: $ROOT/Inferno/$OBJTYPE/include/ureg.h csr.h
```

```
75
76 csr.h csr.s: generate_csr.sh csrregs.h
77     sh generate_csr.sh csrregs.h
78
79 csr.$0: csr.h csr.s
80
81 devuart.$0: ../port/devuart.c ../port/uart.h
82     $CC $CFLAGS ../port/devuart.c
```

C /os/virtriscv/main.c

```
1  #include "u.h"
2  #include "../port/lib.h"
3  #include "dat.h"
4  #include "mem.h"
5  #include "fns.h"
6  #include "../port/uart.h"
7  #include "sbi.h"
8  #include "virtio.h"
9  #include "version.h"
10
11 #define MAXCONF      32
12
13 Conf conf;
14 Mach *m = (Mach*)MACHADDR;
15 Proc *up = 0;
16
17 char *confname[MAXCONF];
18 char *confval[MAXCONF];
19 int nconf;
20
21 extern int main_pool_pcnt;
22 extern int heap_pool_pcnt;
23 extern int image_pool_pcnt;
24
25 extern freginit(void);
26
27 /* Unimplemented functions */
28 void fpinit(void) {}
29 void FPsave(void*) {}
30 void FPrestore(void*) {}
31 int segflush(void*, ulong) { return 0; }
32 void idlehands(void) { return; }
33 void setpanic(void) { return; }
34
35 int
36 pcmspecial(char *idstr, ISAConf *isa)
37 {
38     return -1;
39 }
40
41 void
42 exit(int panic)
43 {
44     if (panic) {
45         iprint("PANIC\n");
46     }
47
48     SBI_SHUTDOWN();
49     for (;;)
50 }
51
52 void
53 reboot(void)
54 {
55     spllo();
56     print("Rebooting\n");
57     (*(volatile unsigned char*)(0x0000)) = 1;
58 }
59
60 void
61 halt(void)
62 {
63     spllo();
64     print("CPU halted\n");
65     while (1) {
66         wait_for_interrupt();
67     }
68 }
69
70 void
71 addconf(char *name, char *val)
72 {
73     if (nconf >= MAXCONF)
74         return;
```

```

75     confname[nconf] = name;
76     confval[nconf] = val;
77     nconf++;
78 }
79
80 char*
81 getconf(char *name)
82 {
83     int i;
84
85     for(i = 0; i < nconf; i++)
86         if(cistrncmp(confname[i], name) == 0)
87             return confval[i];
88     return 0;
89 }
90
91 void
92 confinit(void)
93 {
94     ulong base;
95     conf.topofmem = 128*MiB + RAMBOOT;
96
97     base = PGROUND((ulong)end);
98     conf.base0 = base;
99
100    conf.npage1 = 0;
101    conf.npage0 = (conf.topofmem - base)/BY2PG;
102    conf.npage = conf.npage0 + conf.npage1;
103    conf.ialloc = (((conf.npage*(main_pool_pcmt))/100)/2)*BY2PG;
104
105    conf.nproc = 100 + ((conf.npage*BY2PG)/MB)*5;
106    conf.nmach = MAXMACH;
107
108    print("Conf: top=0x%lux, npage0=0x%lux, ialloc=0x%lux, nproc=0x%lux\n",
109          conf.topofmem, conf.npage0,
110          conf.ialloc, conf.nproc);
111 }
112
113 void
114 poolsizeinit(void)
115 {
116     u64int nb;
117     nb = conf.npage*BY2PG;
118     poolsize(mainmem, (nb*main_pool_pcmt)/100, 0);
119     poolsize(heapmem, (nb*heap_pool_pcmt)/100, 0);
120     poolsize(imagemem, (nb*image_pool_pcmt)/100, 1);
121 }
122
123 void
124 init0(void)
125 {
126     Osend *o;
127     char buf[2*KNAMELEN];
128
129     up->nerrlab = 0;
130
131     print("Starting init0()\n");
132     spllo();
133
134     if(waserror())
135         panic("init0 %r");
136
137     o = up->env;
138     o->pgrp->slash = namec("#/", Atodir, 0, 0);
139     cnameclose(o->pgrp->slash->name);
140     o->pgrp->slash->name = newcname("/");
141     o->pgrp->dot = cclone(o->pgrp->slash);
142
143     chandevinit();
144
145     if(!waserror()){
146         ksetenv("cputype", "riscv", 0);
147         snprintf(buf, sizeof(buf), "riscv %s", conffile);
148         ksetenv("terminal", buf, 0);
149         poperror();
150     }

```

```

151         poperror();
152
153         disinit("/osinit.dis");
154     }
155 }
156
157 void
158 userinit(void)
159 {
160     Proc *p;
161     Osend *o;
162
163     p = newproc();
164     o = p->env;
165
166     o->fggrp = newfggrp(nil);
167     o->pggrp = newpggrp();
168     o->eggrp = neweggrp();
169     kstrdup(&o->user, eve);
170
171     strcpy(p->text, "interp");
172
173     p->fpstate = FPINIT;
174
175     p->sched.pc = (ulong)init0;
176     p->sched.sp = (ulong)p->kstack+KSTACK-8;
177
178     ready(p);
179 }
180
181 int
182 main() {
183     char input;
184
185     memset(edata, 0, end-edata);
186     memset(m, 0, sizeof(Mach));
187
188     freginit();
189     confinit();
190     xinit();
191     poolinit();
192     poolsizeinit();
193
194     trapinit();
195     clockinit();
196     printinit();
197     i8250console();
198     serwrite = uartputs;
199     virtio_init();
200     input_init();
201     screeninit();
202
203     print("\nRISC-V QEMU\n");
204     print("Inferno OS %s Vita Nuova\n\n", VERSION);
205
206     procinit();
207     links();
208     chandevreset();
209
210     eve = strdup("inferno");
211
212     userinit();
213     schedinit();
214
215     halt();
216     return 0;
217 }

```

D /os/virtriscv/sdvirtblk.c

```
1  #include "u.h"
2  #include "../port/lib.h"
3  #include "mem.h"
4  #include "dat.h"
5  #include "fns.h"
6  #include "io.h"
7  #include "virtio.h"
8
9  #include "../port/sd.h"
10
11 extern SDifc sdvirtblkifc;
12
13 SDev *head;
14
15 static int
16 blk_virtq_init(virtio_dev *dev)
17 {
18     if (dev->queues == 0) {
19         dev->queues = malloc(sizeof(virtq));
20         dev->numqueues = 1;
21
22         if (dev->queues == 0) {
23             panic("Virtio blk: Could not allocate queues. Malloc failed\n");
24         }
25
26         if (virtq_alloc(dev, 0, 0) != 0) {
27             panic("Virtio blk: Failed to create event queue");
28             return -1;
29         }
30
31         dev->queues[0].default_handler = nil;
32         dev->queues[0].default_handler_data = dev;
33     }
34
35     return 0;
36 }
37
38 static int
39 blk_enable(SDev* sdev)
40 {
41     virtio_enable_interrupt(sdev->ctrlr, nil);
42
43     return 1;
44 }
45
46 static int
47 blk_disable(SDev* sdev)
48 {
49     virtio_disable_interrupt(sdev->ctrlr);
50
51     return 1;
52 }
53
54 static SDev*
55 blk_pnp(void)
56 {
57     virtio_dev *dev;
58     SDev *sdev;
59     SDev **next;
60
61     for (next = &head; *next != 0; *next = (*next)->next) {}
62
63     while ((dev = virtio_get_device(VIRTIO_DEV_BLOCK)) != 0) {
64         int err = virtio_setup(dev, "BLK", blk_virtq_init, VIRTIO_F_ANY_LAYOUT
65             | VIRTIO_F_RING_INDIRECT_DESC | VIRTIO_F_RING_EVENT_IDX
66             | VIRTIO_BLK_F_RO | VIRTIO_BLK_F_SIZE_MAX | VIRTIO_BLK_F_SEG_MAX);
67
68         switch (err) {
69             case 0:
70                 sdev = malloc(sizeof(SDev));
71                 sdev->ctrlr = dev;
72                 sdev->ifc = &sdvirtblkifc;
73                 sdev->nunit = 1;
74
```



```

75         *next = sdev;
76         next = &sdev->next;
77
78         blk_enable(sdev);
79
80         break;
81     case -1:
82         iprint("Virtio blk rejected features\n");
83         break;
84     case -2:
85         iprint("Virtio blk queue error\n");
86         break;
87     default:
88         iprint("Virtio blk unknown error during setup %d\n", err);
89         break;
90     }
91 }
92
93     return head;
94 }
95
96 static SDev*
97 blk_id(SDev* sdev)
98 {
99     char name[16];
100    virtio_dev *dev;
101    static char idno[16] = "0123456789";
102
103    for (int i = 0; sdev != nil; sdev = sdev->next) {
104        if (sdev->ifc == &sdvirtblkifc) {
105            sdev->idno = idno[i++];
106
107            snprintf(name, sizeof(name), "virtblk%c", sdev->idno);
108            kstrdup(&sdev->name, name);
109        }
110    }
111
112    return nil;
113 }
114
115 static int
116 blk_verify(SDunit *unit)
117 {
118     virtio_dev *dev = unit->dev->ctrlr;
119
120     snprintf((void*) &unit->inquiry[8], sizeof(unit->inquiry)-8,
121             "VIRTIO port %d Block Device", dev->index);
122
123     unit->inquiry[4] = sizeof(unit->inquiry)-4;
124
125     return 1;
126 }
127
128 static int
129 blk_online(SDunit *unit)
130 {
131     virtio_dev *dev = (virtio_dev*) unit->dev->ctrlr;
132     virtio_blk_config *config = (virtio_blk_config*) &dev->regs->config;
133
134     if (dev->features & VIRTIO_BLK_F_BLK_SIZE) {
135         unit->secsize = config->blk_size;
136     } else {
137         unit->secsize = 512;
138     }
139
140     unit->sectors = config->capacity;
141
142     return 1;
143 }
144
145 static long
146 blk_bio(SDunit* unit, int lun, int write, void* data, long nb, long bno)
147 {
148     ulong len = nb * unit->secsize;
149     virtio_dev *dev = (virtio_dev*) unit->dev->ctrlr;
150     virtio_blk_config *config = (virtio_blk_config*) &dev->regs->config;

```

```

151     virtio_blk_req *req;
152     uchar *status;
153
154     if (write && dev->features & VIRTIO_BLK_F_RO) {
155         // The drive is read-only
156         iprint("VIRTIO block write of read only device\n");
157         return -1;
158     } else if (bno + nb > config->capacity) {
159         // Out of bounds
160         iprint("VIRTIO block device %s out of bounds\n", write ? "Write" : "Read");
161         return -1;
162     }
163
164     req = malloc(sizeof(*req));
165     req->type = write ? VIRTIO_BLK_T_OUT : VIRTIO_BLK_T_IN;
166     req->sector = (bno * unit->secsize) / 512; // VIRTIO always uses sectors of 512, though the device might
167     ↪ not
168     status = &req->status;
169     *status = 255;
170
171     virtq_add_desc_chain(&dev->queues[0], nil, nil, 3,
172                        req, VIRTIO_BLK_HDR_SIZE, 0,
173                        data, len, write ? 0 : 1,
174                        status, VIRTIO_BLK_STATUS_SIZE, 1);
175
176     virtq_make_available(&dev->queues[0]);
177     virtq_notify(dev, 0, 0, -1);
178
179     // Block until the drive responds
180     while (*status == 255) {}
181
182     switch (*status) {
183     case VIRTIO_BLK_S_OK:
184         free(req);
185         return len;
186         break;
187     case VIRTIO_BLK_S_IOERR:
188         iprint("VIRTIO block device IO error\n");
189         break;
190     case VIRTIO_BLK_S_UNSUPP:
191         iprint("VIRTIO block device unsupported operation\n");
192         break;
193     default:
194         iprint("VIRTIO block device returned %d\n", status);
195         error("Unknown VIRTIO block device return code\n");
196     }
197
198     free(req);
199     return -1;
200 }
201
202 SDifc sdvirtblkifc = {
203     "virtblk",           /* name */
204     blk_pnp,            /* pnp */
205     nil,                /* legacy */
206     blk_id,             /* id */
207     blk_enable,        /* enable */
208     blk_disable,       /* disable */
209
210     blk_verify,        /* verify */
211     blk_online,        /* online */
212     nil,              /* rio */
213     nil,              /* rctl */
214     nil,              /* wctl */
215
216     blk_bio,          /* bio */
217 };

```

E /libinterp/comp-riscv.c

```
1  #include "lib9.h"
2  #include "isa.h"
3  #include "interp.h"
4  #include "raise.h"
5
6  /*
7   * JIT compiler to RISC-V.
8   * Assumes that processor supports at least rv32mfd.
9   *
10  * Note that the operand order is different than the JIT compilers
11  * for other architectures, both for instructions and functions.
12  * The general order is rd, rs, imm. The exception is store instructions,
13  * which goes against the instruction operand ordering by having the source
14  * register first.
15  */
16
17  enum {
18      R0      = 0,
19      R1      = 1,
20      R2      = 2,
21      R3      = 3,
22      R4      = 4,
23      R5      = 5,
24      R6      = 6,
25      R7      = 7,
26      R8      = 8,
27      R9      = 9,
28      R10     = 10,
29      R11     = 11,
30      R12     = 12,
31      R13     = 13,
32      R14     = 14,
33      R15     = 15,
34
35      Rlink   = 1,
36      Rsp     = 2,
37      Rarg    = 8,
38
39      // Temporary registers
40      Rtmp    = 4, // Used for building constants and other single-instruction values
41      Rta     = 5, // Used for intermediate addresses for double indirect
42
43      // Permanent registers
44      Rh      = 6, // Contains H, which is used to check if values are invalid
45
46      // Registers for storing arguments and other mid-term values
47      RAO     = 8,
48      RA1     = 9,
49      RA2     = 10,
50      RA3     = 11,
51      RA4     = 12,
52
53      Rfp     = 13, // Frame pointer
54      Rmp     = 14, // Module pointer
55      Rreg    = 15, // Pointer to the REG struct
56
57      // Floating-point registers
58      F0      = 0,
59      F1      = 1,
60      F2      = 2,
61      F3      = 3,
62      F4      = 4,
63      F5      = 5,
64      F6      = 6,
65
66      // Opcodes
67      OP      = 51,           // 0b0110011
68      OPimm   = 19,          // 0b0010011
69      OPfp    = 83,          // 0b1010011
70      OPlui   = 55,          // 0b0110111
71      OPauipc = 23,          // 0b0010111
72      OPjal   = 111,         // 0b1101111
73      OPjalr  = 103,         // 0b1100111
74      OPbranch = 99,         // 0b1100011
```

```

75     OPload      = 3,           // 0b0000011
76     OPloadfp   = 7,           // 0b0000111
77     OPstore    = 35,          // 0b0100011
78     OPstorefp  = 39,          // 0b0100111
79     OPmiscmem  = 15,          // 0b0001111
80     OPsystem   = 115,         // 0b1110011
81     OPpamo     = 47,          // 0b0101111
82     OPmadd     = 67,          // 0b1000011
83     OPnmadd    = 79,          // 0b1001111
84     OPmsub     = 71,          // 0b1000111
85     OPnmsub    = 75,          // 0b1001011
86
87     // Rounding modes
88     RNE        = 0, // Round to nearest, ties to even
89     RTZ        = 1, // Round towards zero
90     RDN        = 2, // Round down
91     RUP        = 3, // Round up
92     RMM        = 4, // Round to nearest, ties to max magnitude
93     RDYN       = 7, // Use default
94
95     RM         = RDYN, // Default rounding mode
96
97     // Flags to mem
98     Ldw = 1,    // Load 32-bit word
99     Ldh,       // Load 16-bit half-word (with sign-extension)
100    Ldb,       // Load 8-bit byte (with sign-extension)
101    Ldhu,      // Load 16-bit unsiged half-word
102    Ldbu,      // Load 8-bit unsigned byte
103    Lds,       // Load 32-bit single-precision float
104    Ldd,       // Load 64-bit double-precision float
105
106    Stw,       // Store 32-bit word
107    Sth,       // Store 16-bit half-word
108    Stb,       // Store 8-bit byte
109    Sts,       // Store 32-bit single-precision float
110    Std,       // Store 64-bit double-precision float
111
112    Laddr,     // Special flag for operand functions
113              // Moves the address of the operand to a register
114
115    // Flags to branch
116    EQ = 1,
117    NE,
118    LT,
119    LE,
120    GT,
121    GE,
122
123    // Flags to punt
124    SRCOP = (1<<0),
125    DSTOP = (1<<1),
126    WRTPC = (1<<2),
127    TCHECK = (1<<3),
128    NEWPC = (1<<4),
129    DBRAN = (1<<5),
130    THREOP = (1<<6),
131
132    // The index of each macro
133    MacFRP = 0,
134    MacRET,
135    MacCASE,
136    MacCOLR,
137    MacMCAL,
138    MacFRAM,
139    MacMFRA,
140    MacRELQ,
141    NMACRO
142 };
143
144 // Masks for the high and low portions of immediate values
145 #define IMMSIGNED    0xFFFF800
146 #define IMMH        0xFFFF000
147 #define IMML        0x00000FFF
148
149 // Check if a immediate has to be split over multiple instructions
150 #define SPLITIMM(imm) (((((ulong)(imm) & IMMSIGNED) != 0)) && (((ulong)(imm)
↪ & IMMSIGNED) != IMMSIGNED))

```

```

151
152 #define SPLITH(imm)                (((ulong)(imm)) + (((ulong)(imm)) & (1<<11)) ? (1<<12) :
    ↪ 0) & IMMH
153 #define SPLITL(imm)              (((ulong)(imm)) & IMML)
154
155 // Extract bits of immediate values, like imm[11:5]. Basically shifts to the right and masks
156 // Examples:
157 // imm[11:5] -> IMM(imm, 11, 5)
158 // imm[11:0] -> IMM(imm, 11, 0)
159 // imm[11]   -> IMM(imm, 11, 11)
160 #define IMM(imm, to, from)        (((ulong)(imm)) >> (from)) & ((1 << ((to)-(from)+1)) -
    ↪ 1)
161
162 // All RISC-V instruction encoding variants. Set up with LSB on the left and MSB on the right, opposite to the
    ↪ tables in the RISC-V specification
163
164 #define Imm(imm)                   (IMM(imm, 11, 0)<<20)
165 #define Simm(imm)                  ((IMM(imm, 4, 0)<<7) | (IMM(imm, 11, 5)<<25))
166 #define Bimm(imm)                  ((IMM(imm, 11, 11)<<7) | (IMM(imm, 4, 1)<<8) | (IMM(imm,
    ↪ 10, 5)<<25) | (IMM(imm, 12, 12)<<30))
167 #define Uimm(imm)                  ((IMM(imm, 31, 12)<<12))
168 #define Jimm(imm)                  ((IMM(imm, 19, 12)<<12) | (IMM(imm, 11, 11)<<20) |
    ↪ (IMM(imm, 10, 1)<<21) | (IMM(imm, 20, 20)<<30))
169
170 #define Rtype(op, funct3, funct7, rd, rs1, rs2)    gen((op) | ((rd)<<7) | ((funct3)<<12) | ((rs1)<<15) |
    ↪ ((rs2)<<20) | ((funct7)<<25))
171 #define R4type(op, funct3, funct2, rd, rs1, rs2, rs3) gen((op) | ((rd)<<7) | ((funct3)<<12) | ((rs1)<<15) |
    ↪ ((rs2)<<20) | ((funct2)<<25) | ((rs3)<<27))
172 #define Itype(op, funct3, rd, rs1, imm)           gen((op) | ((rd)<<7) | ((funct3)<<12) | ((rs1)<<15) |
    ↪ Imm(imm))
173 #define Stype(op, funct3, rs1, rs2, imm)          gen((op) | ((funct3)<<12) | ((rs1)<<15) | ((rs2)<<20) |
    ↪ Simm(imm))
174 #define Btype(op, funct3, rs1, rs2, imm)          gen((op) | ((funct3)<<12) | ((rs1)<<15) | ((rs2)<<20) |
    ↪ Bimm(imm))
175 #define Utype(op, rd, imm)                        gen((op) | ((rd)<<7) | Uimm(imm))
176 #define Jtype(op, rd, imm)                        gen((op) | ((rd)<<7) | Jimm(imm))
177
178 /* Macros for laying down RISC-V instructions. Uses the instruction name from the specification */
179
180 // Upper immediate instructions
181 #define LUI(dest, imm)                        Utype(OPlui, dest, imm)
182 #define AUIPC(dest, imm)                     Utype(OPauipc, dest, imm)
183
184 // Jump instructions
185 #define JAL(dest, offset)                    Jtype(OPjal, dest, offset)
186 #define JALR(dest, base, offset)             Itype(OPjalr, 0, dest, base, offset)
187
188 // Branch instructions
189 #define BEQ(src1, src2, offset)               Btype(OPbranch, 0, src1, src2, offset)
190 #define BNE(src1, src2, offset)               Btype(OPbranch, 1, src1, src2, offset)
191 #define BLT(src1, src2, offset)               Btype(OPbranch, 4, src1, src2, offset)
192 #define BGE(src1, src2, offset)               Btype(OPbranch, 5, src1, src2, offset)
193 #define BLTU(src1, src2, offset)              Btype(OPbranch, 6, src1, src2, offset)
194 #define BGEU(src1, src2, offset)              Btype(OPbranch, 7, src1, src2, offset)
195
196 // Load instructions
197 #define LB(dest, base, imm)                   Itype(OPload, 0, dest, base, imm)
198 #define LH(dest, base, imm)                   Itype(OPload, 1, dest, base, imm)
199 #define LW(dest, base, imm)                   Itype(OPload, 2, dest, base, imm)
200 #define LBU(dest, base, imm)                  Itype(OPload, 4, dest, base, imm)
201 #define LHU(dest, base, imm)                  Itype(OPload, 5, dest, base, imm)
202
203 // Store instructions
204 #define SB(src, base, imm)                    Stype(OPstore, 0, base, src, imm)
205 #define SH(src, base, imm)                    Stype(OPstore, 1, base, src, imm)
206 #define SW(src, base, imm)                    Stype(OPstore, 2, base, src, imm)
207
208 // Arithmetic immediate instructions
209 #define ADDI(dest, src, imm)                  Itype(OPimm, 0, dest, src, imm)
210 #define SLTI(dest, src, imm)                  Itype(OPimm, 2, dest, src, imm)
211 #define SLTIU(dest, src, imm)                 Itype(OPimm, 3, dest, src, imm)
212 #define XORI(dest, src, imm)                  Itype(OPimm, 4, dest, src, imm)
213 #define ORI(dest, src, imm)                   Itype(OPimm, 6, dest, src, imm)
214 #define ANDI(dest, src, imm)                  Itype(OPimm, 7, dest, src, imm)
215
216 #define SLLI(dest, src, shamt)                Itype(OPimm, 1, dest, src, shamt)

```

```

217 #define SRLI(dest, src, shamt)           Itype(OPimm, 5, dest, src, shamt)
218 #define SRAI(dest, src, shamt)         Itype(OPimm, 5, dest, src, shamt | (1<<10))
219
220 // Arithmetic register instructions
221 #define ADD(dest, src1, src2)           Rtype(OP, 0, 0, dest, src1, src2)
222 #define SUB(dest, src1, src2)           Rtype(OP, 0, (1<<5), dest, src1, src2)
223 #define SLL(dest, src1, src2)           Rtype(OP, 1, 0, dest, src1, src2)
224 #define SLT(dest, src1, src2)           Rtype(OP, 2, 0, dest, src1, src2)
225 #define SLTU(dest, src1, src2)          Rtype(OP, 3, 0, dest, src1, src2)
226 #define XOR(dest, src1, src2)           Rtype(OP, 4, 0, dest, src1, src2)
227 #define SRL(dest, src1, src2)           Rtype(OP, 5, 0, dest, src1, src2)
228 #define SRA(dest, src1, src2)           Rtype(OP, 5, (1<<5), dest, src1, src2)
229 #define OR(dest, src1, src2)            Rtype(OP, 6, 0, dest, src1, src2)
230 #define AND(dest, src1, src2)           Rtype(OP, 7, 0, dest, src1, src2)
231
232 // The M extension for multiplication and division
233 #define MUL(dest, src1, src2)           Rtype(OP, 0, 1, dest, src1, src2)
234 #define MULH(dest, src1, src2)          Rtype(OP, 1, 1, dest, src1, src2)
235 #define MULHSU(dest, src1, src2)        Rtype(OP, 2, 1, dest, src1, src2)
236 #define MULHU(dest, src1, src2)         Rtype(OP, 3, 1, dest, src1, src2)
237 #define DIV(dest, src1, src2)           Rtype(OP, 4, 1, dest, src1, src2)
238 #define DIVU(dest, src1, src2)          Rtype(OP, 5, 1, dest, src1, src2)
239 #define REM(dest, src1, src2)           Rtype(OP, 6, 1, dest, src1, src2)
240 #define REMU(dest, src1, src2)          Rtype(OP, 7, 1, dest, src1, src2)
241
242 // The F extension for single-precision floating-point. rm is the rounding mode
243 #define FLW(dest, base, offset)         Itype(OPloadfp, 2, dest, base, offset)
244 #define FSW(src, base, offset)          Stype(OPstorefp, 2, base, src, offset)
245
246 #define FMADDS(rm, dest, src1, src2, src3) R4type(OPmadd, rm, 0, dest, src1, src2, src3)
247 #define FMSUBS(rm, dest, src1, src2, src3) R4type(OPmsub, rm, 0, dest, src1, src2, src3)
248 #define FNMADDS(rm, dest, src1, src2, src3) R4type(OPnmadd, rm, 0, dest, src1, src2, src3)
249 #define FNMSUBS(rm, dest, src1, src2, src3) R4type(OPnmsub, rm, 0, dest, src1, src2, src3)
250
251 #define FADDS(rm, dest, src1, src2)      Rtype(OPfp, rm, 0, dest, src1, src2)
252 #define FSUBS(rm, dest, src1, src2)      Rtype(OPfp, rm, 1<<2, dest, src1, src2)
253 #define FMULS(rm, dest, src1, src2)      Rtype(OPfp, rm, 1<<3, dest, src1, src2)
254 #define FDIVS(rm, dest, src1, src2)      Rtype(OPfp, rm, 3<<2, dest, src1, src2)
255 #define FSQRTS(rm, dest, src)            Rtype(OPfp, rm, 11<<2, dest, src, 0)
256
257 #define FSGNJS(dest, src1, src2)         Rtype(OPfp, 0, 1<<4, dest, src1, src2)
258 #define FSGNJNS(dest, src1, src2)        Rtype(OPfp, 1, 1<<4, dest, src1, src2)
259 #define FSGNJXS(dest, src1, src2)        Rtype(OPfp, 2, 1<<4, dest, src1, src2)
260
261 #define FMINS(dest, src1, src2)          Rtype(OPfp, 0, 5<<2, dest, src1, src2)
262 #define FMAXS(dest, src1, src2)          Rtype(OPfp, 1, 5<<2, dest, src1, src2)
263
264 #define FMVXW(dest, src)                 Rtype(OPfp, 0, 7<<4, dest, src, 0)
265 #define FMVWX(rm, dest, src)             Rtype(OPfp, 0, 15<<3, dest, src, 0)
266
267 #define FEQS(dest, src1, src2)           Rtype(OPfp, 2, 5<<4, dest, src1, src2)
268 #define FLTS(dest, src1, src2)           Rtype(OPfp, 1, 5<<4, dest, src1, src2)
269 #define FLES(dest, src1, src2)           Rtype(OPfp, 0, 5<<4, dest, src1, src2)
270
271 #define FCLASSS(dest, src)               Rtype(OPfp, 1, 7<<4, dest, src, 0)
272
273 #define FCVTWS(rm, dest, src)             Rtype(OPfp, rm, 3<<5, dest, src, 0)
274 #define FCVTWUS(rm, dest, src)           Rtype(OPfp, rm, 3<<5, dest, src, 1)
275 #define FCVTSW(rm, dest, src)            Rtype(OPfp, rm, 13<<3, dest, src, 0)
276 #define FCVTSWU(rm, dest, src)           Rtype(OPfp, rm, 13<<3, dest, src, 1)
277
278 // The D extension for double-precision floating-point
279 #define FLD(dest, base, offset)          Itype(OPloadfp, 3, dest, base, offset)
280 #define FSD(src, base, offset)           Stype(OPstorefp, 3, base, src, offset)
281
282 #define FMADDD(rm, dest, src1, src2, src3) R4type(OPmadd, rm, 1, dest, src1, src2, src3)
283 #define FMSUBD(rm, dest, src1, src2, src3) R4type(OPmsub, rm, 1, dest, src1, src2, src3)
284 #define FNMADDD(rm, dest, src1, src2, src3) R4type(OPnmadd, rm, 1, dest, src1, src2, src3)
285 #define FNMSUBD(rm, dest, src1, src2, src3) R4type(OPnmsub, rm, 1, dest, src1, src2, src3)
286
287 #define FADDD(rm, dest, src1, src2)       Rtype(OPfp, rm, 1, dest, src1, src2)
288 #define FSUBD(rm, dest, src1, src2)       Rtype(OPfp, rm, 5, dest, src1, src2)
289 #define FMULD(rm, dest, src1, src2)       Rtype(OPfp, rm, 9, dest, src1, src2)
290 #define FDIVD(rm, dest, src1, src2)       Rtype(OPfp, rm, 13, dest, src1, src2)
291 #define FSQRD(rm, dest, src)              Rtype(OPfp, rm, 45, dest, src, 0)
292

```

```

293 #define FSGNJD(dest, src1, src2)           Rtype(OPfp, 0, 17, dest, src1, src2)
294 #define FSGNJND(dest, src1, src2)         Rtype(OPfp, 1, 17, dest, src1, src2)
295 #define FSGNJXD(dest, src1, src2)         Rtype(OPfp, 2, 17, dest, src1, src2)
296
297 #define FMIND(dest, src1, src2)           Rtype(OPfp, 0, 21, dest, src1, src2)
298 #define FMAXD(dest, src1, src2)           Rtype(OPfp, 1, 21, dest, src1, src2)
299
300 #define FEQD(dest, src1, src2)            Rtype(OPfp, 2, 81, dest, src1, src2)
301 #define FLTD(dest, src1, src2)            Rtype(OPfp, 1, 81, dest, src1, src2)
302 #define FLED(dest, src1, src2)            Rtype(OPfp, 0, 81, dest, src1, src2)
303
304 #define FCLASSD(dest, src)                Rtype(OPfp, 1, 113, dest, src, 0)
305
306 #define FCVTSD(rm, dest, src)             Rtype(OPfp, rm, 32, dest, src, 1)
307 #define FCVTDS(rm, dest, src)            Rtype(OPfp, rm, 32, dest, src, 0)
308 #define FCVTWD(rm, dest, src)            Rtype(OPfp, rm, 97, dest, src, 0)
309 #define FCVTWUD(rm, dest, src)           Rtype(OPfp, rm, 97, dest, src, 1)
310 #define FCVTDW(rm, dest, src)            Rtype(OPfp, rm, 105, dest, src, 0)
311 #define FCVTDWU(rm, dest, src)           Rtype(OPfp, rm, 105, dest, src, 1)
312
313 // Pseudoinstructions
314 #define MOV(rd, rs)                        ADDI(rd, rs, 0)
315 #define NOT(rd, rs)                        XORI(rd, rs, -1)
316 #define NEG(rd, rs)                        SUB(rd, R0, rs)
317
318 #define BEQZ(rs, offset)                   BEQ(rs, R0, offset)
319 #define BNEZ(rs, offset)                   BNE(rs, R0, offset)
320 #define BLEZ(rs, offset)                   BGE(R0, rs, offset)
321 #define BGEZ(rs, offset)                   BGE(rs, R0, offset)
322 #define BLTZ(rs, offset)                   BLT(rs, R0, offset)
323 #define BGTZ(rs, offset)                   BLT(R0, rs, offset)
324
325 #define BGT(rs1, rs2, offset)              BLT(rs2, rs1, offset)
326 #define BLE(rs1, rs2, offset)              BGE(rs2, rs1, offset)
327 #define BGTU(rs1, rs2, offset)             BLTU(rs2, rs1, offset)
328 #define BLEU(rs1, rs2, offset)            BGEU(rs2, rs1, offset)
329
330 #define JUMP(offset)                       JAL(R0, offset)
331 #define JL(offset)                         JAL(R1, offset)
332 #define JR(rs, offset)                     JALR(R0, rs, offset)
333 #define JRL(rs, offset)                    JALR(R1, rs, offset)
334
335 /* Helper macros */
336
337 // Used to look up the address of an array element relative to base
338 #define IA(s, o)                            (ulong)(base+s[o])
339
340 // The offset from the current code address to the pointer
341 #define OFF(ptr)                            ((ulong)(ptr) - (ulong)(code))
342
343 // Call a function at the given address
344 #define CALL(o)                             (LUI(Rtmp, SPLITH(o)), JRL(Rtmp, SPLITL(o)))
345
346 // Return from a function
347 #define RETURN                              JR(Rlink, 0)
348
349 // Call a macro. Takes the macro idx as the argument
350 #define CALLMAC(idx)                       CALL(IA(macro, idx))
351
352 // Jump to a specific address
353 #define JABS(ptr)                          (LUI(Rtmp, SPLITH(ptr)), JR(Rtmp, SPLITL(ptr)))
354
355 // Jump to a Dis address
356 #define JDIS(pc)                           JABS(IA(patch, pc))
357
358 // Jump to an address in the dst field of an instruction
359 #define JDST(i)                            JDIS((i->d.ins - mod->prog))
360
361 // Set the offset of a branch instruction at address ptr to the current code address
362 // The order is opposite from OFF because it is used where the branch should jump to,
363 // not where it jumps from
364 #define PATCHBRANCH(ptr)                   *ptr |= Bimm((ulong)(code) - (ulong)(ptr))
365
366 // Gets the address of a PC relative to the base
367 #define RELPC(pc)                          (ulong)(base+(pc))
368

```

```

369 // Throw an error if the register is 0
370 #define NOTNIL(r)          (BNE(r, R0, 12), LUI(Rtmp, nullity), JRL(Rtmp, nullity))
371
372 // Array bounds check. Throws an error if the index is out of bounds
373 #define BCK(rindex, rsize) (BLTU(rindex, rsize, 8), /*CALL(bounds)*/ CRASH())
374
375 // Cause an immediate illegal instruction exception, which should
376 // cause a register dump, stack trace, and a halt.
377 // Useful for debugging register state
378 #define CRASH()           gen(0)
379
380
381 static  ulong*  code;
382 static  ulong*  codestart;
383 static  ulong*  codeend;
384 static  ulong*  base;
385 static  ulong*  patch;
386 static  ulong   codeoff;
387 static  int     pass;
388 static  int     puntpc = 1;
389 static  Module* mod;
390 static  uchar*  tinit;
391 static  ulong*  litpool;
392 static  int     nlit;
393 static  ulong   macro[NMACRO];
394 static  void    (*comvec)(void);
395 static  void    macfrp(void);
396 static  void    macret(void);
397 static  void    maccase(void);
398 static  void    maccolr(void);
399 static  void    macmcal(void);
400 static  void    macfram(void);
401 static  void    macmfra(void);
402 static  void    macrelq(void);
403 static  void    movmem(Inst*);
404 static  void    mid(Inst*, int, int);
405
406 extern void    das(ulong*, int);
407 extern void    _d2v(vlong *y, double d);
408
409 // Float constants
410 double double05 = 0.5;
411 double double4294967296 = 4294967296.0;
412
413 #define T(r)      *((void**)(R.r))
414
415 // The macro table. Macros are long sequences of instructions which come up often, like calls and returns,
416 // so they are extracted out into separate blocks. The calling convention is separate for each macro.
417 struct
418 {
419     int     idx;
420     void    (*gen)(void);
421     char*   name;
422 } mactab[] =
423 {
424     MacFRP,      macfrp,      "FRP", /* decrement and free pointer */
425     MacRET,      macret,      "RET", /* return instruction */
426     MacCASE,     maccase,     "CASE", /* case instruction */
427     MacCOLR,     maccolr,     "COLR", /* increment and color pointer */
428     MacMCAL,     macmcal,     "MCAL", /* mcall bottom half */
429     MacFRAM,     macfram,     "FRAM", /* frame instruction */
430     MacMFRA,     macmfra,     "MFRA", /* punt mframe because t->initialize=0 */
431     MacRELQ,     macrelq,     "RELQ", /* reschedule */
432 };
433
434
435 /* Helper functions */
436
437 void
438 urk(char *s)
439 {
440     iprint("urk: %s\n", s);
441     error(exCompile);
442 }
443
444 static void

```



```

445 gen(u32int o)
446 {
447     if (code < codestart || code >= codeend) {
448         iprint("gen: code out of bounds\n");
449         iprint("code:      0x%p\n", code);
450         iprint("codestart: 0x%p\n", codestart);
451         iprint("codeend:   0x%p\n", codeend);
452         //while (1) {}
453     }
454
455     *code++ = o;
456 }
457
458 static void
459 loadi(int reg, ulong val)
460 {
461     // Load a value into a register
462
463     // Check if the upper 20 bits are needed
464     if (SPLITIMM(val)) {
465         // Check if the lower 12 bits are needed
466         LUI(reg, SPLITH(val));
467         ADDI(reg, reg, SPLITL(val));
468     } else {
469         ADDI(reg, R0, val);
470     }
471 }
472
473 static void
474 multiply(int rd, int rs, long c)
475 {
476     // Multiply by a constant, rd = rs * c
477     int shamt;
478
479     if (c < 0) {
480         NEG(rd, rs);
481         rs = rd;
482         c = -c;
483     }
484
485     switch (c) {
486     case 0:
487         MOV(rd, R0);
488         break;
489     case 1:
490         if (rd != rs)
491             MOV(rd, rs);
492         break;
493     case 2:
494         shamt = 1;
495         goto shift;
496     case 3:
497         shamt = 1;
498         goto shiftadd;
499     case 4:
500         shamt = 2;
501         goto shift;
502     case 5:
503         shamt = 2;
504         goto shiftadd;
505     case 7:
506         shamt = 3;
507         goto shiftsub;
508     case 8:
509         shamt = 3;
510         goto shift;
511     case 16:
512         shamt = 4;
513         goto shift;
514     case 32:
515         shamt = 5;
516         goto shift;
517     case 64:
518         shamt = 6;
519         goto shift;
520     case 128:

```

```

521         shamt = 7;
522         goto shift;
523     case 256:
524         shamt = 8;
525         goto shift;
526     case 512:
527         shamt = 9;
528         goto shift;
529     case 1024:
530         shamt = 10;
531         goto shift;
532     shift:
533         SLLI(rd, rs, shamt);
534         break;
535     shiftadd:
536         if (rd == rs) {
537             MOV(Rtmp, rs);
538             rs = Rtmp;
539         }
540
541         SLLI(rd, rs, shamt);
542         ADD(rd, rd, rs);
543         break;
544     shiftsub:
545         if (rd == rs) {
546             MOV(Rtmp, rs);
547             rs = Rtmp;
548         }
549
550         SLLI(rd, rs, shamt);
551         SUB(rd, rd, rs);
552         break;
553     default:
554         loadi(Rtmp, c);
555         MUL(rd, rd, Rtmp);
556 }
557 }
558
559 static void
560 mem(int type, int r, int base, long offset)
561 {
562     // Load or store data at an offset from an address in a register.
563     // - type should be one of Ld* or St*.
564     // - r is the source or destination register.
565     // - base is the register with the base address.
566     // - offset is added to the value of rs to get the
567     // address to load/store from/to
568
569     if (SPLITIMM(offset)) {
570         // The offset is too long. Add the upper part of offset to rs in the tmp register,
571         // and use that as the base instead.
572         LUI(Rtmp, SPLITH(offset));
573         ADD(Rtmp, Rtmp, base);
574         base = Rtmp;
575         offset = SPLITL(offset);
576     }
577
578     switch (type) {
579     case Ldw:
580         LW(r, base, offset);
581         break;
582     case Ldh:
583         LH(r, base, offset);
584         break;
585     case Ldhu:
586         LHU(r, base, offset);
587         break;
588     case Ldb:
589         LB(r, base, offset);
590         break;
591     case Ldbu:
592         LBU(r, base, offset);
593         break;
594     case Lds:
595         FLW(r, base, offset);
596         break;

```

```

597     case Ldd:
598         FLD(r, base, offset);
599         break;
600     case Stw:
601         SW(r, base, offset);
602         break;
603     case Sth:
604         SH(r, base, offset);
605         break;
606     case Stb:
607         SB(r, base, offset);
608         break;
609     case Sts:
610         FSW(r, base, offset);
611         break;
612     case Std:
613         FSD(r, base, offset);
614         break;
615     case Laddr:
616         ADDI(r, base, offset);
617         break;
618     default:
619         if (cflag > 2)
620             iprint("Invalid type argument to mem: %d\n", type);
621         urk("mem");
622         break;
623 }
624 }
625
626 static void
627 operand(int mtype, int mode, Adr *a, int r, int li)
628 {
629     // Load or store the value from a src or dst operand of an instruction
630     // - mtype is the memory access type, as in mem
631     // - mode is the mode bits of the operand fields
632     // - a is the source or dest struct
633     // - r is the register to load the address into
634     int base;
635     long offset;
636
637     switch (mode) {
638     default:
639         urk("operand");
640     case AIMM:
641         // Immediate value
642         loadi(r, a->imm);
643
644         if (mtype == Laddr) {
645             mem(Stw, r, Rreg, li);
646             mem(Laddr, r, Rreg, li);
647         }
648         return;
649     case AFP:
650         // Indirect offset from FP
651         base = Rfp;
652         offset = a->ind;
653         break;
654     case AMP:
655         // Indirect offset from MP
656         base = Rmp;
657         offset = a->ind;
658         break;
659     case AIND|AFP:
660         // Double indirect from FP
661         mem(Ldw, Rta, Rfp, a->i.f);
662         base = Rta;
663         offset = a->i.s;
664         break;
665     case AIND|AMP:
666         // Double indirect from MP
667         mem(Ldw, Rta, Rmp, a->i.f);
668         base = Rta;
669         offset = a->i.s;
670         break;
671 }
672

```

```

673     mem(mtype, r, base, offset);
674 }
675
676 static void
677 op1(int mtype, Inst *i, int r)
678 {
679     // Load or store the source operand
680     operand(mtype, USRC(i->add), &i->s, r, O(REG, st));
681 }
682
683 static void
684 op3(int mtype, Inst *i, int r)
685 {
686     // Load or store the dest operand
687     operand(mtype, UDST(i->add), &i->d, r, O(REG, dt));
688 }
689
690 static void
691 op2(int mtype, Inst *i, int r)
692 {
693     // Load or store the middle operand
694     int ir;
695
696     switch (i->add&ARM) {
697     default:
698         return;
699     case AXIMM:
700         // Short immediate
701         loadi(r, (short) i->reg);
702
703         if (mtype == Laddr) {
704             mem(Stw, r, Rreg, O(REG, t));
705             mem(Laddr, r, Rreg, O(REG, t));
706         }
707         return;
708     case AXINF:
709         // Small offset from FP
710         ir = Rfp;
711         break;
712     case AXINM:
713         // Small offset from MP
714         ir = Rmp;
715         break;
716     }
717
718     // Load indirect
719     mem(mtype, r, ir, i->reg);
720 }
721
722 static void
723 literal(ulong imm, int roff)
724 {
725     // TODO: Why do this?
726     nlit++;
727
728     loadi(Rta, (ulong) litpool);
729     mem(Stw, Rta, Rreg, roff);
730
731     if (pass == 0)
732         return;
733
734     *litpool = imm;
735     litpool++;
736 }
737
738 static void
739 rdestroy(void)
740 {
741     destroy(R.s);
742 }
743
744 static void
745 rmcalls(void)
746 {
747     // Called by the compiled code to transfer control during an mcall
748     Frame *f;

```

```

749     Prog *p;
750
751     if (R.dt == (ulong) H)
752         error(exModule);
753
754     f = (Frame*)R.FP;
755     if (f == H)
756         error(exModule);
757
758     f->mr = nil;
759
760     ((void*)(Frame*))R.dt(f);
761
762     R.SP = (uchar*)f;
763     R.FP = f->fp;
764
765     if (f->t == nil)
766         unextend(f);
767     else
768         freeptrs(f, f->t);
769
770     p = currun();
771     if (p->kill != nil)
772         error(p->kill);
773 }
774
775 static void
776 rmfram(void)
777 {
778     Type *t;
779     Frame *f;
780     uchar *nsp;
781
782     if(R.d == H)
783         error(exModule);
784     t = (Type*)R.s;
785     if(t == H)
786         error(exModule);
787     nsp = R.SP + t->size;
788     if(nsp >= R.TS) {
789         R.s = t;
790         extend();
791         T(d) = R.s;
792         return;
793     }
794     f = (Frame*)R.SP;
795     R.SP = nsp;
796     f->t = t;
797     f->mr = nil;
798     initmem(t, f);
799     T(d) = f;
800 }
801
802 static void
803 bounds(void)
804 {
805     error(exBounds);
806 }
807
808 static void
809 nullity(void)
810 {
811     error(exNilref);
812 }
813
814 static void
815 punt(Inst *i, int m, void (*fn)(void))
816 {
817     ulong pc;
818     ulong *branch;
819
820     if (m & SRCOP) {
821         // Save the src operand in R->s
822         op1(Laddr, i, RA1);
823         mem(Stw, RA1, Rreg, 0(REG, s));
824     }

```

```

825
826     if (m & DSTOP) {
827         // Save the dst operand in R->d
828         op3(Laddr, i, RA3);
829         mem(Stw, RA3, Rreg, 0(REG, d));
830     }
831
832     if (m & WRTPC) {
833         // Store the PC in R->PC
834         loadi(RAO, RELPC(patch[i - mod->prog+1]));
835         mem(Stw, RAO, Rreg, 0(REG, PC));
836     }
837
838     if (m & DBRAN) {
839         // TODO: What does this do?
840         pc = patch[i->d.ins - mod->prog];
841         literal((ulong) (base+pc), 0(REG, d));
842     }
843
844     if ((i->add & ARM) == AXNON) {
845         if (m & THREOP) {
846             // R->m = R->d
847             mem(Ldw, RA2, Rreg, 0(REG, d));
848             mem(Stw, RA2, Rreg, 0(REG, m));
849         }
850     } else {
851         // R->m = middle operand
852         op2(Laddr, i, RA2);
853         mem(Stw, RA2, Rreg, 0(REG, m));
854     }
855
856     // R->FP = Rfp
857     mem(Stw, Rfp, Rreg, 0(REG, FP));
858
859     CALL(fn);
860
861     loadi(Rreg, (ulong) &R);
862
863     if (m & TCHECK) {
864         mem(Ldw, RAO, Rreg, 0(REG, t));
865
866         branch = code;
867         BEQZ(RAO, 0);
868
869         // If R->t != 0
870         mem(Ldw, Rlink, Rreg, 0(REG, xpc)); // Rlink = R->xpc
871         RETURN;
872
873         PATCHBRANCH(branch); // endif
874     }
875
876     mem(Ldw, Rfp, Rreg, 0(REG, FP));
877     mem(Ldw, Rmp, Rreg, 0(REG, MP));
878
879     if (m & NEWPC) {
880         // Jump to R->PC
881         mem(Ldw, RAO, Rreg, 0(REG, PC));
882         JR(RAO, 0);
883     }
884 }
885
886 static void
887 movloop(uint s)
888 {
889     // Move a section of memory in a loop.
890     // s is the size of each value, and should be 1, 2, or 4.
891     // The source address should be in RA1.
892     // The destination address should be in RA2.
893     // The amount of values to transfer should be in RA3
894     // All registers will be altered
895
896     ulong *loop;
897
898     if (s > 4 && s == 3) {
899         // Unnatural size. Transfer byte for byte
900         s = 1;

```

```

901     }
902
903     loop = code;
904     BEQZ(RA3, 0);
905
906     switch (s) {
907     case 0:
908         MOV(RA3, R0);
909         break;
910     case 1:
911         mem(Ldb, RA0, RA1, 0);
912         mem(Stb, RA0, RA2, 0);
913         break;
914     case 2:
915         mem(Ldh, RA0, RA1, 0);
916         mem(Ldh, RA0, RA2, 0);
917         break;
918     case 4:
919         mem(Ldw, RA0, RA1, 0);
920         mem(Ldw, RA0, RA2, 0);
921         break;
922     default:
923         urk("movloop");
924     }
925
926     ADDI(RA1, RA2, s);
927     ADDI(RA1, RA2, s);
928     ADDI(RA3, RA3, -s);
929
930     JABS(loop);
931
932     PATCHBRANCH(loop);
933 }
934
935 static void
936 movmem(Inst *i)
937 {
938     // Move a region of memory. Makes small transfers efficient, while defaulting
939     // to a move loop for larger transfers.
940     // The source address should be in RA1
941     ulong *branch;
942
943     if ((i->add & ARM) != AXIMM) {
944         op2(Ldw, i, RA3);
945
946         branch = code;
947         BEQ(RA3, R0, 0);
948
949         // if src2 != 0
950         movloop(1);
951         // endif
952
953         PATCHBRANCH(branch);
954         return;
955     }
956
957     switch (i->reg) {
958     case 0:
959         break;
960     case 4:
961         mem(Ldw, RA2, RA1, 0);
962         op3(Stw, i, RA2); // Save directly, don't bother loading the address
963         break;
964     case 8:
965         mem(Ldw, RA2, RA1, 0);
966         mem(Ldw, RA3, RA1, 4);
967
968         op3(Laddr, i, RA4);
969         mem(Stw, RA2, RA4, 0);
970         mem(Stw, RA3, RA4, 4);
971         break;
972     default:
973         op3(Laddr, i, RA2);
974
975         if ((i->reg & 3) == 0) {
976             loadi(RA3, i->reg >> 2);

```

```

977         movloop(4);
978     } else if ((i->reg & 1) == 0) {
979         loadi(RA3, i->reg >> 1);
980         movloop(2);
981     } else {
982         loadi(RA3, i->reg);
983         movloop(1);
984     }
985     break;
986 }
987 }
988
989 static void
990 movptr(Inst *i)
991 {
992     // Arguments:
993     // - RA1: The address to move from
994     // - op3: The address to move to
995
996     ulong *branch;
997
998     branch = code;
999     BEQ(RA1, Rh, 0);
1000
1001     // if RA1 != H
1002     CALLMAC(MacCOLR);           // colour if not H
1003     // endif
1004
1005     PATCHBRANCH(branch);
1006
1007     op3(Laddr, i, RA2);
1008     NOTNIL(RA2);
1009
1010     mem(Ldw, RA0, RA2, 0);
1011     mem(Stw, RA1, RA2, 0);
1012     CALLMAC(MacFRP);
1013 }
1014
1015 static void
1016 branch(Inst *i, int mtype, int btype)
1017 {
1018     // Insert a branch comparing integers
1019     // mtype should be the mtype to pass to mem to get the correct width
1020     // btype should be a constant like EQ, NE, LT, etc
1021     ulong *branch;
1022
1023     op2(mtype, i, RA1);
1024     op1(mtype, i, RA2);
1025
1026     branch = code;
1027
1028     // Invert the condition to skip the jump
1029     switch (btype) {
1030     case EQ:
1031         BNE(RA1, RA2, 0);
1032         break;
1033     case NE:
1034         BEQ(RA1, RA2, 0);
1035         break;
1036     case GT:
1037         BLE(RA1, RA2, 0);
1038         break;
1039     case LT:
1040         BGE(RA1, RA2, 0);
1041         break;
1042     case LE:
1043         BGT(RA1, RA2, 0);
1044         break;
1045     case GE:
1046         BLT(RA1, RA2, 0);
1047         break;
1048     }
1049
1050     iprint("branch: pc %d, branch to %d\n", ((ulong)i-(ulong)mod->prog), ((ulong)i->d.ins -
1051     ↵ (ulong)mod->prog));

```



```

1052         JDST(i);
1053
1054         PATCHBRANCH(branch);
1055     }
1056
1057     static void
1058     branchl(Inst *i, int btype)
1059     {
1060         // Insert a branch comparing 64-bit integers
1061         // btype should be a constant like EQ, NE, LT, etc
1062         ulong *branch;
1063
1064         op1(Laddr, i, RAO);
1065         mem(Ldw, RA1, RAO, 0);
1066         mem(Ldw, RA2, RAO, 4);
1067
1068         op2(Laddr, i, RAO);
1069         mem(Ldw, RA3, RAO, 0);
1070         mem(Ldw, RA4, RAO, 4);
1071
1072         // Set RA1 and RA2 to 1 if the condition holds
1073         switch (btype) {
1074             case EQ:
1075             case NE:
1076                 // RA1 = RA1 - RA3 == 0
1077                 // RA2 = RA2 - RA4 == 0
1078                 SUB(RA1, RA1, RA3);
1079                 SUB(RA2, RA2, RA4);
1080                 SLTU(RA1, R0, RA1);
1081                 SLTU(RA2, R0, RA2);
1082                 break;
1083             case LT:
1084             case GE:
1085                 // RA1 = RA1 < RA3
1086                 // RA2 = RA2 < RA4
1087                 SLT(RA1, RA1, RA3);
1088                 SLT(RA2, RA2, RA4);
1089                 break;
1090             case GT:
1091             case LE:
1092                 // RA1 = RA3 < RA1
1093                 // RA2 = RA4 < RA2
1094                 SLT(RA1, RA3, RA1);
1095                 SLT(RA2, RA4, RA2);
1096                 break;
1097         }
1098
1099         AND(RA1, RA1, RA2);
1100
1101         // Insert the branch. Negate to skip the jump
1102         // Have to negate again for NE, GE and LE
1103         branch = code;
1104         switch (btype) {
1105             case NE:
1106             case GE:
1107             case LE:
1108                 // If the negated condition holds, skip the jump
1109                 BNE(RA1, R0, 0);
1110                 break;
1111             default:
1112                 // If the condition doesn't hold, skip the jump
1113                 BEQ(RA1, R0, 0);
1114                 break;
1115         }
1116
1117         JDST(i);
1118
1119         PATCHBRANCH(branch);
1120     }
1121
1122     static void
1123     branchfd(Inst *i, int btype)
1124     {
1125         // Insert a branch comparing double-precision floats
1126         // btype should be a constant like EQ, NE, LT, etc
1127         ulong *branch;

```

```

1128
1129     op2(Ldd, i, F1);
1130     op1(Ldd, i, F2);
1131
1132     // Float compare instructions don't branch, so the branch
1133     // instruction has to check the result
1134     switch (btype) {
1135     case EQ:
1136     case NE:
1137         FEQD(RAO, F1, F2);
1138         break;
1139     case LT:
1140     case GE:
1141         FLTD(RAO, F1, F2);
1142         break;
1143     case LE:
1144     case GT:
1145         FLED(RAO, F1, F2);
1146         break;
1147     }
1148
1149     // Branch if the result is negative, skipping the jump
1150     branch = code;
1151     switch (btype) {
1152     case NE:
1153     case GE:
1154     case GT:
1155         BNE(RAO, R0, 0);
1156         break;
1157     default:
1158         BEQ(RAO, R0, 0);
1159         break;
1160     }
1161
1162     JDST(i);
1163
1164     PATCHBRANCH(branch);
1165 }
1166
1167 /* Macros */
1168 static void
1169 macfram(void)
1170 {
1171     // Allocate a mframe
1172     // Arguments:
1173     // - RA3: src1->links[src2]->t
1174
1175     ulong *branch;
1176
1177     mem(Ldw, RA2, Rreg, 0(REG, SP));           // RA2 = f = R.SP
1178     mem(Ldw, RA1, RA3, 0(Type, size));        // RA1 = src1->links[src2]->t->size
1179     ADD(RAO, RA2, RA1);                        // RAO = nsp = R.SP + t->size
1180     mem(Ldw, RA1, Rreg, 0(REG, TS));          // RA1 = R->TS
1181
1182     branch = code;
1183     BGEU(RAO, RA1, 0);
1184
1185     // nsp < R.TS
1186     mem(Stw, RA2, Rreg, 0(REG, SP));          // R.SP = nsp
1187
1188     mem(Stw, RA3, RA2, 0(Frame, t));          // f->t = RA3
1189     mem(Stw, R0, RA2, 0(Frame, mr));          // f->mr = 0
1190     mem(Ldw, Rta, RA3, 0(Type, initialize));
1191     JRL(Rta, 0);                               // call t->init(RA2)
1192
1193     // nsp >= R.TS; must expand
1194     PATCHBRANCH(branch);
1195     // Call extend. Store registers
1196     mem(Stw, RA3, Rreg, 0(REG, s));
1197     mem(Stw, Rlink, Rreg, 0(REG, st));
1198     mem(Stw, Rfp, Rreg, 0(REG, FP));
1199     CALL(extend);
1200
1201     // Restore registers
1202     loadi(Rreg, (ulong) &R);
1203     mem(Ldw, Rlink, Rreg, 0(REG, st));

```

```

1204     mem(Ldw, Rfp, Rreg, 0(REG, FP));
1205     mem(Ldw, Rmp, Rreg, 0(REG, MP));
1206     mem(Ldw, RA2, Rreg, 0(REG, s));
1207     RETURN;
1208 }
1209
1210 static void
1211 macmfra(void)
1212 {
1213     mem(Stw, Rlink, Rreg, 0(REG, st));
1214     mem(Stw, RA3, Rreg, 0(REG, s)); // Save type
1215     mem(Stw, RA0, Rreg, 0(REG, d)); // Save destination
1216     mem(Stw, Rfp, Rreg, 0(REG, FP));
1217
1218     CALL(rmfram);
1219
1220     loadi(Rreg, (ulong)&R);
1221     mem(Ldw, Rlink, Rreg, 0(REG, st));
1222     mem(Ldw, Rfp, Rreg, 0(REG, FP));
1223     mem(Ldw, Rmp, Rreg, 0(REG, MP));
1224
1225     RETURN;
1226 }
1227
1228 static void
1229 macmcal(void)
1230 {
1231     // The bottom half of a mcall instruction
1232     // Calling convention:
1233     // - RA0: The address of the function to jump to
1234     // - RA2: The frame address, src1 to mcall
1235     // - RA3: The module reference, src3 to mcall
1236
1237     ulong *branch1, *branch2, *branch3;
1238
1239     branch1 = code;
1240     BEQ(RA0, Rh, 0);
1241     // If RA0 != H
1242
1243     mem(Ldw, RA1, RA3, 0(Modlink, prog)); // Load m->prog into RA1
1244
1245     branch2 = code;
1246     BNEZ(RA1, 0);
1247     // If m->prog == 0
1248
1249     mem(Stw, Rlink, Rreg, 0(REG, st)); // Store link register
1250     mem(Stw, RA2, Rreg, 0(REG, FP)); // Store FP register
1251     mem(Stw, RA0, Rreg, 0(REG, dt)); // Store destination address
1252
1253     CALL(rmcall);
1254
1255     // After the call has returned
1256     loadi(Rreg, (ulong)&R); // Load R
1257     mem(Ldw, Rlink, Rreg, 0(REG, st)); // Load link register
1258     mem(Ldw, Rfp, Rreg, 0(REG, FP)); // Load FP register
1259     mem(Ldw, Rmp, Rreg, 0(REG, MP)); // Load MP register
1260     RETURN;
1261
1262     // else
1263     PATCHBRANCH(branch1); // If RA0 == H
1264     PATCHBRANCH(branch2); // If m->prog != 0
1265
1266     MOV(Rfp, RA2); // Rfp = RA2
1267     mem(Stw, RA3, Rreg, 0(REG, M)); // R.M = RA3
1268
1269     // D2H(RA3)->ref++
1270     ulong heapref = 0(Heap, ref) - sizeof(Heap);
1271     mem(Ldw, RA1, RA3, heapref);
1272     ADDI(RA1, RA1, 1);
1273     mem(Stw, RA1, RA3, heapref);
1274
1275     mem(Ldw, Rmp, RA3, 0(Modlink, MP)); // Rmp = R.M->mp
1276     mem(Stw, Rmp, Rreg, 0(REG, MP)); // R.MP = Rmp
1277
1278     mem(Ldw, RA1, RA3, 0(Modlink, compiled));
1279     branch3 = code;

```

```

1280     BNEZ(RA1, 0);
1281
1282     // if M.compiled == 0
1283     mem(Stw, Rfp, Rreg, 0(REG, FP)); // R.FP = Rfp
1284     mem(Stw, RAO, Rreg, 0(REG, PC)); // R.PC = Rpc
1285     mem(Ldw, Rlink, Rreg, 0(REG, xpc));
1286     RETURN; // Leave it to the interpreter to handle
1287
1288     // else
1289     PATCHBRANCH(branch3);
1290     JR(RAO, 0); // Jump to the compiled module
1291 }
1292
1293 static void
1294 maccase(void)
1295 {
1296     /*
1297     * RA1 = value (input arg), v
1298     * RA2 = count, n
1299     * RA3 = table pointer (input arg), t
1300     * RA0 = n/2, n2
1301     * RA4 = pivot element t+n/2*3, l
1302     */
1303
1304     ulong *loop, *found, *branch;
1305
1306     mem(Ldw, RA2, RA3, 0); // get count from table
1307     MOV(Rlink, RA3); // initial table pointer
1308
1309     loop = code;
1310     BLEZ(RA2, 0); // n <= 0? goto out
1311
1312     SRAI(RA0, RA2, 1); // n2 = n>>1
1313
1314     // l = t + n/2*3
1315     ADD(RA4, RA0, RA2); // l = n/2 + n
1316     ADD(RA4, RA3, RA1); // l += t
1317
1318     mem(Ldw, Rta, RA4, 4); // Rta = l[1]
1319     branch = code;
1320     BGE(RA1, Rta, 0);
1321
1322     // if v < l[1]
1323     MOV(RA2, RA0); // n = n2
1324     JABS(loop); // continue
1325
1326     // if v >= l[1]
1327     PATCHBRANCH(branch);
1328     mem(Ldw, Rta, RA4, 8); // Rta = l[2]
1329     found = code;
1330     BLT(RA1, Rta, 0); // branch to found
1331
1332     // if v >= l[2]
1333     ADDI(RA3, RA4, 12); // t = l+3
1334     SUB(RA2, RA2, RA0); // n -= n2
1335     ADDI(RA2, RA2, -1); // n -= 1
1336
1337     JABS(loop); // goto loop
1338
1339     // endloop
1340
1341     // found: v >= l[1] && v < l[2]
1342     // jump to l[3]
1343     PATCHBRANCH(found);
1344     JR(RA4, 12);
1345
1346     // out: Loop ended
1347     PATCHBRANCH(loop);
1348     mem(Ldw, RA2, Rlink, 0); // load initial n
1349     ADD(Rtmp, RA2, RA2); // Rtmp = 2*n
1350     ADD(RA2, RA2, Rtmp); // n = 3*n
1351
1352     // goto (initial t)[n*3+1]
1353     SLLI(RA2, RA2, 2); // RA2 = n*sizeof(long)
1354     ADD(Rlink, Rlink, RA2); // Rlink = t[n*3]
1355     JR(Rlink, 4); // goto Rlink+4 = t[n*3+1]

```

```

1356 }
1357
1358 static void
1359 maccolr(void)
1360 {
1361     // Color a pointer
1362     // Arguments:
1363     // - RA1: The pointer to color
1364     ulong *branch;
1365
1366     // h->ref++
1367     mem(Ldw, RA0, RA1, 0(Heap, ref) - sizeof(Heap));
1368     ADDI(RA0, RA0, 1);
1369     mem(Stw, RA0, RA1, 0(Heap, ref) - sizeof(Heap));
1370
1371     // RA0 = mutator
1372     mem(Ldw, RA0, RA1, 0(Heap, color) - sizeof(Heap));
1373
1374     // RA2 = h->color
1375     loadi(RA2, (ulong) &mutator);
1376     mem(Ldw, RA2, RA2, 0);
1377
1378     branch = code;
1379     BEQ(RA0, RA2, 0);
1380
1381     // if h->color != mutator
1382
1383     // h->color = propagator
1384     loadi(RA2, propagator);
1385     mem(Stw, RA2, RA1, 0(Heap, color) - sizeof(Heap));
1386
1387     // nprop = RA1
1388     loadi(RA2, (ulong) &nprop);
1389     mem(Stw, RA1, RA2, 0);
1390
1391     // endif
1392     PATCHBRANCH(branch);
1393     RETURN;
1394 }
1395
1396 static void
1397 macfrp(void)
1398 {
1399     // Destroy a pointer
1400     // Arguments:
1401     // - RA0: The pointer to destroy
1402     ulong *branch1, *branch2;
1403
1404     branch1 = code;
1405     BEQ(RA0, Rh, 0);
1406
1407     // if RA0 != H
1408     mem(Ldw, RA2, RA0, 0(Heap, ref) - sizeof(Heap));
1409     ADDI(RA2, RA2, -1);
1410
1411     branch2 = code;
1412     BEQ(RA2, R0, 0);
1413
1414     // if --h->ref != 0
1415     mem(Stw, RA2, RA0, 0(Heap, ref) - sizeof(Heap));
1416     RETURN;
1417     // endif
1418
1419     PATCHBRANCH(branch2);
1420     mem(Stw, Rfp, Rreg, 0(REG, FP));
1421     mem(Stw, Rlink, Rreg, 0(REG, st));
1422     mem(Stw, RA0, Rreg, 0(REG, s));
1423     CALL(rdestroy);
1424
1425     loadi(Rreg, (ulong) &R);
1426     mem(Ldw, Rlink, Rreg, 0(REG, st));
1427     mem(Ldw, Rfp, Rreg, 0(REG, FP));
1428     mem(Ldw, Rmp, Rreg, 0(REG, MP));
1429
1430     // endif
1431     PATCHBRANCH(branch1);

```

```

1432     RETURN;
1433 }
1434
1435 static void
1436 macret(void)
1437 {
1438     Inst i;
1439     ulong *branch1, *branch2, *branch3, *branch4, *branch5, *branch6;
1440
1441     branch1 = code;
1442     BEQ(RA1, R0, 0);
1443
1444     // if t(Rfp) != 0
1445     mem(Ldw, RA0, RA1, 0(Type, destroy));
1446     branch2 = code;
1447     BEQ(RA0, R0, 0);
1448
1449     // if destroy(t(fp)) != 0
1450     mem(Ldw, RA2, Rfp, 0(Frame, fp));
1451     branch3 = code;
1452     BEQ(RA2, R0, 0);
1453
1454     // if fp(Rfp) != 0
1455     mem(Ldw, RA3, Rfp, 0(Frame, mr));
1456     branch4 = code;
1457     BEQ(RA3, R0, 0);
1458
1459     // if mr(Rfp) != 0
1460     mem(Ldw, RA2, Rreg, 0(REG, M));
1461     mem(Ldw, RA3, RA2, 0(Heap, ref) - sizeof(Heap));
1462     ADDI(RA3, RA3, -1);
1463
1464     branch5 = code;
1465     BEQ(RA3, R0, 0);
1466
1467     // if --ref(arg) != 0
1468     mem(Stw, RA3, RA2, 0(Heap, ref) - sizeof(Heap));
1469     mem(Ldw, RA1, Rfp, 0(Frame, mr));
1470     mem(Stw, RA1, Rreg, 0(REG, M));
1471     mem(Ldw, Rmp, RA1, 0(Modlink, MP));
1472     mem(Stw, Rmp, Rreg, 0(REG, MP));
1473
1474     mem(Ldw, RA3, RA1, 0(Modlink, compiled));
1475     branch6 = code;
1476     BEQ(RA3, R0, 0);
1477
1478     // This part is a bit weird, because it should be the innermost
1479     // if-statement (in C terms), but the else of branch4 also ends up here.
1480     // This could be a mistake, but it's in at least the ARM and MIPS version.
1481
1482     // if R.M->compiled != 0
1483     // if mr(Rfp) == 0
1484     PATCHBRANCH(branch4);
1485     JRL(RA0, 0); // Call destroy(t(fp))
1486
1487     mem(Stw, Rfp, Rreg, 0(REG, SP)); // R->SP = Rfp
1488     mem(Ldw, RA1, Rfp, 0(Frame, lr)); // RA1 = Rfp->lr
1489     mem(Ldw, Rfp, Rfp, 0(Frame, fp)); // Rfp = Rfp->fp
1490     mem(Stw, Rfp, Rreg, 0(REG, FP)); // R->FP = Rfp
1491
1492     JR(RA1, 0); // goto RA1, if compiled
1493     // does not continue past here
1494
1495     // if R.M->compiled == 0
1496     PATCHBRANCH(branch6);
1497     JRL(RA0, 0); // Call destroy(t(fp))
1498
1499     mem(Stw, Rfp, Rreg, 0(REG, SP)); // R->SP = Rfp
1500     mem(Ldw, RA1, Rfp, 0(Frame, lr)); // RA1 = Rfp->lr
1501     mem(Ldw, Rfp, Rfp, 0(Frame, fp)); // Rfp = Rfp->fp
1502     mem(Stw, RA1, Rreg, 0(REG, PC)); // R.PC = RA1
1503     mem(Ldw, Rlink, Rreg, 0(REG, xpc)); // Rlink = R->xpc
1504     RETURN; // return to xec uncompiled code
1505
1506     // endif
1507     PATCHBRANCH(branch5);

```

```

1508     PATCHBRANCH(branch3);
1509     PATCHBRANCH(branch2);
1510     PATCHBRANCH(branch1);
1511
1512     i.add = AXNON;
1513     punt(&i, TCHECK|NEWPC, optab[IRET]);
1514 }
1515
1516 static void
1517 macrelq(void)
1518 {
1519     // Store frame pointer and link register, then return to xev
1520     mem(Stw, Rfp, Rreg, 0(REG, FP));
1521     mem(Stw, Rlink, Rreg, 0(REG, PC));
1522     mem(Ldw, Rlink, Rreg, 0(REG, xpc));
1523     RETURN;
1524 }
1525
1526 /* Main compilation functions */
1527 static void
1528 comi(Type *t)
1529 {
1530     // Compile a type initializer
1531     int i, j, m, c;
1532
1533     for (i = 0; i < t->np; i++) {
1534         c = t->map[i];
1535         j = i << 5;
1536
1537         for (m = 0x80; m != 0; m >>= 1) {
1538             if (c & m)
1539                 mem(Stw, Rh, RA2, j);
1540
1541             j += sizeof(WORD*);
1542         }
1543     }
1544
1545     RETURN;
1546 }
1547
1548 static void
1549 comd(Type *t)
1550 {
1551     // Compile a type destructor
1552     int i, j, m, c;
1553
1554     mem(Stw, Rlink, Rreg, 0(REG, dt));
1555
1556     for (i = 0; i < t->np; i++) {
1557         c = t->map[i];
1558         j = i << 5;
1559
1560         for (m = 0x80; m != 0; m >>= 1) {
1561             if (c & m) {
1562                 mem(Ldw, RA0, Rfp, j);
1563                 CALL(base+macro[MacFRP]);
1564             }
1565
1566             j += sizeof(WORD*);
1567         }
1568     }
1569
1570     mem(Ldw, Rlink, Rreg, 0(REG, dt));
1571     RETURN;
1572 }
1573
1574 static void
1575 typecom(Type *t)
1576 {
1577     // Compile a type
1578     int n;
1579     ulong *tmp, *start;
1580
1581     if (t == nil | t->initialize != 0)
1582         return;
1583

```

```

1584     tmp = mallocz(4096*sizeof(ulong), 0);
1585     if (tmp == nil)
1586         error(exNomem);
1587
1588     codestart = tmp;
1589     codeend = tmp + 4096;
1590     iprint("Typecom np %d, size %d\n", t->np, t->size);
1591     code = tmp;
1592     comi(t);
1593     n = code - tmp;
1594     code = tmp;
1595     comd(t);
1596     n += code - tmp;
1597     free(tmp);
1598
1599     n *= sizeof(*code);
1600     code = mallocz(n, 0);
1601     if (code == nil)
1602         return;
1603
1604     codestart = code;
1605     codeend = code + n;
1606
1607     start = code;
1608     t->initialize = code;
1609     comi(t);
1610     t->destroy = code;
1611     comd(t);
1612
1613     segflush(start, n);
1614
1615     if (cflag > 3)
1616         iprint("typ= %.8p %4d i %.8p d %.8p asm=%d\n",
1617             t, t->size, t->initialize, t->destroy, n);
1618
1619     if (cflag > 6) {
1620         das(start, code-start);
1621     }
1622 }
1623
1624 static void
1625 patchex(Module *m, ulong *p)
1626 {
1627     // Apply patches for a module. p is the patch array
1628     Handler *h;
1629     Except *e;
1630
1631     for (h = m->htab; h != nil && h->etab != nil; h++) {
1632         h->pc1 = p[h->pc1];
1633         h->pc2 = p[h->pc2];
1634
1635         for (e = h->etab; e->s != nil; e++)
1636             e->pc = p[e->pc];
1637
1638         if (e->pc != -1)
1639             e->pc = p[e->pc];
1640     }
1641 }
1642
1643 static void
1644 commframe(Inst *i)
1645 {
1646     // Compile a mframe instruction
1647     ulong *branch1, *branch2;
1648     loadi(R7, 0);
1649
1650     op1(Ldw, i, RAO);
1651     branch1 = code;
1652     BEQ(RAO, Rh, 0);
1653
1654     // if RAO != H
1655
1656     // RA3 = src->links[src2]->frame
1657     if ((i->add & ARM) == AXIMM) {
1658         mem(Ldw, RA3, RAO, 0A(Modlink, links) + i->reg*sizeof(Modl) + 0(Modl, frame));
1659     } else {

```



```

1660         // RA1 = src->links[src2]
1661         op2(Ldw, i, RA1);
1662         multiply(RA1, RA1, sizeof(Modl));
1663         ADD(RA1, RA1, RA0);
1664
1665         // RA3 = src->links[src2]->frame
1666         mem(Ldw, RA3, RA1, 0(Modl, frame));
1667     }
1668
1669     mem(Ldw, RA1, RA3, 0(Type, initialize));
1670     branch2 = code;
1671     BNEZ(RA1, 0);
1672
1673     // if frame->initialize == 0
1674     op3(Laddr, i, RA0);
1675     // endif
1676
1677     // if RA0 == H || frame->initialize == 0
1678     PATCHBRANCH(branch1);
1679     loadi(Rlink, RELPC(patch[i - mod->prog + 1]));
1680     loadi(R7, 7);
1681     CALLMAC(MacMFRA);
1682
1683     // if frame->inititalize != 0
1684     PATCHBRANCH(branch2);
1685     loadi(R7, 8);
1686     CALLMAC(MacFRAM);
1687     op3(Stw, i, RA2);
1688 }
1689
1690 static void
1691 commcall(Inst *i)
1692 {
1693     // Compile a mcall instruction
1694     ulong *branch;
1695
1696     op1(Ldw, i, RA2); // RA2 = src1 = frame
1697     loadi(RA0, RELPC(patch[i - mod->prog+1])); // RA0 = pc
1698     mem(Stw, RA0, RA2, 0(Frame, lr)); // frame.lr = RA0 = pc
1699     mem(Stw, Rfp, RA2, 0(Frame, fp)); // frame.fp = fp
1700     mem(Ldw, RA3, Rreg, 0(REG, M)); // RA3 = R.M
1701     mem(Stw, RA3, RA2, 0(Frame, mr)); // frame.mr = RA3 = R.M
1702
1703     op3(Ldw, i, RA3); // RA3 = src3 = Modlink
1704
1705     branch = code;
1706     BEQ(RA3, Rh, 0);
1707     // If RA3 != H
1708
1709     // RA0 = Modlink->links[src2]->pc
1710     if ((i->add&ARM) == AXIMM) {
1711         // i->reg contains the immediate of src2, don't have to store it in a register
1712         mem(Ldw, RA0, RA3, 0A(Modlink, links) + i->reg*sizeof(Modl) + 0(Modl, u.pc));
1713     } else {
1714         op2(Ldw, i, RA1); // RA1 = src2
1715
1716         // RA1 *= sizeof(Modl)
1717         multiply(RA1, RA1, sizeof(Modl));
1718
1719         ADDI(RA1, RA1, RA3);
1720         mem(Ldw, RA0, RA1, 0A(Modlink, links) + 0(Modl, u.pc));
1721     }
1722
1723     PATCHBRANCH(branch); // endif
1724
1725     CALLMAC(MacMCAL);
1726 }
1727
1728 static void
1729 comcase(Inst *i, int w)
1730 {
1731     // Compile a case instruction
1732     int l;
1733     WORD *t, *e;
1734
1735     if (w != 0) {

```

```

1736         // Use the MacCASE macro
1737         op1(Ldw, i, RA1);
1738         op3(Laddr, i, RA3);
1739         CALLMAC(MacCASE);
1740     }
1741
1742     // Get a pointer to the table
1743     t = (WORD*)(mod->origmp + i->d.ind+4);
1744
1745     // Get the flag right before the table
1746     l = t[-1];
1747
1748     /* have to take care not to relocate the same table twice -
1749     * the limbo compiler can duplicate a case instruction
1750     * during its folding phase
1751     */
1752
1753     if (pass == 0) {
1754         if (l >= 0)
1755             t[-1] = -1-1; /* Mark it not done */
1756         return;
1757     }
1758
1759     if (l >= 0) { /* Check pass 2 done */
1760         return;
1761     }
1762
1763     t[-1] = -1-1; /* Set real count */
1764     e = t + t[-1]*3;
1765
1766     while (t < e) {
1767         t[2] = RELPC(patch[t[2]]);
1768         t += 3;
1769     }
1770
1771     t[0] = RELPC(patch[t[0]]);
1772 }
1773
1774 static void
1775 comcasel(Inst *i)
1776 {
1777     // Same as comcase, but with double words
1778     int l;
1779     WORD *t, *e;
1780
1781     t = (WORD*) (mod->origmp + i->d.ind + 8);
1782     l = t[-2];
1783
1784     if (pass == 0) {
1785         if (l >= 0)
1786             t[-2] = -1-1; /* Mark it not done */
1787         return;
1788     }
1789
1790     if (l >= 0) /* Check pass 2 done */
1791         return;
1792
1793     t[-2] = -1-1; /* Set real count */
1794     e = t + t[-2]*6;
1795
1796     while (t < e) {
1797         t[4] = RELPC(patch[t[4]]);
1798         t += 6;
1799     }
1800
1801     t[0] = RELPC(patch[t[0]]);
1802 }
1803
1804 static void
1805 comgoto(Inst *i)
1806 {
1807     // Compile a goto instruction
1808     WORD *t, *e;
1809
1810     op1(Ldw, i, RA1); /* RA1 = src
1811     op3(Laddr, i, RA0); /* RA0 = &dst

```

```

1812     SLLI(RA1, RA1, 2);           // RA1 = src*sizeof(int)
1813     ADD(RA1, RA1, RAO);         // RA1 += RAO
1814     mem(Ldw, RAO, RA1, 0);     // RAO = dst[src]
1815     JR(RAO, 0);                // goto dst[src]
1816
1817     if (pass == 0)
1818         return;
1819
1820     t = (WORD*)(mod->origmp+i->d.ind);
1821     e = t + t[-1];
1822     t[-1] = 0;
1823
1824     while (t < e) {
1825         t[0] = RELPC(patch[t[0]]);
1826         t++;
1827     }
1828 }
1829
1830 static void
1831 comp(Inst *i)
1832 {
1833     // Compile a single DIS instruction
1834     char buf[64];
1835     ulong *branch1, *branch2, *loop;
1836
1837     switch (i->op) {
1838     default:
1839         sprintf(buf, sizeof buf, "%s compile, no '%D'", mod->name, i);
1840         error(buf);
1841         break;
1842     case IMCALL:
1843         commcall(i);
1844         break;
1845     case ISEND:
1846     case IRECV:
1847     case IALT:
1848         punt(i, SRCOP|DSTOP|TCHECK|WRTPC, optab[i->op]);
1849         break;
1850     case ISPAWN:
1851         punt(i, SRCOP|DBRAN, optab[i->op]);
1852         break;
1853     case IBNEC:
1854     case IBEQC:
1855     case IBLTC:
1856     case IBLEC:
1857     case IBGTC:
1858     case IBGEC:
1859         punt(i, SRCOP|DBRAN|NEWPC|WRTPC, optab[i->op]);
1860         break;
1861     case ICASEC:
1862         comcase(i, 0);
1863         punt(i, SRCOP|DSTOP|NEWPC, optab[i->op]);
1864         break;
1865     case ICASEL:
1866         comcase1(i);
1867         punt(i, SRCOP|DSTOP|NEWPC, optab[i->op]);
1868         break;
1869     case IADDC:
1870     case IMULL:
1871     case IDIVL:
1872     case IMODL:
1873     case IMNEWZ:
1874     case ILSRW:
1875     case ILSRL:
1876         punt(i, SRCOP|DSTOP|THREOP, optab[i->op]);
1877         break;
1878     case IMODW:
1879         op1(Ldw, i, RA1);
1880         op2(Ldw, i, RAO);
1881         REM(RAO, RAO, RA1);
1882         op3(Stw, i, RAO);
1883         break;
1884     case IMODB:
1885         op1(Ldb, i, RA1);
1886         op2(Ldb, i, RAO);
1887         REM(RAO, RAO, RA1);

```

```

1888         op3(Stb, i, RAO);
1889         break;
1890     case IDIVW:
1891         op1(Ldw, i, RA1);
1892         op2(Ldw, i, RAO);
1893         DIV(RAO, RAO, RA1);
1894         op3(Stw, i, RAO);
1895         break;
1896     case IDIVE:
1897         op1(Ldb, i, RA1);
1898         op2(Ldb, i, RAO);
1899         DIV(RAO, RAO, RA1);
1900         op3(Stb, i, RAO);
1901         break;
1902     case ILOAD:
1903     case INEW:
1904     case INEWAZ:
1905     case INEW:
1906     case INEWZ:
1907     case ISLICEA:
1908     case ISLICELA:
1909     case ICONSB:
1910     case ICONSW:
1911     case ICONSL:
1912     case ICONSF:
1913     case ICONSM:
1914     case ICONSMP:
1915     case ICONSP:
1916     case IMOVMP:
1917     case IHEADMP:
1918     case IHEADB:
1919     case IHEADW:
1920     case IHEADL:
1921     case IINSC:
1922     case ICVTAC:
1923     case ICVTCW:
1924     case ICVTWC:
1925     case ICVTLC:
1926     case ICVTCL:
1927     case ICVTFC:
1928     case ICVTCF:
1929     case ICVTRF:
1930     case ICVTFR:
1931     case ICVTWS:
1932     case ICVTSW:
1933     case IMSPAWN:
1934     case ICVTCA:
1935     case ISLICEC:
1936     case INBALT:
1937         punt(i, SRCOP|DSTOP, optab[i->op]);
1938         break;
1939     case INEWCM:
1940     case INEWCMP:
1941         punt(i, SRCOP|DSTOP|THREOP, optab[i->op]);
1942         break;
1943     case IMFRAME:
1944         commframe(i);
1945         break;
1946     case ICASE:
1947         comcase(i, 1);
1948         break;
1949     case IGOTO:
1950         comgoto(i);
1951         break;
1952     case IMOVF:
1953         op1(Ldd, i, F1);
1954         op3(Std, i, F1);
1955         break;
1956     case IMOVL:
1957         op1(Laddr, i, RAO);
1958         mem(Ldw, RA1, RAO, 0);
1959         mem(Ldw, RA2, RAO, 4);
1960
1961         op3(Laddr, i, RAO);
1962         mem(Stw, RA1, RAO, 0);
1963         mem(Stw, RA2, RAO, 4);

```

```

1964         break;
1965     case IHEADM:
1966         punt(i, SRCOP|DSTOP, optab[i->op]);
1967         break;
1968         op1(Laddr, i, RA1);
1969         NOTNIL(RA1);
1970
1971         if(OA(List, data) != 0) {
1972             ADDI(RA1, RA1, OA(List, data));
1973         }
1974
1975         movmem(i);
1976         break;
1977     case IMOVM:
1978         punt(i, SRCOP|DSTOP|THREOP, optab[i->op]);
1979         break;
1980         op1(Laddr, i, RA1);
1981         movmem(i);
1982         break;
1983     case IFRAME:
1984         if(UXSRC(i->add) != SRC(AIMM)) {
1985             punt(i, SRCOP|DSTOP, optab[i->op]);
1986             break;
1987         }
1988         tinit[i->s.imm] = 1;
1989         loadi(RA3, (ulong) mod->type[i->s.imm]);
1990         CALL(base+macro[MacFRAM]);
1991         op3(Stw, i, RA2);
1992         break;
1993     case INEWCB:
1994     case INEWCW:
1995     case INEWCF:
1996     case INEWCP:
1997     case INEWCL:
1998         punt(i, DSTOP|THREOP, optab[i->op]);
1999         break;
2000     case IEXIT:
2001         punt(i, 0, optab[i->op]);
2002         break;
2003     case ICVTBW:
2004         op1(Ldbu, i, RA0);
2005         op3(Stw, i, RA0);
2006         break;
2007     case ICVTWB:
2008         op1(Ldw, i, RA0);
2009         op3(Stb, i, RA0);
2010         break;
2011     case ILEA:
2012         op1(Laddr, i, RA0);
2013         op3(Stw, i, RA0);
2014         break;
2015     case IMOVW:
2016         op1(Ldw, i, RA0);
2017         op3(Stw, i, RA0);
2018         break;
2019     case IMOVb:
2020         op1(Ldb, i, RA0);
2021         op3(Stb, i, RA0);
2022         break;
2023     case ITAIL:
2024         punt(i, SRCOP|DSTOP, optab[i->op]);
2025         break;
2026         op1(Ldw, i, RA0);
2027         NOTNIL(RA0);
2028         mem(Ldw, RA1, RA0, 0(List, tail));
2029         movptr(i);
2030         break;
2031     case IMOVP:
2032         punt(i, SRCOP|DSTOP, optab[i->op]);
2033         break;
2034         op1(Ldw, i, RA1);
2035         NOTNIL(RA1);
2036         movptr(i);
2037         break;
2038     case IHEADP:
2039         punt(i, SRCOP|DSTOP, optab[i->op]);

```

```

2040         break;
2041         op1(Ldw, i, RAO);
2042         NOTNIL(RAO);
2043         mem(Ldw, OA(List, data), RAO, RA1);
2044         movptr(i);
2045         break;
2046     case ILENA:
2047         punt(i, SRCOP|DSTOP, optab[i->op]);
2048         break;
2049         op1(Ldw, i, RA1);
2050         MOV(RAO, RO);
2051
2052         branch1 = code;
2053         BEQ(RA1, Rh, 0);
2054
2055         // if src != H
2056         mem(Ldw, RAO, RA1, O(Array, len));
2057         // endif
2058
2059         PATCHBRANCH(branch1);
2060         op3(Stw, i, RAO);
2061         break;
2062     case ILENC:
2063         punt(i, SRCOP|DSTOP, optab[i->op]);
2064         break;
2065         op1(Ldw, i, RA1);
2066         MOV(RAO, RO);
2067
2068         branch1 = code;
2069         BEQ(RA1, Rh, 0);
2070
2071         // if RA1 != H
2072         mem(Ldw, RAO, RA1, O(String, len));
2073
2074         branch2 = code;
2075         BGE(RAO, 0, 0);
2076
2077         // if string->len < 0
2078
2079         // RAO = abs(string->len)
2080         NEG(RAO, RAO);
2081
2082         // endif
2083
2084         PATCHBRANCH(branch1);
2085         PATCHBRANCH(branch2);
2086         op3(Stw, i, RAO);
2087         break;
2088     case ILENL:
2089         punt(i, SRCOP|DSTOP, optab[i->op]);
2090         break;
2091         MOV(RAO, RO); // RAO = 0
2092         op1(Ldw, i, RA1); // RA1 = src
2093
2094         // while RA1 != H
2095         loop = code;
2096         BEQ(RA1, Rh, 0);
2097
2098         mem(Ldw, RA1, RA1, O(List, tail)); // RAO = RAO->tail
2099         ADDI(RAO, RAO, 1); // RA1++
2100         JABS(loop);
2101         // endwhile
2102
2103         PATCHBRANCH(loop);
2104         op3(Stw, i, RAO); // return RA1
2105         break;
2106     case ICALL:
2107         op1(Ldw, i, RAO);
2108         loadi(RA1, RELPC(patch[i - mod->prog + 1]));
2109         mem(Stw, RA1, RAO, O(Frame, lr));
2110         mem(Stw, Rfp, RAO, O(Frame, fp));
2111         MOV(Rfp, RAO);
2112         JDST(i);
2113         break;
2114     case IJMP:
2115         JDST(i);

```

```

2116         break;
2117     case IBEQW:
2118         branch(i, Ldw, EQ);
2119         break;
2120     case IBNEW:
2121         branch(i, Ldw, NE);
2122         break;
2123     case IBLTW:
2124         branch(i, Ldw, LT);
2125         break;
2126     case IBLEW:
2127         branch(i, Ldw, LE);
2128         break;
2129     case IBGTW:
2130         branch(i, Ldw, GT);
2131         break;
2132     case IBGEW:
2133         branch(i, Ldw, GE);
2134         break;
2135     case IBEQB:
2136         branch(i, Ldb, EQ);
2137         break;
2138     case IBNEB:
2139         branch(i, Ldb, NE);
2140         break;
2141     case IBLTB:
2142         branch(i, Ldb, LT);
2143         break;
2144     case IBLEB:
2145         branch(i, Ldb, LE);
2146         break;
2147     case IBGTB:
2148         branch(i, Ldb, GT);
2149         break;
2150     case IBGEB:
2151         branch(i, Ldb, GE);
2152         break;
2153     case IBEQF:
2154         branchfd(i, EQ);
2155         break;
2156     case IBNEF:
2157         branchfd(i, NE);
2158         break;
2159     case IBLTF:
2160         branchfd(i, LT);
2161         break;
2162     case IBLEF:
2163         branchfd(i, LE);
2164         break;
2165     case IBGTF:
2166         branchfd(i, GT);
2167         break;
2168     case IBGEF:
2169         branchfd(i, GE);
2170         break;
2171     case IRET:
2172         mem(Ldw, RA1, Rfp, 0(Frame, t));
2173         CALLMAC(MacRET);
2174         break;
2175     case IMULW:
2176         op1(Ldw, i, RA1);
2177         op2(Ldw, i, RAO);
2178         MUL(RAO, RAO, RA1);
2179         op3(Stw, i, RAO);
2180         break;
2181     case IMULB:
2182         op1(Ldb, i, RA1);
2183         op2(Ldb, i, RAO);
2184         MUL(RAO, RAO, RA1);
2185         op3(Stb, i, RAO);
2186         break;
2187     case IORW:
2188         op1(Ldw, i, RA1);
2189         op2(Ldw, i, RA2);
2190         OR(RAO, RA1, RA2);
2191         op3(Stw, i, RAO);

```

```

2192         break;
2193     case IANDW:
2194         op1(Ldw, i, RA1);
2195         op2(Ldw, i, RA2);
2196         AND(RA0, RA1, RA2);
2197         op3(Stw, i, RA0);
2198         break;
2199     case IXORW:
2200         op1(Ldw, i, RA1);
2201         op2(Ldw, i, RA2);
2202         XOR(RA0, RA1, RA2);
2203         op3(Stw, i, RA0);
2204         break;
2205     case ISUBW:
2206         op1(Ldw, i, RA2);
2207         op2(Ldw, i, RA1);
2208         SUB(RA0, RA1, RA2);
2209         op3(Stw, i, RA0);
2210         break;
2211     case IADDW:
2212         op1(Ldw, i, RA1);
2213         op2(Ldw, i, RA2);
2214         ADD(RA0, RA1, RA2);
2215         op3(Stw, i, RA0);
2216         break;
2217     case ISHRW:
2218         op1(Ldw, i, RA1);
2219         op2(Ldw, i, RA2);
2220         SRL(RA0, RA2, RA1); // Shift order is switched
2221         op3(Stw, i, RA0);
2222         break;
2223     case ISHLW:
2224         op1(Ldw, i, RA1);
2225         op2(Ldw, i, RA2);
2226         SLL(RA0, RA2, RA1); // Shift order is switched
2227         op3(Stw, i, RA0);
2228     case IORB:
2229         op1(Ldb, i, RA1);
2230         op2(Ldb, i, RA2);
2231         OR(RA0, RA1, RA2);
2232         op3(Stb, i, RA0);
2233         break;
2234     case IANDB:
2235         op1(Ldb, i, RA1);
2236         op2(Ldb, i, RA2);
2237         AND(RA0, RA1, RA2);
2238         op3(Stb, i, RA0);
2239         break;
2240     case IXORB:
2241         op1(Ldb, i, RA1);
2242         op2(Ldb, i, RA2);
2243         XOR(RA0, RA1, RA2);
2244         op3(Stb, i, RA0);
2245         break;
2246     case ISUBB:
2247         op1(Ldb, i, RA1);
2248         op2(Ldb, i, RA2);
2249         SUB(RA0, RA1, RA2);
2250         op3(Stb, i, RA0);
2251         break;
2252     case IADDB:
2253         op1(Ldb, i, RA1);
2254         op2(Ldb, i, RA2);
2255         ADD(RA0, RA1, RA2);
2256         op3(Stb, i, RA0);
2257         break;
2258     case ISHRB:
2259         op1(Ldb, i, RA1);
2260         op2(Ldb, i, RA2);
2261         SRL(RA0, RA2, RA1); // Shift order is switched
2262         op3(Stb, i, RA0);
2263     case ISHLB:
2264         op1(Ldb, i, RA1);
2265         op2(Ldb, i, RA2);
2266         SLL(RA0, RA2, RA1); // Shift order is switched
2267         op3(Stb, i, RA0);

```



```

2268     case IINDC:
2269         op1(Ldw, i, RA1);           // RA1 = src1 = string
2270         NOTNIL(RA1);
2271
2272         op2(Ldw, i, RA2);           // RA2 = src2 = index
2273
2274         mem(Ldw, RA0, RA1, 0(String, len)); // RA0 = string->len
2275
2276         if(bflag){
2277             MOV(RA3, RA0);
2278             branch1 = code;
2279             BGE(RA3, R0, 0);
2280
2281             // if string->len < 0
2282             NEG(RA3, RA3);
2283             // endif
2284
2285             PATCHBRANCH(branch1);
2286             BCK(RA2, RA3);
2287         }
2288
2289         ADDI(RA1, RA1, 0(String, data));
2290
2291         branch2 = code;
2292         BGE(RA0, R0, 0);
2293
2294         // if string->len < 0
2295         SLLI(RA2, RA2, 2);           // index = index << 2; in words, not bytes
2296         // endif
2297
2298         PATCHBRANCH(branch2);
2299         mem(Ldw, RA3, RA1, RA2);     // RA3 = string[index]
2300         op3(Stw, i, RA3);
2301         break;
2302     case IINDL:
2303     case IINDF:
2304     case IINDW:
2305     case IINDB:
2306         op1(Ldw, i, RA1);           // RA1 = src1 = array
2307         NOTNIL(RA1);
2308         op3(Ldw, i, RA2);           // RA2 = src2 = index
2309
2310         if(bflag) {
2311             mem(Ldw, RA3, RA1, 0(Array, len)); // RA3 = array->len
2312             BCK(RA2, RA3);
2313         }
2314
2315         mem(Ldw, RA1, RA1, 0(Array, data)); // RA1 = array->data
2316
2317         // Modify the index to match the data width
2318         switch(i->op) {
2319             case IINDL:
2320             case IINDF:
2321                 SLLI(RA2, RA2, 3);
2322                 break;
2323             case IINDW:
2324                 SLLI(RA2, RA2, 2);
2325                 break;
2326         }
2327
2328         ADD(RA1, RA1, RA2);
2329         op2(Stw, i, RA1);
2330         break;
2331     case IINDX:
2332         op1(Ldw, i, RA1);           // RA1 = src1 = array
2333         NOTNIL(RA0);
2334         op3(Ldw, i, RA2);           // RA2 = src2 = index
2335
2336         if(bflag){
2337             mem(Ldw, RA3, RA1, 0(Array, len)); // RA3 = array->len
2338             BCK(RA2, RA3);
2339         }
2340
2341         mem(Ldw, RA3, RA1, 0(Array, t)); // RA3 = array->t
2342         mem(Ldw, RA3, RA3, 0(Type, size)); // RA3 = array->t->size
2343         mem(Ldw, RA1, RA1, 0(Array, data)); // RA1 = array->data

```

```

2344
2345     MUL(RA2, RA2, RA3);           // RA2 = index*size
2346     ADD(RA1, RA1, RA0);         // RA1 = array->data + index*size
2347     op2(Stw, i, RA1);
2348     break;
2349 case IADDL:
2350 case ISUBL:
2351 case IORL:
2352 case IANDL:
2353 case IXORL:
2354     // The Dis instructions uses the format "src3 = src2 op src1",
2355     // which is opposite to RISC-V. To make the code more intuitive the order
2356     // is switched here, so the operations are "src3 = RA1.RA2 op RA3.RA4"
2357
2358     // RA1, RA2 = src2
2359     op2(Laddr, i, RA0);
2360     mem(Ldw, RA1, RA0, 0);
2361     mem(Ldw, RA2, RA0, 4);
2362
2363     // RA3, RA4 = src1
2364     op1(Laddr, i, RA0);
2365     mem(Ldw, RA3, RA0, 0);
2366     mem(Ldw, RA4, RA0, 4);
2367
2368     switch (i->op) {
2369     case IADDL:
2370         ADD(RA0, RA1, RA3);           // RA0 = src2[31:0] + src1[31:0]
2371         ADD(RA2, RA2, RA4);         // RA2 = src2[63:32] + src1[63:32]
2372
2373         // Check for overflow
2374         SLTU(RA1, RA0, RA1);         // RA1 = RA0 < src2[31:0] ? 1 : 0
2375
2376         // Add the overflow to the upper bits
2377         ADD(RA2, RA2, RA1);
2378
2379         // Move the lower result to RA1
2380         MOV(RA1, RA0);
2381         break;
2382     case ISUBL:
2383         SUB(RA0, RA1, RA3);           // RA0 = src2[31:0] - src1[31:0]
2384         SUB(RA2, RA2, RA4);         // RA2 = src2[63:32] - src1[63:32]
2385
2386         // Check for underflow
2387         SLTU(RA1, RA1, RA0);         // RA1 = src2[31:0] < RA0 ? 1 : 0
2388
2389         // Add the underflow to the upper bits
2390         SUB(RA2, RA2, RA1);
2391
2392         // Move the lower result to RA1
2393         MOV(RA1, RA0);
2394         break;
2395     case IORL:
2396         OR(RA1, RA1, RA3);
2397         OR(RA2, RA2, RA4);
2398         break;
2399     case IANDL:
2400         AND(RA1, RA1, RA3);
2401         AND(RA2, RA2, RA4);
2402         break;
2403     case IXORL:
2404         XOR(RA1, RA1, RA3);
2405         XOR(RA2, RA2, RA4);
2406         break;
2407     }
2408
2409     // dst = RA1, RA2
2410     op3(Laddr, i, RA0);
2411     mem(Stw, RA1, RA0, 0);
2412     mem(Stw, RA2, RA0, 4);
2413     break;
2414 case ICVTWL:
2415     op1(Ldw, i, RA1);
2416     op2(Laddr, i, RA0);
2417     SRAI(RA2, RA1, 31);           // Shift right 31 places to sign-extend
2418     mem(Stw, RA1, RA0, 0);
2419     mem(Stw, RA2, RA0, 4);

```

```

2420         break;
2421     case ICVTLW:
2422         op1(Ldw, i, RAO);
2423         op3(Stw, i, RAO);
2424         break;
2425     case IBEQL:
2426         branchl(i, EQ);
2427         break;
2428     case IBNEL:
2429         branchl(i, NE);
2430         break;
2431     case IBLEL:
2432         branchl(i, LE);
2433         break;
2434     case IBGTL:
2435         branchl(i, GT);
2436         break;
2437     case IBLTL:
2438         branchl(i, LT);
2439         break;
2440     case IBGEL:
2441         branchl(i, GE);
2442         break;
2443     case ICVTFL:
2444         ADDI(Rsp, Rsp, -16);
2445
2446         op1(Ldd, i, F1);           // Load the double to convert
2447         op3(Laddr, i, Rarg);      // Load the destination as the first argument to _d2v
2448
2449         // Round F1 by adding 0.5 or -0.5
2450
2451         // F2 = 0.5
2452         LUI(Rta, SPLITH(&double05));
2453         mem(Ldd, F2, Rta, SPLITL(&double05));
2454
2455         FSGNJD(F2, F2, F1);       // F2 = F1 >= 0 ? F2 : -F2
2456         FADDD(RM, F1, F1, F2);   // F1 += F2
2457
2458         mem(Std, F1, Rsp, 8);     // Store F1 as the second argument, and call _d2v
2459
2460         // Call _d2v
2461         mem(Stw, Rfp, Rreg, 0(REG, FP));
2462         CALL(_d2v);
2463         loadi(Rreg, (ulong) &R);
2464         mem(Ldw, Rfp, Rreg, 0(REG, FP));
2465         mem(Ldw, Rmp, Rreg, 0(REG, MP));
2466
2467         ADDI(Rsp, Rsp, 16);
2468         break;
2469     case ICVTLF:
2470         op1(Laddr, i, Rta);
2471         mem(Ldw, RAO, Rta, 0);
2472         mem(Ldw, RA1, Rta, 4);
2473
2474         FCVTDWU(RM, F0, RAO);     // F0 = float(unsigned src[0:31])
2475         FCVTDW(RM, F1, RA1);     // F1 = float(src[32:63])
2476
2477         // F2 = 4294967296
2478         LUI(Rta, SPLITH(&double4294967296));
2479         mem(Ldd, F2, Rta, SPLITL(&double4294967296));
2480
2481         FMADDD(RM, F0, F1, F2, F0); // F0 = F1 * F2 + F0
2482
2483         // Store the result
2484         op3(Std, i, F0);
2485         break;
2486     case IDIVF:
2487         op1(Ldd, i, F1);
2488         op2(Ldd, i, F2);
2489         FDIVD(RM, F1, F2, F1);
2490         op3(Std, i, F1);
2491         break;
2492     case IMULF:
2493         op1(Ldd, i, F1);
2494         op2(Ldd, i, F2);
2495         FMULD(RM, F1, F2, F1);

```

```

2496         op3(Std, i, F1);
2497         break;
2498     case ISUBF:
2499         op1(Ldd, i, F1);
2500         op2(Ldd, i, F2);
2501         FSUBD(RM, F1, F2, F1);
2502         op3(Std, i, F1);
2503         break;
2504     case IADDF:
2505         op1(Ldd, i, F1);
2506         op2(Ldd, i, F2);
2507         FADDD(RM, F1, F2, F1);
2508         op3(Std, i, F1);
2509         break;
2510     case INEGF:
2511         op1(Ldd, i, F1);
2512         FSGNJND(F1, F1, F1);
2513         op3(Std, i, F1);
2514         break;
2515     case ICVTWF:
2516         op1(Ldw, i, RAO);
2517         FCVTDW(RM, F1, RAO);
2518         op3(Std, i, F1);
2519         break;
2520     case ICVTFW:
2521         op1(Ldd, i, F1);
2522         FCVTWD(RM, RAO, F1);
2523         op3(Stw, i, RAO);
2524         break;
2525     case ISHLL:
2526         /* should do better */
2527         punt(i, SRCOP|DSTOP|THREOP, optab[i->op]);
2528         break;
2529     case ISHRL:
2530         /* should do better */
2531         punt(i, SRCOP|DSTOP|THREOP, optab[i->op]);
2532         break;
2533     case IRAISE:
2534         punt(i, SRCOP|WRTPC|NEWPC, optab[i->op]);
2535         break;
2536     case IMULX:
2537     case IDIVX:
2538     case ICVTXX:
2539     case IMULX0:
2540     case IDIVX0:
2541     case ICVTXX0:
2542     case IMULX1:
2543     case IDIVX1:
2544     case ICVTXX1:
2545     case ICVTFX:
2546     case ICVTXF:
2547     case IEXPW:
2548     case IEXPL:
2549     case IEXPF:
2550         punt(i, SRCOP|DSTOP|THREOP, optab[i->op]);
2551         break;
2552     case ISELF:
2553         punt(i, DSTOP, optab[i->op]);
2554         break;
2555     }
2556 }
2557
2558 static void
2559 preamble(void)
2560 {
2561     if(comvec)
2562         return;
2563
2564     comvec = malloc(20 * sizeof(*code));
2565     if(comvec == nil)
2566         error(exNomem);
2567     code = (ulong*)comvec;
2568     codestart = code;
2569     codeend = code + 10;
2570
2571     loadi(Rh, (ulong) H);

```

```

2572     loadi(Rreg, (ulong) &R);
2573     mem(Stw, Rlink, Rreg, 0(REG, xpc));
2574     mem(Ldw, Rfp, Rreg, 0(REG, FP));
2575     mem(Ldw, Rmp, Rreg, 0(REG, MP));
2576     mem(Ldw, RAO, Rreg, 0(REG, PC));
2577     JR(RAO, 0);
2578
2579     if (cflag > 4) {
2580         iprint("preamble\n");
2581         das(codestart, code-codestart);
2582     }
2583
2584     segflush(comvec, ((ulong)code-(ulong)comvec) * sizeof(*code));
2585 }
2586
2587 int
2588 compile(Module *m, int size, Modlink *ml)
2589 {
2590     Link *l;
2591     Modl *e;
2592     int i, n;
2593     ulong *s, *tmp;
2594
2595     iprint("compile\n");
2596
2597     base = nil;
2598     patch = mallocz(size*sizeof(*patch), 0);
2599     tinit = malloc(m->ntype*sizeof(*tinit));
2600     tmp = mallocz(2048*sizeof(ulong), 0);
2601
2602     if (patch == nil || tinit == nil || tmp == nil)
2603         goto bad;
2604
2605     // Set base so that addresses are at the same order of magnitude in both passes
2606     base = tmp;
2607
2608     preamble();
2609     codestart = tmp;
2610     codeend = tmp + 2048;
2611
2612     mod = m;
2613     n = 0;
2614     pass = 0;
2615     nlit = 0;
2616
2617     // Do the first pass
2618     for (i = 0; i < size; i++) {
2619         codeoff = n;
2620         code = tmp;
2621         comp(&m->prog[i]);
2622         patch[i] = n;
2623         n += code - tmp;
2624     }
2625     iprint("first pass used %d instructions\n", n);
2626
2627     // Generate macros at the end
2628     for (i = 0; i < NMACRO; i++) {
2629         codeoff = n;
2630         code = tmp;
2631         mactab[i].gen();
2632         macro[mactab[i].idx] = n;
2633         n += code - tmp;
2634     }
2635
2636     iprint("first pass and macros used %d instructions\n", n);
2637
2638     free(tmp);
2639     base = mallocz((n+nlit)*sizeof(*code), 0);
2640     codestart = base;
2641     codeend = base + n + nlit;
2642     if (base == nil)
2643         goto bad;
2644
2645     iprint("base address: 0x%p\n", base);
2646     iprint("mod->prog: 0x%p\n", mod->prog);
2647     iprint("size: %d\n", size);

```

```

2648
2649     if (cflag > 3)
2650         iprint("dis=%5d %5d risc-v=%5d asm=%.8p: %s\n",
2651             size, size*sizeof(Inst), n, base, m->name);
2652
2653     // Prepare for the next pass
2654     pass++;
2655     nlit = 0;
2656     litpool = base + n;
2657     code = base;
2658     n = 0;
2659     codeoff = 0;
2660
2661     // Translate the instructions
2662     iprint("compile second pass\n");
2663     for (i = 0; i < size; i++) {
2664         s = code;
2665         comp(&m->prog[i]);
2666
2667         if (patch[i] != n) {
2668             // The previous instruction used a different number of instructions
2669             // than in the first pass, messing up the offsets
2670             if (cflag <= 4)
2671                 iprint("%3d %D\n", i, &m->prog[i-1]);
2672             iprint("First and second pass instruction count doesn't match\n");
2673             iprint("first pass: %lud\nsecond pass: %d\n", patch[i], n);
2674             urk("phase error");
2675         }
2676
2677         if (cflag > 4) {
2678             iprint("%3d %D\n", i, &m->prog[i]);
2679             das(s, code-s);
2680         }
2681
2682         n += code - s;
2683     }
2684
2685     // Insert the macros
2686     iprint("compile second macro\n");
2687     for (i = 0; i < NMACRO; i++) {
2688         s = code;
2689         mactab[i].gen();
2690
2691         if (macro[mactab[i].idx] != n) {
2692             iprint("mac phase err: %lud != %d\n", macro[mactab[i].idx], n);
2693             urk("phase error");
2694         }
2695
2696         n += code - s;
2697
2698         if (cflag > 5) {
2699             iprint("%s:\n", mactab[i].name);
2700             das(s, code-s);
2701         }
2702     }
2703
2704     iprint("compile m->ext types\n");
2705     for (l = m->ext; l->name; l++) {
2706         l->u.pc = (Inst*) RELPC(patch[l->u.pc - m->prog]);
2707         typecom(l->frame);
2708     }
2709
2710     if (ml != nil) {
2711         e = &ml->links[0];
2712
2713         iprint("compile ml->links types\n");
2714         for (i = 0; i < ml->nlinks; i++) {
2715             e->u.pc = (Inst*) RELPC(patch[e->u.pc - m->prog]);
2716             typecom(e->frame);
2717             e++;
2718         }
2719     }
2720
2721     iprint("compile m->type types\n");
2722     for (i = 0; i < m->ntype; i++) {
2723         if (tinit[i] != 0)

```

```
2724         typecom(m->type[i]);
2725     }
2726
2727     iprint("compile patches\n");
2728     patchex(m, patch);
2729     m->entry = (Inst*) RELPC(patch[mod->entry - mod->prog]);
2730
2731     iprint("compile done\n");
2732     free(patch);
2733     free(tinit);
2734     free(m->prog);
2735     m->prog = (Inst*) base;
2736     m->compiled = 1;
2737     segflush(base, n*sizeof(*base));
2738     return 1;
2739
2740     bad:
2741     iprint("compile failed\n");
2742     free(patch);
2743     free(tinit);
2744     free(base);
2745     free(tmp);
2746     return 0;
2747 }
```

