

ForwardCom: An open-standard instruction set for high-performance microprocessors

Agner Fog

November 3, 2017

Contents

1	Introduction	2
1.1	Highlights	2
1.2	Background	3
1.3	Design goals	3
1.4	Comparison with other open instruction sets	4
1.5	References and links	5
2	Basic architecture	6
2.1	A fully orthogonal instruction set	6
2.2	Instruction size	6
2.3	Register set	7
2.4	Vector support	8
2.5	Vector loops	8
2.6	Maximum vector length	9
2.7	Instruction masks	10
2.8	Addressing modes	10
3	Instruction formats	12
3.1	Formats and templates	12
	Maximum number of input operands	19
3.2	Coding of operands	19
	Operand type	19
	Register type	19
	Pointer register	19
	Index register	19
	Offsets	20
	Limit on index	20
	Vector length	20
	Combining vectors with different lengths	20
	Immediate constants	20
	Mask register and fallback register	21
3.3	Coding of masks	21
3.4	Format for jump, call and branch instructions	23
3.5	Assignment of opcodes	26
4	Instruction lists	28
4.1	List of multi-format instructions	33
4.2	List of tiny instructions	34
4.3	List of single-format instructions	35
4.4	List of control transfer instructions	43

5	Description of instructions	44
	Data move and conversion instructions	44
	Data read and write instructions	52
	General arithmetic instructions	55
	Arithmetic instructions with carry, overflow check, or saturation	65
	Logic and bit manipulation instructions	68
	Combined ALU and branch instructions	76
	Unconditional and indirect jump, call, and return instructions	82
	Miscellaneous instructions	84
	System instructions	85
5.1	Common operations that have no dedicated instruction	90
5.2	Unused instructions	91
6	Other implementation details	92
6.1	Endianness	92
6.2	Implementation of call stack	92
6.3	Floating point errors and exceptions	94
6.4	Detecting integer overflow	95
6.5	Multithreading	96
6.6	Security features	96
	How to improve the security of applications and systems	96
7	Programmable application-specific instructions	99
8	Microarchitecture and pipeline design	100
8.1	Proposal for reducing branch misprediction delay	102
9	Memory model	104
9.1	Thread memory protection	105
9.2	Memory management	105
10	System programming	109
10.1	Memory map	109
10.2	Call stack	110
10.3	System calls and system functions	110
10.4	Inter-process calls	112
10.5	Error message handling	112
11	Support for multiple instruction sets	114
11.1	Transitions between ForwardCom and x86-64	114
11.2	Transitions between ForwardCom and ARM	115
11.3	Transitions between ForwardCom and RISC-V	115
12	Standardization of ABI and software ecosystem	117
12.1	Compiler support	117
12.2	Binary data representation	118
12.3	Further conventions for object-oriented languages	119
12.4	Function calling convention	119
12.5	Register usage convention	121
12.6	Name mangling for function overloading	122
12.7	Binary format for object files and executable files	122
12.8	Function libraries and link methods	123
12.9	Predicting the stack size	124
12.10	Exception handling, stack unrolling and debug information	125
12.11	Assembly language syntax	126

13 Binary tools	127
13.1 Assembler	127
Introduction	127
Command line	129
File format	129
Names of symbols	130
Sections	130
Constant expressions	131
Data types	131
Data definitions	132
Function definitions and labels	132
Instructions	133
Unconditional and indirect jumps, calls, and returns	134
Conditional jumps and loops	135
Imports and exports	137
Other directives	137
13.2 Metaprogramming	138
Metaprogramming variables	138
13.3 Disassembler	139
13.4 Linker	139
13.5 Library manager	139
13.6 Emulator	140
13.7 Dump utility	140
13.8 Compiling the forw tools	140
14 Code examples	141
Horizontal vector add	141
Horizontal vector minimum	141
Switch-case statement	142
Boolean operations	143
Virtual functions	144
Memory copying	146
String length	147
High precision arithmetic	148
Matrix multiplication	148
14.1 Optimization tricks	149
15 Conclusion	151
16 Revision history	154
17 Copyright notice	156

Chapter 1

Introduction

ForwardCom stands for Forward Compatible Computer system.

This document proposes a new open instruction set architecture designed for optimal performance, flexibility and scalability. The ForwardCom project includes both a new instruction set architecture and the corresponding ecosystem of software standards, application binary interface (ABI), memory management, development tools, library formats and system functions. This project illustrates the improvements that can be obtained by a complete vertical redesign of hardware and software based on an open, collaborative process.

A short introduction to ForwardCom is provided at <http://www.forwardcom.info>.

This manual and all associated code is maintained at <https://github.com/ForwardCom>.

1.1 Highlights

- The ForwardCom instruction set is neither RISC nor CISC, but a new paradigm with the advantages of both. ForwardCom has few instructions, but many variants of each instruction. A consistent template system with few instruction sizes combines the fast and streamlined decoding and pipeline design of RISC systems with the compactness and more work-done-per-instruction of CISC systems.
- The ForwardCom design is scalable to support small embedded systems as well as large supercomputers and vector processors without losing binary compatibility.
- Vector registers of variable length are provided for efficient handling of large data sets.
- Array loops are implemented in a new flexible way that automatically uses the maximum vector length supported by the microprocessor in all but the last iteration of a loop. The last iteration automatically uses a vector length that fits the remaining number of elements. No extra code is needed to deal with remaining data and special cases. There is no need to compile the code separately for different microprocessors with different vector lengths.
- No recompilation or update of software is needed when a new microprocessor with a different vector register length becomes available. The software is guaranteed to be forward compatible and take advantage of the longer vectors of new microprocessor models.
- Strong security features are a fundamental part of the hardware and software design.
- Memory management is simpler and more efficient than in traditional systems. Various techniques are used for avoiding memory fragmentation. There is no memory paging and no translation lookaside buffer (TLB). Instead, there is a memory map with a limited number of sections with variable size.

- There are no dynamic link libraries (DLLs) or shared objects. Instead, there is only one type of function libraries that can be used for both static and dynamic linking. Only the part of the library that is actually used is loaded and linked. The library code is kept contiguous with the main program code to improve caching and reduce memory fragmentation. Executable files can be re-linked to replace or update library functions and plug-ins and to support multiple user interface frameworks.
- A mechanism for calculating the required stack size is provided. This can prevent stack overflow in most cases without making the stack bigger than necessary.
- A mechanism for optimal register allocation across program modules and function libraries is provided. This makes it possible to keep most variables in registers without spilling to memory. Vector registers can be saved in an efficient way that stores only the part of the register that is actually used.

This can be useful as a sandbox for experiments aiming at improving many different aspects of computer design, as discussed at <http://www.forwardcom.info>.

1.2 Background

An instruction set architecture is a standardized set of machine instructions that a computer can run. There are many instruction set architectures in use.

Some commonly used instruction sets are poorly designed from the beginning. These systems have been augmented many times with extensions and patches. One of the worst cases is the widely used x86 instruction set and its many extensions. The x86 instruction set is the result of a long history of short-sighted extensions and patches. The result of this development history is a very complicated architecture with thousands of different instruction codes, which is very difficult and costly to decode in a microprocessor. We need to learn from past mistakes in order to make better choices when designing a new instruction set architecture and the software that supports it.

The design should be based on an open process. Krste Asanović and David Patterson have presented compelling arguments for why an open instruction set should be preferred. Openness can be crucial for the success of a technical design. For example, the original IBM PC in the early 1980's had an advantage over competing computers because the open architecture allowed other hardware and software producers to make compatible equipment. IBM lost their market dominance when they switched to the proprietary Micro Channel Architecture in 1987. The successes of open source software are well known and need no further discussion here. The only thing that is missing for a complete computer ecosystem based on open standards is an open microprocessor architecture. This will open the market also for smaller microprocessor producers and niche products.

This project is based on discussions in various Internet forums. The specifications are preliminary. The development of a new standard should benefit from a long experimental phase, and it would be unwise to make it a fixed standard at this initial stage.

1.3 Design goals

Previously published open instruction sets are suitable for small, cheap microprocessors for embedded systems, system-on-a-chip designs, FPGA implementations for scientific experiments, etc. The proposed ForwardCom architecture takes the idea further and aims at a design that can outperform existing high-end processors.

The ForwardCom instruction set architecture is based on the following priorities:

- The instruction set should have a simple and consistent modular design.

- The instruction set should represent a suitable compromise between the RISC principle that enables fast decoding, and the CISC principle that makes it possible to do more work per instruction and to use the code cache more efficiently.
- The design should be extensible so that new instructions and extensions can be added in a consistent and predictable way.
- The design should be scalable so that it is suitable for both small computers with on-chip RAM and large supercomputers with very long vectors.
- The design should be competitive over current commercial designs with a focus on the high-end applications of tomorrow rather than the low-end applications of yesterday.
- Vector support and other features that have proven essential for high performance should be a fundamental part of the design, not a clumsy appendix.
- Security should be a fundamental part of the design, not patches added ad hoc.
- The instruction set should be designed through an open process with the participation of the international hardware and software community, similar to the standardization work in other technical areas.
- The entire vertical design should be non-proprietary and allow anybody to make compatible software, hardware, and equipment for test, debugging and emulation.
- Decisions about instructions and extensions should not be determined by the short term marketing considerations of an oligopolistic microprocessor industry but by the long term needs of the entire hardware and software community and organizations.
- The design should allow the construction of forward compatible software that will run optimally without recompilation on future processors with larger vector registers.
- The design should allow application-specific extensions.
- The basic aspects of the entire ecosystem of ABI standard, assembler, compilers, function libraries, system functions, user interface framework, etc. should also be standardized for maximum compatibility.

A new instruction set will not easily get success on a commercial market, even if it is better than legacy systems, because the market prefers backward compatibility with existing software and hardware. It is unlikely that the ForwardCom instruction set will make a successful commercial product within a short time frame, but the discussion about what an ideal instruction set, microprocessor design, and software ecosystem might look like is still useful. The ForwardCom project has already generated so many important new ideas that it is worth pursuing further, even if we don't know where this process will end. The present work can be useful if the need for introducing a new instruction set architecture should arise for other reasons. It will be particularly useful for large vector processors, for applications where security is important, for real-time operating systems, as well as for projects where the patent and license restrictions of other architectures would be an obstacle.

The proposals in this document will also be useful as a source of inspiration and for scientific experiments. Many of the ideas are independent of the design details and may be implemented in existing systems.

1.4 Comparison with other open instruction sets

A few other open instruction sets have been proposed, most notably RISC-V and OpenRISC. Both are pure RISC designs with mostly fixed 32-bit instruction word sizes. These instruction sets are suitable

for small systems where the use of silicon space is economized, but they are not designed for high performance superscalar processors and they do not focus on details that are critical for achieving maximum performance in bigger systems. The present proposal is thought as the next step towards making an open instruction set that is actually more efficient than the best commercial instruction sets today.

A typical RISC design with the instruction size limited to 32 bits leaves only limited space for immediate constants and addresses of memory operands. A medium size program will need 32-bit relative addresses of static memory operands to avoid overflow during the relocation process in the linker. A 32-bit relative address requires several instructions in the pure RISC designs. For example, to add a memory operand to the value of a register, you need five instructions in a RISC design with only 32-bit instruction words: (1) load the lower part of the 32-bit address offset, (2) add the upper part of the 32-bit address offset, (3) add the reference pointer or instruction pointer to this value, (4) read the memory operand from the calculated address, (5) do the desired addition. The ForwardCom design does all this in a single instruction with double word size. The speed advantage is obvious. The address calculation, load, and execution are done at each their stage in the pipeline in order to achieve a smooth throughput of one instruction per clock cycle in each pipeline lane.

Another important difference is that the previous RISC designs have limited support for vector operations. The ForwardCom design introduces a new system of variable-length vector registers that is more efficient and flexible than the best current commercial designs. Efficient vector operations are essential for obtaining maximum performance, and this has been an important priority in the design of the ForwardCom architecture proposed here.

1.5 References and links

- Krste Asanović and David Patterson: "The Case for Open Instruction Sets. Open ISA Would Enable Free Competition in Processor Design". Microprocessor Report, August 18, 2014. www.linleygroup.com/mpr/article.php?id=11267
- RISC-V: The Free and Open RISC Instruction Set Architecture riscv.org
- OpenRISC: openrisc.io
- Open Cores: opencores.org
- Agner Fog: Proposal for an ideal extensible instruction set, 2015. A blog discussion thread that initiated the ForwardCom project. www.agner.org/optimize/blog/read.php?i=421
- Agner Fog: Stop the instruction set war, 2009. Blog post about the problems with the x86 instruction set. www.agner.org/optimize/blog/read.php?i=25
- Darek Mihocka: Standard Need To Be Forward Looking, 2007. Blog post criticizing the x86 instruction set standard. www.emulators.com/docs/nx02_standards.htm. See also the following pages.

Chapter 2

Basic architecture

This chapter gives an overview of the most important features of the ForwardCom instruction set architecture. Details are given in the subsequent chapters.

2.1 A fully orthogonal instruction set

The ForwardCom instruction set is fully orthogonal in all respects. Where other instruction sets have a large number of different instructions for different register types, operand types, operand sizes, addressing modes, etc., ForwardCom has fewer instructions, but many variants of each instruction. This modular design makes the hardware implementation much simpler. The same instruction can use integer operands of all sizes and floating point operands of all precisions. It can use register operands, memory operands or immediate operands. It can use many different addressing modes. Instructions can be coded in short forms with two operands where the same register is used for destination and source operand, or longer forms with three operands. It can work with scalars or vectors of any size. It can have predication or masks for conditional execution at the vector element level, and it can have optional flag inputs for deciding rounding mode, exception control and other details, where appropriate. Data constants of all types can be included in the instructions and compressed in various ways to reduce the instruction size.

Rationale

The orthogonality is implemented by a standardized modular design that makes the hardware implementation simpler. It also makes compilation simpler and more flexible and makes it easier for the compiler to convert linear code to vector code.

The support for immediate constants of all types is an improvement over current systems. Most current systems store floating point constants in a data segment and access them through a 32-bit address in the instruction code. This is a waste of data cache space and causes many cache misses because the data are scattered around in different sections. Replacing a 32-bit address with a 32-bit immediate constant makes the code more efficient without increasing the code size. Extensions to allow 64-bit immediate constants are possible at the cost of having instructions with triple length.

2.2 Instruction size

The ForwardCom instruction set uses a 32-bit word size for code. An instruction can consist of one, two, or optionally three 32-bit words. The code density can be increased by using tiny instructions of half the size, but the 32-bit unit size is preserved by pairing tiny instructions two-by-two. It is not possible to jump to the second tiny instruction in such a pair of tiny instructions. It is possible to add future extensions with instruction sizes of four or more words.

Rationale

A CISC architecture with many different instruction sizes is inefficient in superscalar processors where we want to execute several instructions per clock cycle. The decoding front end is often a bottleneck. You have to determine the length of the first instruction before you know where the next instruction begins. The “instruction length decoding” is a fundamentally serial process which makes it difficult to decode multiple instructions per clock cycle. Some microprocessors have an extra “micro-operations cache” after the decoder in order to circumvent this bottleneck.

Here, it is desired to have as few different instruction lengths as possible and to make it easy to determine the length of each instruction. We want a small instruction size for the most common simple instructions, but we also need a larger instruction size in order to accommodate things like a larger register set, instructions with multiple operands, vector operations with advanced features, 32-bit address offsets, and large immediate constants. This proposal is a compromise between code compactness, easy decoding, and space for advanced features. The instruction length is indicated by only two bits. A decoder can find the instruction boundaries in n words by means of a simple Boolean function of $2n$ inputs.

2.3 Register set

There are 32 general purpose registers ($r0$ – $r31$) of 64 bits each, and 32 vector registers ($v0$ – $v31$) of variable length. The maximum vector length is different for different hardware implementations. The general purpose registers can be used for integers of up to 64 bits as well as for pointers. The vector registers can be used for scalars or vectors of integers and floating point numbers.

The following special registers are defined and visible at the application program level. All have 64 bits:

- Instruction pointer (IP)
- Data section pointer (DATAP)
- Thread data pointer (THREADP)
- Stack pointer (SP)
- Numeric control register (NUMCONTR)

The stack pointer is identical to $r31$. The other special registers cannot be accessed as ordinary registers.

There is no dedicated flags register. Registers $r0$ – $r6$ and $v0$ – $v6$ can be used for masks, predicates and floating point option flags to control attributes such as rounding mode and exception control.

The unused part of a register is always set to zero. This means that integer operations with an operand size smaller than 64 bits and vector operations with a vector length smaller than the maximum will always set the unused bits of the destination register to zero.

Rationale

The number of registers is a compromise between code density and flexibility. The cost of spilling registers to memory is usually important only in the critical innermost loop, which is unlikely to need more than 32 registers.

We can avoid false dependencies on the previous value of a register by setting all unused register bits to zero rather than leaving them unchanged. The hardware can save power by disabling the unused parts of execution units and data buses.

A dedicated flags register is unfeasible for code that schedules multiple calculations in between each other and for vector code.

The reason for handling floating point scalars in the vector registers rather than in separate registers is to make it easy for a compiler to convert scalar code including function calls to vector code. Floating point code often contains calls to mathematical library functions. A library function with variable-length vectors as input and output can be used for both scalars and vectors, and the compiler can easily vectorize code that contains such library function calls.

2.4 Vector support

A vector register can contain signed or unsigned integers of 8, 16, 32, 64, and optionally 128 bits, or floating point numbers of single and double precision. There is limited support for floating point numbers in half precision and optional support for quadruple precision. All elements of a vector must have the same type. The elements of a vector are processed in parallel. For example, a vector addition will produce the sum of two vectors in a single operation.

The vector registers have variable length. Each vector register has extra bits for storing the length of the vector. The maximum vector length depends on the hardware. For example, if the hardware supports a maximum vector length of 64 bytes and a particular application needs only 16 bytes, then the vector length is set to 16.

Some instructions need to specify the length of a vector explicitly, for example when reading a vector from memory. These instructions use a general purpose register for specifying the vector length. The length is usually indicated as the number of bytes, not the number of vector elements.

The maximum length supported by the processor must be a power of 2. The actual length specified does not need to be a power of 2. If the specified length is longer than the maximum length, then the maximum length is used.

The contents of a vector register can arbitrarily be interpreted as any of the types and element sizes supported. For example, the hardware does not prevent the application of integer instructions on a vector that contains floating point data. It is the responsibility of the programmer that the code makes sense.

2.5 Vector loops

A special addressing mode is provided to make vector loops more compact and efficient. It uses a pointer P to the end of an array, and a negative index J, and calculates the address of a memory operand as P-J, where P and J are general purpose registers. This makes it possible to make a loop through an array as illustrated by the following pseudocode:

```
P = address of array
J = size of array (in bytes)
L = maximum vector length (depends on processor)
X = a vector register
P += J; // point to end of array
while (J > 0) {
    X = whatever_operation(X, [P-J], vector_length = J)
    J -= L;
}
```

This loop works in the following way: P points to the end of the array. J is the remaining number of array bytes; counting down until the loop is finished. The loop reads one vector at a time from the array at the address (P-J). J is larger than the maximum vector length L in all but the last iteration of the loop. This makes the processor use the maximum vector length. If the array size is not divisible by the maximum vector length then the last iteration of the loop will use a smaller vector length that

fits the remaining number of elements. Obviously, the loop can contain any number of vector read, vector write, and vector arithmetic instructions, using the same principle.

This loop will work on different processors with different maximum vector lengths *without knowing the maximum vector length at compile time*. Thus, the same piece of software will work on different microprocessors with different vector lengths without the need to compile separately for each microprocessor.

A further advantage is that no extra code is needed after the loop to handle remaining elements in the case that the array size is not divisible by the vector length. The loop overhead can be reduced to a single instruction (`sub_maxlen/jump_pos`) which subtracts L from J and jumps back if the result is positive.

Rationale

Most current systems have fixed vector lengths. If different processors have different vector lengths then you have to compile the code separately for each vector length. Every time a new processor with longer vectors comes on the market, you have to compile a new version of the code for the new vector length, using newly defined extensions to the instruction set. It usually takes several years for the new software to be developed and to penetrate the mainstream market. It is so costly for software producers to develop, test, and maintain different versions of their code for each vector length that this is rarely done.

A further problem with current systems is that it is impossible to save a vector register in a way that is guaranteed to be compatible with future processors with longer vectors. This is no problem with the ForwardCom design because the vector length is stored in the vector register itself. Instructions are provided for saving and restoring vectors of variable length and for storing only the part of a vector register that is actually used.

The ForwardCom design makes it possible to take advantage of a new processor with longer vector registers immediately without recompiling the code. The loop method described above makes this easy and very efficient. You don't need different versions of the code for different processors.

It is possible to obtain the same effect without the special negative addressing mode by inverting the sign of J and allowing a negative value in the register that specifies the vector length while using the absolute value for the actual vector length. This solution is less elegant and more confusing, but it may possibly be included in other instruction sets by allowing negative values when specifying a vector length.

Loop unrolling is generally not necessary. The loop overhead is already reduced to a single instruction and a superscalar processor will execute multiple iterations in parallel if dependency chains are not too long. Loop unrolling with multiple accumulators may be useful for hiding a loop-carried dependency. In this case, you will either insert a loop control instruction after each section in the unrolled code or calculate the loop iteration count before the loop.

The ForwardCom design has no practical limit to the vector length that a microprocessor can support. A large microprocessor with very long vectors can be useful for calculations with a high amount of data parallelism. Other solutions to obtain high performance on parallel data processing have been discussed, such as rolling register stacks and software pipelining, but it was concluded that long vectors is the method that can be implemented most efficiently in the microprocessor as well as in the compiler.

2.6 Maximum vector length

The maximum length of vector registers will be different for different processors. The maximum length must be a power of 2. It can be as large as desired and must be at least 16 bytes. Each instruction can use a smaller length, which does not need to be a power of 2.

The maximum length may be different for different element sizes. For example, the maximum length for 32-bit integers can be 32 bytes to contain eight integers, while the maximum length for 8-bit integers could be 16 bytes to contain 16 smaller numbers. However, the maximum length must be the same for different types with the same element size. For example, the maximum length for double precision floating point numbers must be the same as for 64-bit integers because loops are likely to contain both types when integer vectors are used as masks for floating point vectors. The maximum length for a 32-bit element size cannot be less than for any other element size or operand type. This rule guarantees that it is possible to save a complete vector using a 32-bit operand type.

The maximum vector length should generally be the same for all instructions for the same data type, but there may be exceptions for instructions that are particularly expensive to implement.

It is possible for an application program or the operating system to reduce the maximum vector length. This can be useful if a smaller vector length is more appropriate for a particular purpose.

It is also possible to increase the apparent maximum vector length for purposes of emulation. Virtual vector registers that are bigger than what the hardware supports can be emulated through traps (synchronous interrupts) in order to verify the functionality of a program on processors with a longer maximum vector length than is currently available.

When an instruction specifies a longer vector than the maximum, then the maximum length is used (unless the emulation of larger vectors is activated). This is necessary for the efficient implementation of vector loops as described above on page 8. If the specified vector length is zero or negative then the result will be a vector of zero length.

2.7 Instruction masks

Most instructions can have a mask register which can be used for conditional execution and for specifying various options. Instructions with general purpose registers use one of the registers r0–r6 as a mask register or predicate. Bit 0 of the mask register indicates whether the operation is executed or not.

The instruction will produce the normal result when bit 0 of the mask is one, and a fallback value when this bit is zero. The fallback value can be the value of the first source operand, a separate register, or zero.

This mechanism can be vectorized. Instructions with vector registers use one of the vector registers v0–v6 as mask register. The calculation of each vector element is conditional on the corresponding element in the mask register.

Additional bits in the mask register are used for various options, overriding the values in the numeric control register. See page 22 for details.

2.8 Addressing modes

All memory addressing is relative to a base pointer. Memory operands are addressed in this general form:

$$\text{Address} = \text{BP} + \text{IX} * \text{SF} + \text{OF}$$

Where BP is a 64-bit base pointer, IX is a 64-bit index register, SF is a scale factor, and OF is a direct offset. A base pointer is always present; the other elements are optional.

The base pointer, BP, can be a general purpose register, or it can be the instruction pointer (IP), data section pointer (DATAP), thread data pointer (THREADP), or stack pointer (SP).

The index register, IX, can be one of the registers r0–r30. A value of 31 in the index register field means no index register.

A limit can be applied to the index register in the form of an integer constant. A trap is generated if the index register is bigger than the limit in an unsigned comparison.

The scale factor, SF, is equal to the operand size (in bytes) for scalar operands and broadcasts. The scale factor is 1 for vector operands. A special addressing mode with $SF = -1$ is also available, as explained on page 8.

The offset, OF, is a sign-extended integer of 8, 16 or 32 bits. 8-bit offsets are multiplied by the operand size. Offsets of 16 and 32 bits have no multiplier.

Support for addressing modes with both index and offset is optional.

Memory operands in vector instructions can load a vector of a specified length, a scalar, or a broadcast scalar. The length of the loaded or broadcast vector is specified by a general purpose register. The specified length is the number of bytes. The number of vector elements is the number of bytes divided by the operand size. Register r31, which is the stack pointer, cannot be used for specifying vector length. Instead, a value of 31 in the length register fields will give a scalar.

Jumps and calls specify a target address relative to the instruction pointer. The relative address is specified with a signed offset of 8, 16, 24, or 32 bits, multiplied by the code word size which is 4. This will cover an address range of ± 8 gigabytes with the 32-bit offset.

Rationale

A 64-bit address space is used. Relative addressing is used in order to avoid 64-bit addresses in the instruction code. In the rare case that a 64-bit absolute address is needed, it must be loaded into a register which is then used as a pointer.

Addressing with an index scaled by the operand size is useful for arrays. A limit can be applied to the index so that array bounds can be checked without any extra instructions.

Addressing with a negative index is useful for the efficient implementation of vector loops described on page 8.

The addressing modes specified here will cover all common applications, including arrays, vectors, structures, classes, and stack frames.

Support for addressing modes with both base pointer, index and direct offset is optional because this requires two adders in the address-calculation stage in the pipeline which might limit the maximum clock frequency.

Chapter 3

Instruction formats

3.1 Formats and templates

All instructions use one of the general format templates shown below (the most significant bits are shown to the left). The basic layout of the 32-bit code word is shown in template A. Template B, C and D are derived from template A by replacing 8, 16 or 24 bits, respectively, with immediate data. Double-size and triple-size instructions can be constructed by adding one or two 32-bit words to one of these templates. For example, template A with an extra 32-bit word containing data is called A2. Template E2 is an extension to template A where the second code word contains an extra register field, extra opcode bits, mode bits, option bits, and data.

Some small, often-used instructions can be coded in a tiny format that uses a half code word. Two such tiny instructions can be packed into a single code word, using template T. An unpaired tiny instruction must be combined with a tiny NOP to fill a whole code word.

Bits	2	3	6	5	1	2	5	3	5
Field	IL	Mode	OP1	RD	M	OT	RS	Mask	RT
Template A. Has three operand registers and a mask register.									

Bits	2	3	6	5	1	2	5	8
Field	IL	Mode	OP1	RD	M	OT	RS	IM1
Template B. Has two operand registers and an 8-bit immediate constant.								

Bits	2	3	6	5	8	8
Field	IL	Mode	OP1	RD	IM2	IM1
Template C. Has one operand register two 8-bit immediate constants.						

Bits	2	3	3	24
Field	IL	Mode	OP1	IM2
Template D. Has no register and a 24-bit immediate constant.				

Bits	2	3	6	5	1	2	5	3	5
Field	IL	Mode	OP1	RD	M	OT	RS	Mask	RT
Field	OP2	Mode2	IM3	RU	IM2				
Template E2. Has 4 register operands, mask, a 16-bit immediate constant and extra bits for opcode, mode, and options.									

Bits	2	3	6	5	1	2	5	3	5
Field	IL	Mode	OP1	RD	M	OT	RS	Mask	RT
Field	IM2								
Template A2. 2 words. As A, with an extra 32-bit immediate constant.									

Bits	2	3	6	5	1	2	5	3	5
Field	IL	Mode	OP1	RD	M	OT	RS	Mask	RT
Field	IM2								
Field	IM3								
Template A3. 3 words. As A, with two extra 32-bit immediate constants.									

Bits	2	3	6	5	1	2	5	8
Field	IL	Mode	OP1	RD	M	OT	RS	IM1
Field	IM2							
Template B2. As B, with an extra 32-bit immediate constant.								

Bits	2	3	6	5	1	2	5	8
Field	IL	Mode	OP1	RD	M	OT	RS	IM1
Field	IM2							
Field	IM3							
Template B3. As B, with two extra 32-bit immediate constants.								

Bits	2	3	6	5	8	8
Field	IL	Mode	OP1	RD	IM2	IM1
Field	IM3					
Template C2. As C, with an extra 32-bit immediate constant.						

Bits	2	3	6	5	1	2	5	3	5
Field	IL	Mode	OP1	RD	M	OT	RS	Mask	RT
Field	OP2	Mode2	IM3	RU	IM2				
Field	IM4								
Template E3. Has 4 register operands, mask, 16-bit and 32-bit immediate constants and extra bits for opcode, mode, and options.									

Bits	4	14	14
Field	0111	Tiny instruction 2	Tiny instruction 1
Template T. 1 word containing two tiny instructions.			

Bits	5	5	4
Field	OP1	RD	RS
Format for each tiny instruction			

The meaning of each field is described in the following table.

Table 3.14: Fields in instruction templates

Field name	Meaning	Values
IL	Instruction length	0 or 1: 1 word = 32 bits 2: 2 words = 64 bits 3: 3 or more words
Mode	Format	Determines the format template and the use of each field. Extended with the M bit when needed. See details below.
Mode2	Format	Extension to Mode.
OT	Operand type and size (OS)	0: 8 bit integer, OS = 1 byte 1: 16 bit integer, OS = 2 bytes 2: 32 bit integer, OS = 4 bytes 3: 64 bit integer, OS = 8 bytes 4: 128 bit integer, OS = 16 bytes (optional) 5: single precision float, OS = 4 bytes 6: double precision float, OS = 8 bytes 7: quadruple precision float, OS = 16 bytes (optional) The OT field is extended with the M bit when needed.
M	Operand type or mode	Extends the mode field when bit 1 and bit 2 of Mode are both zero (general purpose registers). Extends the OT field otherwise (vector registers).
OP1	Opcode	Decides the operation, for example add or move.
OP2	Opcode	Opcode extension for single-format instructions.
RD	Destination register	r0 – r31 or v0 – v31. Also used for first source operand and fallback if the instruction format does not specify enough operands.
RS	Source register	r0 – r31 or v0 – v31. Source register, pointer, index, fallback, or vector length register.
RT	Source register	r0 – r31 or v0 – v31. Source register or pointer.
RU	Source register	r0 – r31 or v0 – v31. Source register or fallback.
Mask	mask register	0-6 means that a general purpose register or vector register is used for mask and option bits. 7 means no mask.
IM1 IM2 IM3 IM4	Immediate data	8, 16, 24, or 32 bits immediate operand or address offset or option bits. Adjacent IM fields can be merged to make a larger constant.

Instructions have several different formats, defined by the IL and mode bits, according to table 3.15 below. The different formats specify different sizes of immediate data or memory operands with different addressing modes.

Instructions can have up to three source operands (input), one destination operand (output), and a mask. The destination operand always uses the RD field, except where the destination is a memory operand. The source operands are using the available operand fields according to the following algorithm: The required source operands are assigned to the available operand fields defined by table 3.15 in the following order of priority: immediate data field, memory operand, RT, RS, RU, RD - assigning the last source operands first. Thus, RD is used for both destination and the first source operand only if there are no other vacant register fields. Any remaining register field is used as fallback value if there is a mask. Otherwise, the fallback value will be the first source operand.

This principle can be illustrated with the following example. Format 2.2.1 uses RT for pointer in a memory operand and RS for vector length, leaving only RU and RD for register operands. An instruction with one input operand will have the memory operand as input and RU as fallback. An instruction with two input operands will have RU and memory as the operands and RU as fallback. An instruction with three input operands will have RD, RU, and memory as the operands and RD as fallback.

The coding of instructions with two or three source operands is indicated in the table in the following way:

$RD = f2(RS, RT)$ means that instructions with two input operands ($f2$) use the register specified in RD as destination operand and RS and RT as source operands.

$RD = f3(RD, RU, [RT + RS * OS + IM2])$ means that instructions with three input operands ($f3$) use the register specified in RD as both destination and the first source operand. The second source operand is RU. The third source operand is a memory operand with RT as base pointer, RS as index scaled by the operand size, and the constant IM2 as offset.

Instructions with only one input operand are coded as $f2$ with the first source operand omitted.

Table 3.15: List of instruction formats

Format name	IL	Mode. Mode2	Template	Use
0.0	0	0	A	Three general purpose register operands. $RD = f2(RS, RT)$. $RD = f3(RD, RS, RT)$.
0.1	0	1	B	Two general purpose registers and 8-bit immediate operand. $RD = f2(RS, IM1)$. $RD = f3(RD, RS, IM1)$.
0.2	0	2	A	Three vector register operands. $RD = f2(RS, RT)$. $RD = f3(RD, RS, RT)$.
0.3	0	3	B	Two vector registers and a broadcast 8-bit immediate operand. $RD = f2(RS, IM1)$. $RD = f3(RD, RS, IM1)$.
0.4	0	4	A	One vector register and memory operand. Vector length specified by general purpose register. $RD = f2(RD, [RT])$. $length = RS$.
0.5	0	5	A	One vector register and a memory operand with base pointer and negative index. This is used for vector loops as explained on page 8. $RD = f2(RD, [RT - RS])$. $length = RS$.
0.6	0	6	A	One vector register and a scalar memory operand with base pointer and scaled index. $RD = f2(RD, [RT + RS * OS])$.
0.7	0	7	B	One vector register and a scalar memory operand with base pointer and 8-bit offset. $RD = f2(RD, [RS + IM1 * OS])$.
0.8	0	M=1	A	One general purpose register and a memory operand with base pointer and scaled index. $RD = f2(RD, [RT + RS * OS])$.
0.9	0	M=1	B	One general purpose register and a memory operand with base pointer and 8-bit offset. $RD = f2(RD, [RS + IM1 * OS])$.
1.0	1	0	A	Single-format instructions. Three general purpose register operands. $RD = f2(RS, RT)$. $RD = f3(RD, RS, RT)$.

1.1	1	1	C	Single-format instructions. One general purpose register and a 16-bit immediate operand. $RD = f2(RD, IM1-2)$.
1.2	1	2	A	Single-format instructions. Three vector register operands. $RD = f2(RS, RT)$. $RD = f3(RD, RS, RT)$.
1.3B	1	3	B	Single-format instructions. Two vector registers and a broadcast 8-bit immediate operand. $RD = f2(RS, IM1)$. $RD = f3(RD, RS, IM1)$.
1.3C	1	3	C	Single-format instructions. One vector register and a broadcast 16-bit immediate operand. $RD = f2(RD, IM1-2)$.
1.4	1	4	B	Jump instructions with two register operands and 8 bit offset.
1.5C	1	5	C	Jump instructions with one register operand, 8 bit constant (IM2) and 8 bit offset (IM1).
1.5D	1	5	D	Jump instructions with no register and 24 bit offset.
T	1	6-7	T	Two tiny instructions.
1.8	1	0 M=1	B	Single-format instructions. Two general purpose registers and an 8-bit immediate operand. $RD = f2(RS, IM1)$. $RD = f3(RD, RS, IM1)$.
1.9				There is no format 1.9 because 1.1 has no M bit.
2.0.0	2	0.0	E2	Three general purpose registers and a memory operand with base and 16 bit offset. $RD = f2(RS, [RT+IM2])$. $RD = f3(RU, RS, [RT+IM2])$.
2.0.1	2	0.1	E2	Two general purpose registers and a memory operand with base and index, no scale. Optional support for a 16 bit offset. $RD = f2(RU, [RT+RS+IM2])$. $RD = f3(RD, RU, [RT+RS+IM2])$. IM2 must be zero if offset not supported.
2.0.2	2	0.2	E2	Two general purpose registers and a memory operand with base and scaled index. Optional support for a 16 bit offset. $RD = f2(RU, [RT+RS*OS+IM2])$. $RD = f3(RD, RU, [RT+RS*OS+IM2])$. IM2 must be zero if offset not supported.
2.0.3	2	0.3	E2	Two general purpose registers and a memory operand with base, scaled index, and 16-bit limit. $RD = f2(RU, [RT+RS*OS])$. $RD = f3(RD, RU, [RT+RS*OS])$. Limit $RS \leq IM2$ (unsigned).
2.0.6	2	0.6	E2	Four general purpose registers. $RD = f2(RS, RT)$. $RD = f3(RU, RS, RT)$.
2.0.7	2	0.7	E2	Three general purpose registers and a 16-bit integer with left shift. $RD = f2(RS, IM2)$. $RD = f3(RS, RT, IM2)$. IM2 (signed) is shifted left by the 6-bit unsigned value of IM3, or without shift if IM3 is used for other purposes.

2.1	2	1	A2	Two general purpose registers and a memory operand with base and 32 bit offset (IM2). RD = f2(RS, [RT+IM2]). RD = f3(RD, RS, [RT+IM2]).
2.2.0	2	2.0	E2	Two vector registers and a broadcast scalar memory operand with base and 16 bit offset. RD = f2(RU, [RT+IM2]). RD = f3(RD, RU, [RT+IM2]). Broadcast to length RS.
2.2.1	2	2.1	E2	Two vector registers and a memory operand with base and 16 bit offset. RD = f2(RU, [RT+IM2]). RD = f3(RD, RU, [RT+IM2]). Length=RS.
2.2.2	2	2.2	E2	Two vector registers and a scalar memory operand with base and scaled index. Optional support for a 16-bit offset. RD = f2(RU, [RT+RS*OS+IM2]). RD = f3(RD, RU, [RT+RS*OS+IM2]). IM2 must be zero if offset not supported.
2.2.3	2	2.3	E2	Two vector registers and a scalar memory operand with base, scaled index, and 16-bit limit. RD = f2(RU, [RT+RS*OS]). RD = f3(RD, RU, [RT+RS*OS]). Limit $RS \leq IM2$ (unsigned).
2.2.4	2	2.4	E2	Two vector registers and a memory operand with base and negative index. Optional support for a 16-bit offset. RD = f2(RU, [RT-RS+IM2]). RD = f3(RD, RU, [RT-RS+IM2]). Length=RS. IM2 must be zero if offset not supported.
2.2.6	2	2.6	E2	Four vector registers. RD = f2(RS, RT). RD = f3(RU, RS, RT).
2.2.7	2	2.7	E2	Three vector registers and a broadcast immediate half-precision float or 16-bit integer with left shift. RD = f2(RT, IM2). RD = f3(RS, RT, IM2). Floating point operands: IM2 is half precision. Integer operands: IM2 (signed) is shifted left by the 6-bit unsigned value of IM3, or without shift if IM3 is used for other purposes.
2.3	2	3	A2	Three vector registers and a broadcast 32-bit immediate operand. RD = f2(RT, IM2). RD = f3(RS, RT, IM2).
2.4	2	4	A2	One vector register and a memory operand with base and 32 bit offset. RD = f2(RD, [RT+IM2]). length=RS.
2.5	2	5	A2, B2, C2	Jump instructions for $OP1 < 8$. Single format instructions with memory operands or mixed register types for $OP1 \geq 8$.
2.6	2	6	A2	Single-format instructions. Three vector registers and a 32-bit immediate operand. RD = f2(RT, IM2). RD = f3(RS, RT, IM2).
2.7	2	7		Currently unused.

2.8	2	0 M=1	A2	Three general purpose registers and a 32-bit immediate operand. RD = f2(RT, IM2). RD = f3(RS, RT, IM2).
2.9	2	1 M=1	A2	Single-format instructions. Three general purpose registers and a 32-bit immediate operand. RD = f2(RT, IM2). RD = f3(RS, RT, IM2).
3.0.0	3	0.0	E3	Three general purpose registers and a memory operand with base and 32 bit offset. RD = f2(RS, [RT+IM4]). RD = f3(RU, RS, [RT+IM4]).
3.0.2	3	0.2	E3	Two general purpose registers and a memory operand w. base, scaled index, and 32 bit offset. RD = f2(RU, [RT+RS*OS+IM4]). RD = f3(RD, RU, [RT+RS*OS+IM4]). Optional.
3.0.3	3	0.3	E3	Two general purpose registers and a memory operand with base, scaled index, and 32-bit limit. RD = f2(RU, [RT+RS*OS]). RD = f3(RD, RU, [RT+RS*OS]). Limit $RS \leq IM4$ (unsigned).
3.0.7	3	0.7	E3	Three general purpose registers and a 32-bit integer with left shift. RD = f2(RS, IM4 << IM2). RD = f3(RS, RT, IM4 << IM2). IM4 (signed) is shifted left by the unsigned value of IM2.
3.1	3	1	A3, B3	Jump instructions for $OP1 < 8$. Single format instructions with memory operands or mixed register types for $OP1 \geq 8$.
3.2.0	3	2.0	E3	Two vector registers and a broadcast scalar memory operand with base and 32 bit offset. RD = f2(RU, [RT+IM4]). RD = f3(RD, RU, [RT+IM4]). Broadcast to length RS.
3.2.1	3	2.1	E3	Two vector registers and a memory operand with base and 32 bit offset. RD = f2(RU, [RT+IM4]). RD = f3(RD, RU, [RT+IM4]). Length=RS.
3.2.2	3	2.2	E3	Two vector registers and a scalar memory operand w. base, scaled index, and 32-bit offset. RD = f2(RU, [RT+RS*OS+IM4]). RD = f3(RD, RU, [RT+RS*OS+IM4]). Optional.
3.2.3	3	2.3	E3	Two vector registers and a scalar memory operand with base, scaled index, and 32-bit limit. RD = f2(RU, [RT+RS*OS]). RD = f3(RD, RU, [RT+RS*OS]). Limit $RS \leq IM4$ (unsigned).
3.2.7	3	2.7	E3	Three vector registers and a broadcast single precision float or 32-bit integer with left shift. RD = f2(RT, IM4). RD = f3(RS, RT, IM4). Floating point operands: IM4 is single precision. Integer operands: IM4 (signed) is shifted left by the unsigned value of IM2.
3.3	3	3	A3	Three vector registers and a broadcast 64-bit immediate operand. RD = f2(RT, IM2-3). RD = f3(RS, RT, IM2-3).

3.8	3	0 M=1	A3	Three general purpose registers and a 64-bit immediate operand. RD = f2(RT, IM2-3). RD = f3(RS, RT, IM2-3).
3.9				There is no format 3.9 because 3.1 uses the M bit.
4.x	3	4-7		Reserved for future 4-word instructions and longer.

Maximum number of input operands

The hardware supports a maximum of four, or preferably five, input dependencies. Three-input instructions cannot have both a mask and a memory operand with base and index or vector length if the hardware has a limit of four input dependencies. For example the mul_add instruction cannot have a mask in format 2.2.0 if the hardware does not support five inputs.

3.2 Coding of operands

Operand type

The type and size of operands is determined by the OT field as indicated above. The operand type is 64 bit integer by default if there is no OT field. The operand size (OS) is the size in bytes of a scalar operand or a vector element. This is equal to the number of bits divided by 8.

Register type

The instructions can use either general purpose registers or vector registers. General purpose registers are used for source and destination operands and for masks if the Mode field is 0 or 1 (with M = 0 or 1). Vector registers are used for source and destination operands and for masks if Mode is 2-7. Jump instructions use vector registers if M = 1. A few single-format instructions deviate from this rule and use mixed register types.

Pointer register

Instructions with a memory operand always use an address relative to a base pointer. The base pointer can be a general purpose register, the data section pointer, the thread data pointer, the instruction pointer, or the stack pointer. The pointer is determined by the RS or RT field. This field is interpreted as follows.

Single-size instructions (formats 0.4 - 0.9) can use any of the registers r0-r31 as base pointer. r31 is the stack pointer.

Double and triple-size instructions (formats 2.0.x, 2.1, 2.2.x, 2.4, 3.0.x, 3.2.x) can use the same registers, except r28 - r30, which are replaced by the thread pointer (THREADP), data section pointer (DATAP), and instruction pointer (IP), respectively.

Tiny instructions with a memory operand can use r0-r14 or the stack pointer (r31) as pointer in the 4-bit RS field. A value of 15 in the RS field indicates the stack pointer.

Index register

Instruction formats with an index can use r0 - r30 as index. A value of 31 in the index field (RS) means no index. The signed index is multiplied by the operand size (OS) for formats 0.6, 0.8, 2.0.2, 2.0.3, 2.2.3, 3.0.3, 3.2.3; by 1 for format 2.0.1; or by -1 for format 0.5 and 2.2.2. The result is added to the value of the base pointer.

Offsets

Offsets can be 8, 16 or 32 bits. The value is sign-extended to 64 bits. An 8-bit offset is multiplied by the operand size OS, as given by the OT field. An offset of 16 or 32 bits is not scaled. The result is added to the value of the base pointer.

Formats with 8 bits offset can use any general purpose register as base pointer. Formats with 16 or 32 bits offset use THREADP, DATAP, and IP as base pointers instead of r28, r29, and r30.

Support for addressing modes with both index and offset is optional (format 2.0.1, 2.0.2, 2.2.2, 2.2.4, 3.0.2, 3.2.2). If this kind of addressing involving two additions is not supported then the offset must be zero.

Limit on index

Formats 2.0.3, 2.2.3, 3.0.3, and 3.2.3 have a 16-bit or 32-bit limit on the index register. This is useful for checking array limits. A trap is generated if the value of the index register, interpreted as unsigned, is bigger than the unsigned limit. A negative index will also generate a trap because the negative index will be converted to a very large positive value when interpreted as unsigned.

Vector length

The vector length of memory operands is specified by r0-r30 in the RS field for formats 0.4, 0.5, 2.2.0, 2.2.1, 2.2.2, 2.4, 3.2.0, 3.2.1. A value of 31 in the RS field indicates a scalar with the same length as the operand size (OS).

The value of the vector length register gives the vector length in bytes (not the number of elements). If the value is bigger than the maximum vector length then the maximum vector length is used. The value may be zero. The behavior for negative values is implementation dependent: either interpret the value as unsigned or use the absolute value.

The vector length must be a multiple of the operand size OS, as indicated by the OT field. If the vector length is not a multiple of the operand size then the behavior of the partial vector element is implementation dependent.

The vector length for source operands in vector registers is included in the register itself.

Combining vectors with different lengths

The length of the destination register of a vector instruction will be the same as the vector length of the first register source operand (even if the first source operand uses the RD field).

A consequence of this is that the length of the result is determined by the order of the operands when vectors of different lengths are combined.

If the source operands have different lengths then the lengths will be adjusted as follows. If a vector source operand is too long then the extra elements will be ignored. If a vector source operand is too short then the missing elements will be zero.

A scalar memory operand (format 0.6 and 0.7) is not broadcast but treated as a short vector. It is padded with zeroes to the vector length of the destination.

A broadcast memory operand (format 2.2.0 and 3.2.0) will use the vector length given by the vector length register in the RS field.

A broadcast immediate operand will use the same vector length as the destination.

Immediate constants

Immediate constants can be 4, 8, 16, 32, and optionally 64 bits. Immediate fields are generally aligned to natural addresses. They are interpreted as follows.

If OT specifies an integer type then the field is interpreted as an integer. If the field is smaller than the operand size then it is sign-extended to the appropriate size, except for 4-bit fields which are zero-extended. If the field is larger than the operand size then the superfluous upper bits are ignored. The truncation of a too large immediate operand will not trigger any overflow condition.

If OT specifies a floating point type then the field is interpreted as follows. Immediate fields of 4 or 8 bits are interpreted as integers and converted to floating point numbers of the desired precision. A 16-bit field is interpreted as a half precision floating point number (subnormal values are not supported). A 32-bit field is interpreted as a single precision floating point number. It is converted to the desired precision if necessary. A 64-bit field (if supported) is interpreted as a double precision floating point number. A 64-bit field is not allowed with a single precision operand type.

Some instruction formats allow immediate integer constants with a left shift. Large integer constants with a limited number of significant bits can be represented with fewer bits in this way.

Format 2.0.7 and 2.2.7 allow a 16-bit immediate constant in IM2 to be shifted left by the unsigned value of IM3 to give a 64-bit signed value, except for instructions that use IM3 for other purposes.

Format 3.0.7 and 3.2.7 allow a 32-bit immediate constant in IM4 to be shifted left by the unsigned value of IM2. Any overflow beyond 64 bits is ignored.

Some single-format instructions also use shifted constants.

An instruction can be made compact by using the smallest size that fits the actual value of the constant.

Mask register and fallback register

The 3-bit mask field in formats with templates of type A or E indicates a mask register. Register r0-r6 can be used as masks if the destination is a general purpose register. Vector register v0-v6 can be used as masks if the destination is a vector register. A value of 7 in the mask field means no mask and unconditional execution using the options specified in the numeric control register.

If the mask is a vector register then it is interpreted as a vector with the same element size as indicated by the OT field. Each element of the mask register is applied to the corresponding element of the result.

The mask has multiple purposes. The primary purpose is for conditional execution. An instruction is not executed if bit 0 of the mask is zero. In this case, the destination will get a fallback value instead of the result of the calculation, and any numerical error condition will be suppressed. Vector instructions are executed conditionally for each vector element separately, so that each vector element is enabled if bit 0 of the corresponding vector element of the mask register is one.

The fallback value is taken from an extra register if the format has a vacant register, or from the first source register operand otherwise. If there is no source register operand and no vacant register field in the format template then the fallback value is zero.

Register r31 (stack pointer) and v31 cannot be used for fallback value. Instead, the fallback value will be zero if a register number of 31 is indicated. This also applies if the first source operand is used for fallback value. If the first source operand is r31 (stack pointer) or v31 and there is no other fallback register then the fallback value will be zero.

The remaining bits of the mask are used for specifying various options. The meanings of these mask bits are described in the next section.

3.3 Coding of masks

A mask register can be a general purpose register r0-r6 or a vector register v0-v6. A value of 7 in the mask field means no mask.

The bits in the mask register are coded as follows.

Table 3.16: Bits in mask register and numeric control register

Bit number	Meaning
0	Predicate or mask. The operation is executed only if this bit is one, as explained above.
6	Generate a trap if unsigned integer overflow.
7	Generate a trap if signed integer overflow.
10-15	Instruction-specific option bits.
18-19	Floating point rounding mode: 00 = nearest or even 01 = down 10 = up 11 = towards zero
20	Support subnormal numbers. Subnormal floating point numbers are treated as zero or flushed to zero when this bit is 0 (this is generally faster).
21	Better NAN propagation. If this bit is zero then the IEEE Standard 754-2008 (or later) is followed strictly for NAN values. A value of one in bit 21 improves NAN propagation and the use of NANs for tracing floating point errors. The details are described on page 94.
23	Constant execution time. This bit makes the instruction take the same number of clock cycles regardless of the values of mask and operands. Most or all instructions have constant execution time anyway, but the guarantee provided by this bit is useful for cryptographic applications. This feature is optional.
26	Generate a trap if floating point overflow or division by zero.
27	Generate a trap if floating point invalid operation.
28	Generate a trap if floating point underflow and precision loss.
29	Generate a trap if NAN inputs to compare instructions and floating point to integer conversion instructions.

Bits 8, 16, 24, etc. in a vector mask register can be used like bit 0 for 8-bit and 16-bit operand sizes. All other bits are reserved for future use.

Vector instructions treat the mask register as a vector with the same element size (OS) as the operands. Each element of the mask vector has the bit codes as listed above. The different vector elements can have different mask bits.

The numeric control register (NUMCONTR) is used as mask when the mask field is 7 or absent. The NUMCONTR register is broadcast to all elements of a vector, using as many bits of NUMCONTR as indicated by the operand size, when an instruction has no mask register. The same mask is applied to all vector elements in this case. Bit 0 in NUMCONTR must be 1.

The instruction-specific option bits (bit 10-15) are used for various options in compare, add_add, mul_add, and a few other instructions. The instruction-specific options are determined by IM3 in format 2.0.x, 2.2.x, 3.0.x, 3.2.x and by the mask register in other formats. The option bits in the mask are considered zero in vector operands with an 8-bit operand type because each mask element has only 8 bits in this case. The instruction-specific options are also zero when there is no IM3 and no mask. Unlike other options, the instruction-specific options are not taken from the numeric control

register. The reason for this is that frequent changes of the numeric control register would be inefficient (renaming of the numeric control register is not necessarily supported).

3.4 Format for jump, call and branch instructions

Most branches in ordinary code are based on the result of an arithmetic or logic instruction (ALU). The ForwardCom design combines the ALU instruction and the conditional jump into a single instruction. For example, a loop control can be implemented with a single instruction that counts down and jumps until it reaches zero or counts up until it reaches a certain limit.

The jumps, calls, branches and multiway branches will use the following formats.

Table 3.17: List of formats for control transfer instructions

Format	IL	Mode	OP1	Template	Description
1.4	1	4	OPJ	B	Short version with two register operands (RD, RS) and 8 bit offset (IM1).
1.5 C	1	5	OPJ	C	Short version with one register operand (RD), an 8-bit immediate constant (IM2) and 8 bit offset (IM1), or a 16-bit offset (IM2+IM1 combined).
1.5 D	1	5	0-15	D	Jump or call with 24-bit offset.
2.5.0	2	5	0	B2	Double size version with two register operands and 32 bit offset (IM2). IM1 = OPJ.
2.5.1	2	5	1	B2	Double size version with a register destination operand, a register source operand, a 16-bit offset (IM2 lower half) and a 16-bit immediate operand (IM2 upper half).
2.5.2	2	5	2	C2	Double size version with one register operand (RD), one 8-bit immediate constant (IM2) and 32 bit offset (IM3).
2.5.3	2	5	3	C2	Double size version with one register operand (RD), an 8-bit offset (IM2) and a 32-bit immediate constant (IM3).
2.5.4	2	5	4	C2	Double size system call, no OPJ, 16 bit constant (IM1,IM2) and 32-bit constant (IM3).
3.1.0	3	1	0	B3	Triple size version with a register destination operand, a register source operand, a 32-bit immediate operand (IM2) and a 32-bit offset (IM3). Optional.

The jump, call and branch instructions have signed offsets of 8, 16, 24 or 32 bits relative to the instruction pointer. Or, more precisely, relative to the end of the instruction. This offset is multiplied by the instruction word size (= 4) to cover an address range of \pm a half kilobyte for short conditional jumps with 8 bits offset, \pm 128 kbytes for jumps and calls with 16 bits offset, \pm 32 megabytes for 24 bits offset, and \pm 8 gigabytes for 32 bits offsets. The optional triple-size format includes unconditional jump and call with a 64 bits absolute address.

The versions with template C and C2 have no OT field. The operand type is 64-bit integer when there is no OT field. It is not possible to use formats with template C or C2 with floating point types. The instructions will use vector registers (first element only) when there is an OT field and $M = 1$. In other words, the combined ALU-and-branch instructions will use vector registers only when a floating

point type is specified (or 128-bit integer type, if supported). General purpose registers are used in all other cases. The logical instructions will interpret the value in a vector register as an integer, when a floating point type is specified. Only the compare instructions interpret the operands as floating point when a floating point type is indicated.

The OPJ field defines the operation and jump condition. This field has 6 bits in the single size version and 8 bits in the longer versions. The two extra bits in the longer versions are used as follows. Bit 6 is reserved for future extensions, and must be zero. Bit 7 may be used for indicating loop behavior as a hint for choosing the optimal branch prediction algorithm.

The lower 6 bits of the OPJ field contains the following codes.

Table 3.18: Condition codes for control transfer instructions

OPJ	bit 0 of OPJ	Instruction	Comment
0-7	part of offset	Unconditional jump with 24-bit offset (jump)	Format 1.5 D. Bit 0-2 of OPJ are part of offset
8-15	part of offset	Unconditional call with 24-bit offset (call)	Format 1.5 D. Bit 0-2 of OPJ are part of offset
0-1	invert	sub/jump_zero, sub/jump_nzero	Format 1.4 and 2.5.0. No floating point.
2-3	invert	sub/jump_neg, sub/jump_nneg	Format 1.4 and 2.5.0. No floating point.
4-5	invert	sub/jump_pos, sub/jump_npos	Format 1.4 and 2.5.0. No floating point.
6-7	invert	sub/jump_overfl, sub/jump_noverfl	Format 1.4 and 2.5.0. No floating point.
8-9	invert	sub/jump_borrow, sub/jump_nborrow	Format 1.4 and 2.5.0. No floating point.
10-11	invert	shift_left/jump_zero shift_left/jump_nzero	Not format 1.5
12-13	invert	shift_right_u/jump_zero shift_right_u/jump_nzero	Not format 1.5
14-15	invert	rotate/jump_carry, rotate/jump_ncarry	Not format 1.5
16-17	invert	add/jump_zero, add/jump_nzero	No floating point
18-19	invert	add/jump_neg, add/jump_nneg	No floating point
20-21	invert	add/jump_pos, add/jump_npos	No floating point
22-23	invert	add/jump_overfl, add/jump_noverfl	No floating point
24-25	invert	add/jump_carry, add/jump_ncarry	No floating point
24-25	invert	compare/jump_nfinite, compare/jump_finite	Floating point
26-27	invert	and/jump_zero, and/jump_nzero	
28-29	invert	or/jump_zero, or/jump_nzero	
30-31	invert	xor/jump_zero, xor/jump_nzero	

32-33	invert	compare/jump_equal, compare/jump_nequal	
34-35	invert	compare/jump_sbelow, compare/jump_saboveeq	
36-37	invert	compare/jump_sabove, compare/jump_sbeloweq	
38-39	invert	compare/jump_ubelow, compare/jump_uaboveeq	No floating point
38-39	invert	compare/jump_absbelow, compare/jump_absaboveeq	Floating point
40-41	invert	compare/jump_uabove, compare/jump_ubeloweq	No floating point
40-41	invert	compare/jump_absabove, compare/jump_absbeloweq	Floating point
42-43	invert	test_bit/jump_zero, test_bit/jump_nzero	
44-45	invert	test/jump_all1, test/jump_nall1	
46-49	invert	Reserved for future use.	
50-51	invert	increment_compare/jump_above, increment_compare/jump_beloweq	No floating point
52-53	invert	sub_maxlen/jump_pos, sub_maxlen/jump_npos	No floating point
54	0	sub/jump	No floating point
55	1	add/jump	No floating point
56-57		Reserved for future use.	
58-59	0 jump 1 call	Indirect jump or call with memory operand.	Format 1.4 and 2.5.0.
58-59	0 jump 1 call	Unconditional direct jump or call	Format 1.5 C, 2.5.2, and 3.1.0.
60-61	0 jump 1 call	Multiway jump or call with table of relative addresses	Format 1.4 A and 2.5.0
60-61	0 jump 1 call	Indirect jump or call to value of register	Format 1.5 C
62	0	return	Format 1.4 B
62	0	sys_return	Format 1.5 C
63	1	sys_call. ID in register	Format 1.4 A
63	1	sys_call. ID in constants	Format 2.5.1, 2.5.4 and 3.1.0.
63	1	trap or filler	Format 1.5 C
63	1	Conditional traps	Format 2.5.3.

The combined ALU and conditional jump instructions can be coded in the formats 1.4, 1.5 C, 2.5.0, 2.5.1, 2.5.2, 2.5.3, and 3.1.0, except subtraction which cannot be coded in format 1.5 C. Subtraction with an immediate constant can be replaced by addition with the negative constant. The code space that would have been used by subtraction in format 1.5 C is instead used for coding direct jump and call instructions with a 24-bit offset using format 1.5 D, where the lower three bits of OP1 are used as part of the 24-bit offset.

Unconditional and indirect jumps and calls use the formats indicated above, where unused fields must be zero. Bit 0 of the OPJ field is zero for jump instructions and one for call instructions.

See page 76 for detailed descriptions of control transfer instructions.

3.5 Assignment of opcodes

The opcodes and formats for new instructions can be assigned according to the following rules.

- Multi-format instructions. Often-used instructions that need to support many different operand types, addressing modes and formats use all or most of the following formats: 0.0 - 0.9, 2.0.x, 2.1, 2.2.x, 2.3, 2.4, 2.8, and optionally 3.0.x, 3.2.x, 3.3, and 3.8 if triple-size instructions are supported. The same value of OP1 is used in all these formats. OP2 must be 0. Instructions with few source operands should have the lowest values of OP1.
- Tiny instructions. Only some of the most common instructions are available in tiny versions, as there is only space for 32 tiny instructions. The instructions are ordered according to the number and type of operands, as shown in table 4.6 page 35.
- Control transfer instructions, i. e. jumps, branches, calls and returns, can be coded as short instructions with $IL = 1$, mode = 4 - 5, and $OP1 = 0 - 63$ or as double-size instructions with $IL = 2$, mode = 5, $OP1 = 0 - 7$, and optionally as triple-size instructions with $IL = 3$, mode = 1, $OP1 = 0-7$. See page 23.
- Short single-format instructions with general purpose registers. Use mode 1.0, 1.1, and 1.8, with any value of OP1.
- Short single-format instructions with vector registers. Use mode 1.2 and 1.3 with any value of OP1.
- Double-size single-format instructions with general purpose registers can use mode 2.9 with any value of OP1, and mode 2.0.x with any value of OP1 and $OP2 \neq 0$ (give similar instructions the same value of OP2). If more combinations are needed then use IM3 for further subdivision of the code space.
- Double-size single-format instructions with vector registers can use mode 2.6 with with any value of OP1, and mode 2.2.x with any value of OP1 and $OP2 \neq 0$ (give similar instructions the same value of OP2). If more combinations are needed then use IM3 for further subdivision of the code space.
- Double-size single-format instructions with mixed vector and general purpose registers or with memory operands can use mode 2.5 with OP1 in the range 8-63.
- Triple-size single-format instructions with general purpose registers can use mode 3.0.x with with any value of OP1 and $OP2 \neq 0$.
- Triple-size single-format instructions with vector registers can use mode 3.2.x with with any value of OP1 and $OP2 \neq 0$.
- Triple-size single-format instructions with mixed register types can use mode 3.1 with with OP1 in the range 8-63.
- Future instructions longer than three 32-bit words are coded with $IL = 3$, mode = 4-7.
- New options or other modifications to existing instructions can use IM3 bits in template E or mask register bits.
- New addressing modes and formats may be implemented as single-format read and write instructions. Template E formats use Mode2 and OP2 for distinguishing between different formats. Other single-format templates may be divided into groups of eight consecutive OP1 values with the same format. New addressing modes or other formats that apply to all multi-format instructions can use vacant values of Mode2 with E templates.

Application-specific instructions should preferably use E template formats with $OP2 \neq 0$. There are many vacant opcodes in these formats. General multi-purpose instructions may use some of the more crowded formats.

Unused register fields may have the same value as the first source register operand in order to avoid false dependences in a superscalar processor design. Unused mask fields have the value 7 in instructions that can have a mask. All other unused fields must be zero. The instructions with the fewest input operands should preferably have the lowest OP1 codes.

Chapter 4

Instruction lists

The ForwardCom instructions are listed in a comma-separated file `instruction_list.csv`. This file is intended for use by assemblers, disassemblers, debuggers and emulators. The list is preliminary and subject to possible changes. Please remember to keep the lists in this document and the list in the `instruction_list.csv` file synchronized.

The instruction list file has the following fields:

Table 4.1: Fields in instruction list file

Field	Meaning
Name	Name of instruction as used by assembler.
Category	1: single format instruction, 2: tiny instruction, 3: multi-format instruction, 4: jump instruction.
Formats	See table 4.2 below.
Template	Hexadecimal number: 0xA - 0xE for template A - E, 0x1 for tiny template, 0x0 for multiple templates.

Variant	<p>D0: No destination operand, no operand type. D1: No destination operand, but operand type specified. D2: Operand type ignored. I2: Immediate source operand is integer regardless of specified operand type. M0: Memory operand is destination. M1: E formats with a memory operand use IM3 as an extra immediate operand. On: n bits of IM3 in E template format used for options (IM3 can be used for shift count only if it is not used for options). R0: Destination is a general purpose register. R1: First source operand is a general purpose register. R2: Second source operand is a general purpose register. RL: RS is a general purpose register specifying vector length. U0: Integer operands are unsigned. U3: Integer operands are unsigned if option bit 3 is set. (compare instruction). H0: Half precision floating point instruction. X0: Source register can be a special pointer (threadp, datap, ip). X1: Source register is special register. X2: Source register is capabilities register. X3: Source register is performance monitor register. X4: Source register is system register. Y0-4: Destination register is one of the above.</p>
Source operands	Number of source operands, including register, memory and immediate operands, but not including mask, option bits, vector length, and index.
OP1	Operation code OP1.
OP2	Additional operation code OP2. Zero if none.
Operand types general purpose registers	Hexadecimal number indicating required and optional support for each operand type with general purpose registers. See table 4.3 below for meaning of each bit.
Operand types scalar	Hexadecimal number indicating required and optional support for each operand type for scalar operations in vector registers. See table 4.3 below for meaning of each bit.
Operand types vector	Hexadecimal number indicating required and optional support for each operand type for vector operations. See table 4.3 below for meaning of each bit.
Immediate operand type	Type of immediate operand for single-format instructions. See table 4.4 below.
Description	Description of the instruction and comments.

Table 4.2: Meaning of formats field in instruction list file

Category	Interpretation of formats field																										
1. Single format instruction	<p>Number with three hexadecimal digits.</p> <p>The leftmost digit is the value of the IL field (0-3).</p> <p>The middle digit is the value of mode field or the combined M+mode field (0-9).</p> <p>The rightmost digit is the sub-mode defined by OP2 in E template modes or OP1 in mode 2.5.x. Zero otherwise.</p> <p>For example 0x223 means format 2.2.3.</p>																										
2. Tiny instruction	<table> <tr><td>0</td><td>no operands.</td></tr> <tr><td>1</td><td>RD = general purpose destination register, RS = immediate operand.</td></tr> <tr><td>2</td><td>RD = g. p. destination register, RS = g. p. source register.</td></tr> <tr><td>4</td><td>RD = g. p. destination register, RS = pointer to memory source operand.</td></tr> <tr><td>5</td><td>RD = g. p. source register, RS = pointer to memory destination operand.</td></tr> <tr><td>8</td><td>RD = vector destination register, RS unused.</td></tr> <tr><td>9</td><td>RD = vector destination register, RS immediate operand.</td></tr> <tr><td>10</td><td>RD = vector destination register, RS vector source register.</td></tr> <tr><td>11</td><td>RD = vector source register, RS g. p. destination register r0-r14,r31.</td></tr> <tr><td>12</td><td>RD = vector destination register, RS = pointer to memory source operand.</td></tr> <tr><td>13</td><td>RD = vector source register, RS = pointer to memory destination operand.</td></tr> </table>	0	no operands.	1	RD = general purpose destination register, RS = immediate operand.	2	RD = g. p. destination register, RS = g. p. source register.	4	RD = g. p. destination register, RS = pointer to memory source operand.	5	RD = g. p. source register, RS = pointer to memory destination operand.	8	RD = vector destination register, RS unused.	9	RD = vector destination register, RS immediate operand.	10	RD = vector destination register, RS vector source register.	11	RD = vector source register, RS g. p. destination register r0-r14,r31.	12	RD = vector destination register, RS = pointer to memory source operand.	13	RD = vector source register, RS = pointer to memory destination operand.				
0	no operands.																										
1	RD = general purpose destination register, RS = immediate operand.																										
2	RD = g. p. destination register, RS = g. p. source register.																										
4	RD = g. p. destination register, RS = pointer to memory source operand.																										
5	RD = g. p. source register, RS = pointer to memory destination operand.																										
8	RD = vector destination register, RS unused.																										
9	RD = vector destination register, RS immediate operand.																										
10	RD = vector destination register, RS vector source register.																										
11	RD = vector source register, RS g. p. destination register r0-r14,r31.																										
12	RD = vector destination register, RS = pointer to memory source operand.																										
13	RD = vector source register, RS = pointer to memory destination operand.																										
3. Multi-format instruction	<p>Hexadecimal number composed of one bit for each format supported:</p> <table> <tr><td>0x0000001</td><td>Format 0.0: three general purpose registers.</td></tr> <tr><td>0x0000002</td><td>Format 0.1: two general purpose registers, 8-bit immediate.</td></tr> <tr><td>0x0000004</td><td>Format 0.2: Three vector registers.</td></tr> <tr><td>0x0000008</td><td>Format 0.3: Two vectors, 8-bit immediate.</td></tr> <tr><td>0x0000010</td><td>Format 0.4: One vector, memory operand.</td></tr> <tr><td>0x0000020</td><td>Format 0.5: One vector, memory operand with negative index.</td></tr> <tr><td>0x0000040</td><td>Format 0.6: One vector, scalar memory operand with index.</td></tr> <tr><td>0x0000080</td><td>Format 0.7: One vector, scalar memory operand with 8-bit offset.</td></tr> <tr><td>0x0000100</td><td>Format 0.8: One g. p. register, memory operand with index.</td></tr> <tr><td>0x0000200</td><td>Format 0.9: One g. p. register, memory operand with 8-bit offset.</td></tr> <tr><td>0x0001000</td><td>Format 2.8: Three g. p. registers, 32-bit immediate.</td></tr> <tr><td>0x0002000</td><td>Format 2.1: Two g. p. registers, memory with 32-bit offset.</td></tr> <tr><td>0x0004000</td><td>Format 2.3: Three vector registers, 32-bit immediate.</td></tr> </table>	0x0000001	Format 0.0: three general purpose registers.	0x0000002	Format 0.1: two general purpose registers, 8-bit immediate.	0x0000004	Format 0.2: Three vector registers.	0x0000008	Format 0.3: Two vectors, 8-bit immediate.	0x0000010	Format 0.4: One vector, memory operand.	0x0000020	Format 0.5: One vector, memory operand with negative index.	0x0000040	Format 0.6: One vector, scalar memory operand with index.	0x0000080	Format 0.7: One vector, scalar memory operand with 8-bit offset.	0x0000100	Format 0.8: One g. p. register, memory operand with index.	0x0000200	Format 0.9: One g. p. register, memory operand with 8-bit offset.	0x0001000	Format 2.8: Three g. p. registers, 32-bit immediate.	0x0002000	Format 2.1: Two g. p. registers, memory with 32-bit offset.	0x0004000	Format 2.3: Three vector registers, 32-bit immediate.
0x0000001	Format 0.0: three general purpose registers.																										
0x0000002	Format 0.1: two general purpose registers, 8-bit immediate.																										
0x0000004	Format 0.2: Three vector registers.																										
0x0000008	Format 0.3: Two vectors, 8-bit immediate.																										
0x0000010	Format 0.4: One vector, memory operand.																										
0x0000020	Format 0.5: One vector, memory operand with negative index.																										
0x0000040	Format 0.6: One vector, scalar memory operand with index.																										
0x0000080	Format 0.7: One vector, scalar memory operand with 8-bit offset.																										
0x0000100	Format 0.8: One g. p. register, memory operand with index.																										
0x0000200	Format 0.9: One g. p. register, memory operand with 8-bit offset.																										
0x0001000	Format 2.8: Three g. p. registers, 32-bit immediate.																										
0x0002000	Format 2.1: Two g. p. registers, memory with 32-bit offset.																										
0x0004000	Format 2.3: Three vector registers, 32-bit immediate.																										

	0x0008000	Format 2.4: One vector register, memory with 32-bit offset.
	0x0010000	Format 2.0.0: Three g. p. reg., memory with 16-bit offset.
	0x0020000	Format 2.0.1: Two g. p. reg., memory with unscaled index.
	0x0040000	Format 2.0.2: Two g. p. reg., memory with scaled index.
	0x0080000	Format 2.0.3: Two g. p. reg., memory with index and limit.
	0x0400000	Format 2.0.6: Four g. p. reg.
	0x0800000	Format 2.0.7: Three g. p. registers, 16-bit shifted immediate.
	0x1000000	Format 2.2.0: Two vector reg., scalar memory w. 16-bit offset.
	0x2000000	Format 2.2.1: Two vector reg., memory with 16-bit offset.
	0x4000000	Format 2.2.2: Two vector reg., memory with negative index.
	0x8000000	Format 2.2.3: Two vector reg., scalar memory w. index and limit.
	0x40000000	Format 2.2.6: Four vector reg.
	0x80000000	Format 2.2.7: Three vector registers, 16-bit shifted immediate.
	0x100000000	Format 3.8: Three g. p. registers, 64-bit immediate.
	0x400000000	Format 3.3: Three vector registers, 64-bit immediate.
	0x1000000000	Format 3.0.0: Three g. p. reg., memory with 32-bit offset.
	0x8000000000	Format 3.0.3: Two g. p. reg., memory with index and 32-bit limit.
	0x800000000000	Format 3.0.7: Three g. p. registers, 32-bit shifted immediate.
	0x1000000000000	Format 3.2.0: Two vector reg., scalar memory w. 32-bit offset.
	0x2000000000000	Format 3.2.1: Two vector reg., memory with 32-bit offset.
	0x8000000000000	Format 3.2.3: Two vector reg., scalar memory index and 32-bit limit.
	0x8000000000000000	Format 3.2.7: Three vector registers, float or 32-bit shifted immediate.
4. Jump instruction	Hexadecimal number composed of one bit for each format supported:	
	0x001	Format 1.4.0: Two registers, 8-bit address.
	0x002	Format 1.4.1: Memory operand with 8-bit offset.
	0x004	Format 1.4.2: Reg. and memory w. scaled index.
	0x008	Format 1.4.3: Three registers.
	0x010	Format 1.5.0 D: No register, 24-bit address.
	0x020	Format 1.5.1 C: One register, 8-bit immediate, 8-bit address.
	0x040	Format 1.5.2 C: 16-bit address.
	0x080	Format 1.5.3 C: One register.
	0x01000	Format 2.5.0: Two registers, 32-bit address.

0x02000	Format 2.5.0: Mem. w. base and 32-bit offset.
0x04000	Format 2.5.1: Two registers, 16-bit immediate, 16-bit address.
0x08000	Format 2.5.1: Two registers, 2x16-bit immediate.
0x10000	Format 2.5.2: One register, 8-bit immediate, 32-bit address.
0x20000	Format 2.5.3: One register, 32-bit immediate, 8-bit address.
0x40000	Format 2.5.3: Conditional trap. One register, 32-bit immediate, 8-bit immediate.
0x80000	Format 2.5.4: System call, 16-bit function, 32-bit module.
0x1000000	Format 3.1.0: Two registers, 32-bit immediate, 32-bit offset.
0x2000000	Format 3.1.1: Two registers, 2x32-bit immediate.

Table 4.3: Indication of operand types supported for general purpose registers, scalars in vector registers, or vectors. The value is a hexadecimal number composed of one bit for each operand type supported

0x0001	8-bit integer supported.
0x0002	16-bit integer supported.
0x0004	32-bit integer supported.
0x0008	64-bit integer supported.
0x0010	128-bit integer supported.
0x0020	single precision floating point supported.
0x0040	double precision floating point supported.
0x0080	quadruple precision floating point supported.
0x0100	8-bit integer optionally supported.
0x0200	16-bit integer optionally supported.
0x0400	32-bit integer optionally supported.
0x0800	64-bit integer optionally supported.
0x1000	128-bit integer optionally supported.
0x2000	single precision floating point optionally supported.
0x4000	double precision floating point optionally supported.
0x8000	quadruple precision floating point optionally supported.

Table 4.4: Immediate operand type for single-format instructions

0	none or multi-format.
1	4-bit signed integer.
2	8-bit signed integer.
3	16-bit signed integer.
4	32-bit signed integer.
5	64-bit signed integer.
6	8-bit signed integer shifted by specified count.
7	16-bit signed integer shifted by specified count.
8	16-bit signed integer shifted by 16.
9	32-bit signed integer shifted by 32.
17	4-bit unsigned integer.

18	8-bit unsigned integer.
19	16-bit unsigned integer.
20	32-bit unsigned integer.
21	64-bit unsigned integer.
24	two 8-bit unsigned integers.
25	two 8-bit and one 6-bit unsigned integers.
26	two 16-bit unsigned integers.
27	one 16-bit and one 32-bit unsigned integer.
28	two 32-bit unsigned integers.
29	one 16-bit and two 8-bit unsigned integers.
33	4-bit unsigned integer converted to float.
34	8-bit signed integer converted to float.
35	16-bit signed integer converted to float.
64	half precision floating point.
65	single precision floating point.
66	double precision floating point.
100	determined by operand type.

Jump instructions are listed on page 24. All other categories of instructions are listed in the following tables.

4.1 List of multi-format instructions

The following list covers general instructions that can be coded in most or all of the formats assigned to multi-format instructions.

Table 4.5: List of multi-format instructions

Instruction	OP1	Source operands	Description
nop	0	0	No operation.
store	1	1	Store value to memory.
move	2	1	Copy value.
prefetch	3	1	Prefetch from memory.
sign_extend	4	1	Sign-extend smaller integer to 64 bits.
sign_extend_add	5	1	Sign-extend smaller integer to 64 bits and add 64-bit register.
compare	7	2	Compare. Uses condition codes, see p. 57.
add	8	2	src1 + src2.
sub	9	2	src1 - src2.
sub_rev	10	2	src2 - src1.
mul	11	2	src1 · src2.
mul_hi	12	2	(src1 · src2) >> OS, signed (integer only).
mul_hi_u	13	2	(src1 · src2) >> OS, unsigned (integer only).
mul_ex	14	2	Multiply even-numbered signed integer vector elements to double size result.
mul_ex_u	15	2	Multiply even-numbered unsigned integer vector elements to double size result.
div	16	2	src1 / src2, signed division (optional for integer vectors).
div_u	17	2	src1 / src2, unsigned integer division (optional for vectors).

div_rev	18	2	src2 / src1, signed division (optional for integer vectors).
rem	20	2	Modulo or remainder, signed (optional for integer vectors).
rem_u	21	2	Modulo or remainder, unsigned (optional for integer vectors).
min	22	2	Signed minimum.
min_u	23	2	Minimum. unsigned for integers, abs for f.p.
max	24	2	Signed maximum.
max_u	25	2	Maximum. unsigned for integers, abs for f.p.
and	28	2	src1 & src2.
and_not	29	2	src1 & (~src2).
or	30	2	src1 src2.
xor	31	2	src1 ^ src2.
mul_2pow	32	2	src1 * 2 ^{src2} . Floating point multiply by signed integer power of 2.
shift_left	32	2	src1 << src2. Integer shift left.
rotate	33	2	Rotate left if src2 positive, right if negative.
shift_right_s	34	2	src1 >> src2. Integer shift right with sign extension.
shift_right_u	35	2	src1 >> src2. Integer shift right with zero extension.
set_bit	36	2	Set bit. src1 (1 << src2).
clear_bit	37	2	Clear bit. src1 & ~ (1 << src2).
toggle_bit	38	2	Toggle bit. src1 ^ (1 << src2).
and_bit	39	2	Clear all bits except one. src1 & (1 << src2).
test_bit	40	2	Test bit. (src1 >> src2) & 1.
test_bits	41	2	Test if at least one indicated bit is 1. (src1 & src2) != 0
test_bits_all1	42	2	Test if all indicated bits are 1. (src1 & src2) == src2
mul_add	48	3	± src1 ± src2 · src3 (optional).
mul_add2	49	3	± src1 · src2 ± src3 (optional).
add_add	50	3	± src1 ± src2 ± src3 (optional).
userdef55 - userdef62	55-62	2	Reserved for user-defined instructions.
undef	63	2	Undefined code. Guaranteed to generate trap in all future implementations.

4.2 List of tiny instructions

Tiny instructions are fitted two in one 32-bit code word. If a tiny instruction cannot be paired with anything else, it must be paired with a tiny nop.

Tiny instructions have an operand size of 64 bits unless otherwise noted. RD is the destination register, and in most cases also the first source register. RS can be a register r0-r15, v0-v15, or an immediate zero-extended 4-bit constant. Instructions with a pointer in RS use register r0-r14 as pointer when RS is 0-14, and the stack pointer (r31) when RS is 15.

It is not possible to jump to the second instruction in a tiny pair because instruction addresses must be divisible by four. If an interrupt or trap occurs in a tiny instruction then the interrupt handler must remember which of the two tiny instructions in the pair was interrupted.

Table 4.6: List of tiny instructions with general purpose registers

Instruction	OP1	Description
nop	0	No operation.
move	1	RD = unsigned constant RS.
add	2	RD += unsigned constant RS.
sub	3	RD -= unsigned constant RS.
shift_left	4	RD <<= unsigned constant RS.
shift_right_u	5	RD >>= unsigned constant RS (zero extended).
move	8	RD = register operand RS.
add	9	RD += register operand RS.
sub	10	RD -= register operand RS.
and	11	RD &= register operand RS.
or	12	RD = register operand RS.
xor	13	RD ^= register operand RS.
move	14	Read RD from memory operand with pointer RS (RS = r0-r14, r31).
store	15	Write RD to memory operand with pointer RS (RS = r0-r14, r31).

Table 4.7: List of tiny instructions with vector registers

Instruction	OP1	Description
clear	16	Clear register RD by setting the length to zero.
move	17	RD = RS. Copy vector of any type.
move	18	RD = unsigned 4-bit integer RS, converted to single precision scalar.
move	19	RD = unsigned 4-bit integer RS, converted to double precision scalar.
add	20	RD += RS, single precision float vector.
add	21	RD += RS, double precision float vector.
sub	22	RD -= RS, single precision float vector.
sub	23	RD -= RS, double precision float vector.
mul	24	RD *= RS, single precision float vector.
mul	25	RD *= RS, double precision float vector.
add_cps	28	Get size of compressed image for RD and add it to pointer register RS.
sub_cps	29	Get size of compressed image for RD and subtract it from pointer register RS.
restore_cp	30	Restore vector register RD from compressed image pointed to by RS.
save_cp	31	Save vector register RD to compressed image pointed to by RS.

4.3 List of single-format instructions

These instructions are mostly available in only one or a few formats.

Table 4.8: List of single-format instructions with general purpose registers

Instruction	Format	OP1	Description
bitscan_f	1.0	1	Bit scan forward. Find index to lowest set bit.
bitscan_r	1.0	2	Bit scan reverse. Find index to highest set bit.
round_d2	1.0	3	Round down to nearest power of 2.
round_u2	1.0	4	Round up to nearest power of 2.
move	1.1	0	Move 16-bit sign-extended constant to general purpose register.
move_u	1.1	1	Move 16-bit zero-extended constant to general purpose register.
add	1.1	2	Add 16-bit sign-extended constant.
mul	1.1	5	Multiply with 16-bit sign-extended constant.
div	1.1	6	Divide signed integer with 16-bit sign-extended constant.
add	1.1	7	$RD += (IM1, IM2) \ll 16$. Shift 16-bit signed constant left by 16 and add.
move	1.1	16	$RD = IM2 \ll IM1$. Sign-extend IM2 to 64 bits and shift left by the unsigned value IM1.
add	1.1	17	$RD += IM2 \ll IM1$. Sign-extend IM2 to 64 bits, shift left by the unsigned value IM1, add to RD.
and	1.1	18	$RD \&= IM2 \ll IM1$. Sign-extend IM2 to 64 bits, shift left by the unsigned value IM1, AND with RD.
or	1.1	19	$RD = IM2 \ll IM1$. Sign-extend IM2 to 64 bits, shift left by the unsigned value IM1, OR with RD.
xor	1.1	20	$RD \hat{=} IM2 \ll IM1$. Sign-extend IM2 to 64 bits, shift left by the unsigned value IM1, XOR with RD.
abs	1.8	0	Absolute value of integer. IM1 determines handling of overflow: 0: wrap around, 1: saturate, 2: zero.
shift_add	1.8	1	Shift and add. $RD += RS \ll IM1$
read_spec	1.8	32	Read special register RS into g. p. register RD.
write_spec	1.8	33	Write g. p. register RS to special register RD.
read_capabilities	1.8	34	Read capabilities register RS into g. p. register RD.
write_capabilities	1.8	35	Write g. p. register RS to capabilities register RD.
read_perf	1.8	36	Read performance counter.
read_perfs	1.8	37	Read performance counter, serializing.
read_sys	1.8	38	Read system register RS into g. p. register RD.
write_sys	1.8	39	Write g. p. register RS to system register RD.
move_bits	2.0.7	0.1	Replace one or more contiguous bits at one position of RS with contiguous bits from another position of RT. Optional.
move	2.9	0	Load 32-bit constant into the high part of a general purpose register. The low part is zero. $RD = IM2 \ll 32$.

insert_hi	2.9	1	Insert 32-bit constant into the high part of a general purpose register, leaving the low part unchanged. $RD = (RT \& 0xFFFFFFFF) (IM2 \ll 32)$.
add	2.9	2	Add zero-extended 32-bit constant to general purpose register.
sub	2.9	3	Subtract zero-extended 32-bit constant from general purpose register.
add	2.9	4	Add 32-bit constant to high part of general purpose register. $RD = RT + (IM2 \ll 32)$.
and	2.9	5	AND high part of general purpose register with 32-bit constant. $RD = RT \& (IM2 \ll 32)$.
or	2.9	6	OR high part of general purpose register with 32-bit constant. $RD = RT (IM2 \ll 32)$.
xor	2.9	7	XOR high part of general purpose register with 32-bit constant. $RD = RT \wedge (IM2 \ll 32)$.
replace_bits	2.9	9	Replace a group of contiguous bits in RT by a specified constant.
address	2.9	32	$RD = RT + IM2$, RT can be THREADP (28), DATAP (29) or IP (30).

Table 4.9: List of single-format instructions with vector registers and mixed register types

Instruction	Format	OP1. OP2	Description
set_len	1.2	0	$RD =$ vector register RT with length changed to value of RS.
get_len	1.2	1	Get length of vector register RT into general purpose register RD.
set_num	1.2	2	Change the length of vector register RT to $RS \cdot OS$.
get_num	1.2	3	Get length of vector register RT divided by the operand size.
compress	1.2	4	Compress vector RT of length RS to a vector of half the length and half the element size. Double precision \rightarrow single precision, 64-bit integer \rightarrow 32-bit integer, etc.
compress_ss	1.2	5	Compress integer vector RT of length RS to a vector of half the length and half the element size, signed with saturation (optional).
compress_us	1.2	6	Compress integer vector RT of length RS to a vector of half the length and half element size, unsigned with saturation (optional).
expand	1.2	7	Expand vector RT of length $RS/2$ and half the specified element size to a vector of length RS with the specified element size. Half precision \rightarrow single precision, 32-bit integer \rightarrow 64-bit integer with sign extension, etc.

expand_u	1.2	8	Expand integer vector RT of length RS/2 and half the specified element size to a vector of length RS with the specified element size. 32-bit integer → 64-bit integer with zero extension, etc.
compress_sparse	1.2	9	Compress sparse vector elements indicated by mask bits into contiguous vector. RS = length of input vector. (optional).
expand_sparse	1.2	10	Expand contiguous vector into sparse vector with positions indicated by mask bits. RS = length of output vector. (optional).
extract	1.2	11	Extract one element from vector RT, starting at offset RS·OS, with size OS into scalar in vector register RD.
insert	1.2	12	Replace one element in vector RD, starting at offset RS·OS, with scalar RT.
broad	1.2	13	Broadcast first element of vector RT into all elements of RD with length RS.
bits2bool	1.2	14	The lower n bits of RT are unpacked into a boolean vector RD with length RS, with one bit in each element, where $n = RS / OS$.
bool2bits	1.2	15	The boolean vector RT with length RS is packed into the lower n bits of RD, taking bit 0 of each element, where $n = RS / OS$. The length of RD is at least sufficient to contain n bits.
bool_reduce	1.2	16	The boolean vector RT with length RS is reduced by combining bit 0 of all elements. The output is a scalar integer where bit 0 is the AND combination of all the bits, and bit 1 is the OR combination of all the bits. The remaining bits are reserved for future use.
shift_expand	1.2	18	Shift vector RT up by RS bytes and extend the vector length by RS. The lower RS bytes of RD will be zero.
shift_reduce	1.2	19	Shift vector RT down RS bytes and reduce the length by RS. The lower RS bytes of RT are lost.
shift_up	1.2	20	Shift elements of vector RT up RS elements. The lower RS elements of RD will be zero, the upper RS elements of RT are lost.
shift_down	1.2	21	Shift elements of vector RT down RS elements. The upper RS elements of RD will be zero, the lower RS elements of RT are lost.
rotate_up	1.2	22	Rotate vector RT up one element. The length of the vector is RS bytes. Optional.
rotate_down	1.2	23	Rotate vector RT down one element. The length of the vector is RS bytes. Optional.

div_ex	1.2	24	Divide vector of double-size signed integers RS by signed integers RT. RS has element size 2·OS. These are divided by the even numbered elements of RT with size OS. The truncated results are stored in the even-numbered elements of RD. The remainders are stored in the odd-numbered elements of RD. (Optional for vectors).
div_ex_u	1.2	25	Same, with unsigned integers. (Optional for vectors).
sqrt	1.2	26	Square root (floating point, optional).
add_c	1.2	28	Add with carry. Vector has two elements. The upper element is used as carry on input and output (optional).
sub_b	1.2	29	Subtract with borrow. Vector has two elements. The upper element is used as borrow on input and output (optional).
add_ss	1.2	30	Add integer vectors, signed with saturation (optional).
add_us	1.2	31	Add integer vectors, unsigned with saturation (optional).
sub_ss	1.2	32	Subtract integer vectors, signed with saturation (optional).
sub_us	1.2	33	Subtract integer vectors, unsigned with saturation (optional).
mul_ss	1.2	34	Multiply integer vectors, signed with saturation (optional).
mul_us	1.2	35	Multiply integer vectors, unsigned with saturation (optional).
shift_ss	1.2	36	Shift left integer vectors, signed with saturation (optional).
shift_us	1.2	37	Shift left integer vectors, unsigned with saturation (optional).
add_oc	1.2	38	add with overflow check (optional).
sub_oc	1.2	39	subtract with overflow check (optional).
mul_oc	1.2	41	multiply with overflow check (optional).
div_oc	1.2	42	divide with overflow check (optional).
add_h	1.2	48	add half precision floating point (optional).
sub_h	1.2	49	subtract half precision floating point (optional).
mul_h	1.2	50	multiply half precision floating point (optional).
div_h	1.2	51	divide half precision floating point (optional).
mul_add_h	1.2	52	multiply and add half precision floating point (optional).
read_call_stack	1.2	58	read internal call stack. RD = vector register destination of length RS, RT-RS = internal address (privileged instruction).
write_call_stack	1.2	59	write internal call stack. RD = vector register source of length RS, RT-RS = internal address (privileged instruction).

read_memory_map	1.2	60	read memory map. RD = vector register destination of length RS, RT-RS = internal address (privileged instruction).
write_memory_map	1.2	61	write memory map. RD = vector register source of length RS, RT-RS = internal address (privileged instruction).
input	1.2	62	read from input port. RD = vector register, RT = port address, RS = vector length (privileged instruction).
output	1.2	63	write to output port. RD = vector register source operand, RT = port address, RS = vector length (privileged instruction).
gp2vec	1.3 B	0	Move value of general purpose register RS to scalar in vector register RD.
vec2gp	1.3 B	1	Move value of first element of vector register RS to general purpose register RD.
read_spev	1.3 B	2	Read special register RT into vector register RD with length RS.
make_sequence	1.3 B	4	Make a vector with RS sequential numbers. First value is IM1.
float2int	1.3 B	12	Conversion of floating point to integer with the same operand size. The rounding mode is specified in IM1.
int2float	1.3 B	13	Conversion of integer to floating point with same operand size.
round	1.3 B	14	Round floating point to integer in floating point representation. The rounding mode is specified in IM1.
round2n	1.3 B	15	Round to nearest multiple of 2^n . $RD = 2^n \cdot \text{round}(2^{-n} \cdot RS)$. n is a signed integer constant in IM1 (optional).
abs	1.3 B	16	Absolute value of integer. IM1 determines handling of overflow: 0: wrap around, 1: saturate, 2: zero.
fp_category	1.3 B	17	Check if floating point numbers belong to the categories indicated by constant.
broad	1.3 B	18	Broadcast 8-bit constant into all elements of RD with length RS (31 in RS field gives scalar output).
broadcast_max	1.3 B	19	Broadcast 8-bit constant into all elements of RD with maximum vector length.
byte_reverse	1.3 B	20	Reverse the order of bytes in each element of vector.
bit_reverse	1.3 B	21	Reverse the order of bits in each element of vector (optional).
bitscan_f	1.3 B	22	Bit scan forward. Find index to lowest set bit in RS (optional for vectors).
bitscan_r	1.3 B	23	Bit scan reverse. Find index to highest set bit in RS (optional for vectors).
popcount	1.3 B	24	Count the number of bits in RS that are 1 (optional for vectors).
truth_tab2	1.3 B	25	Boolean function of two inputs, given by a truth table.

move	1.3 C	32	Move 16 bit integer constant to 16-bit scalar (optional).
add	1.3 C	33	Add broadcasted 16 bit constant to 16-bit vector elements (optional).
and	1.3 C	34	AND broadcasted 16 bit constant with 16-bit vector elements (optional).
or	1.3 C	35	OR broadcasted 16 bit constant with 16-bit vector elements (optional).
xor	1.3 C	36	XOR broadcasted 16 bit constant with 16-bit vector elements (optional).
add_h	1.3 C	37	add constant to half precision vector (optional).
mul_h	1.3 C	38	multiply half precision vector with constant (optional).
move	1.3 C	40	$RD = IM2 \ll IM1$. Sign-extend IM2 to 32 bits and shift left by the unsigned value IM1 to make 32 bit scalar (optional).
move	1.3 C	41	$RD = IM2 \ll IM1$. Sign-extend IM2 to 64 bits and shift left by the unsigned value IM1 to make 64 bit scalar (optional).
add	1.3 C	42	$RD += IM2 \ll IM1$. Add broadcast shifted signed constant to 32-bit vector elements (optional).
add	1.3 C	43	$RD += IM2 \ll IM1$. Add broadcast shifted signed constant to 64-bit vector elements (optional).
and	1.3 C	44	$RD \&= IM2 \ll IM1$. AND broadcast shifted signed constant with 32-bit vector elements (optional).
and	1.3 C	45	$RD \&= IM2 \ll IM1$. AND broadcast shifted signed constant with 64-bit vector elements (optional).
or	1.3 C	46	$RD = IM2 \ll IM1$. OR broadcast shifted signed constant with 32-bit vector elements (optional).
or	1.3 C	47	$RD = IM2 \ll IM1$. OR broadcast shifted signed constant with 64-bit vector elements (optional).
xor	1.3 C	48	$RD \hat{=} IM2 \ll IM1$. XOR broadcast shifted signed constant with 32-bit vector elements (optional).
xor	1.3 C	49	$RD \hat{=} IM2 \ll IM1$. XOR broadcast shifted signed constant with 64-bit vector elements (optional).
move	1.3 C	56	Move converted half precision floating point constant to single precision scalar (optional).
move	1.3 C	57	Move converted half precision floating point constant to double precision scalar (optional).
add	1.3 C	58	Add broadcast half precision floating point constant to single precision vector (optional).
add	1.3 C	59	Add broadcast half precision floating point constant to double precision vector (optional).

mul	1.3 C	60	Multiply broadcast half precision floating point constant with single precision vector (optional).
mul	1.3 C	61	Multiply broadcast half precision floating point constant with double precision vector (optional).
concatenate	2.2.6	0.1	A vector RU of length RS and a vector RT of length RS are concatenated into a vector RD of length 2·RS.
permute	2.2.6	1.1	The vector elements of RU are permuted within each block of size RS bytes, using indices in RT. Each index is relative to the beginning of a block. An index out of range produces zero. The maximum block size is implementation dependent.
interleave	2.2.6	2.1	Interleave elements of vectors RU and RT of length RS/2 to produce vector RD of length RS. Even-numbered elements of the destination come from RU and odd-numbered elements from RT. (optional).
repeat_block	2.2.6	8.1	Repeat a block of data to make a longer vector. RT is input vector containing data block to repeat. IM2 is length in bytes of the block to repeat (must be a multiple of 4). RS is the length of destination vector RD. (optional).
repeat_within_blocks	2.2.6	9.1	Broadcast the first element of each block of data in a vector to the entire block. RT is input vector containing data blocks. IM2 is length in bytes of each block (must be a multiple of the operand size). RS is length of destination vector RD. The operand size must be at least 4 bytes. (optional).
move_bits	2.2.7	0.1	Replace one or more contiguous bits at one position of RS with contiguous bits from another position of RT. Optional
mask_length	2.2.7	1.1	Make mask with true in the first RS bytes. Option bits in IM2.
truth_tab3	2.2.7	2.1	Boolean function of three inputs, given by a truth table (optional).
load_hi	2.6	0	Make vector of two elements. dest[0] = 0, dest[1] = IM2.
insert_hi	2.6	1	Make vector of two elements. dest[0] = src1[0], dest[1] = IM2.
make_mask	2.6	2	Make vector where bit 0 of each element comes from bits in IM2, the remaining bits come from RT.
replace	2.6	3	Replace elements in RT by constant IM2.
replace_even	2.6	4	Replace even-numbered elements in RT by constant IM2.
replace_odd	2.6	5	Replace odd-numbered elements in RT by constant IM2.

broad	2.6	6	Broadcast 32-bit constant into all elements of RD with length RS (31 in RS field gives scalar output).
permute	2.6	8	The vector elements of RT are permuted within each block of size RS bytes. The 4·n bits of IM2 are used as index with 4 bits for each element in blocks of size n. The same pattern is used in each block. The number of elements in each block, $n = RS / OS \leq 8$.
replace_bits	2.6	9	Replace a group of contiguous bits in RT by a specified constant.
replace	3.1	8	Replace elements in RT by constant IM2,IM3.

Table 4.10: List of single-format instructions with memory operands.

Instruction	Format	OP1, OP2	Description
store	2.5 B	8	Store 32-bit constant IM2 to memory operand [RS+IM1] (optional).
fence	2.5 B	16	Memory fence at address [RS+IM2]. read, write or full indicated by IM1.
compare_swap	2.5 B	18	Atomic compare and exchange with address [RT+IM2].
read_insert	2.5 A	24	Replace one element in vector RD, starting at offset RS·OS, with scalar memory operand [RT+IM2] (optional).
extract_store	2.5 A	32	Extract one element from vector RD, starting at offset RS·OS, with size OS into memory operand [RT+IM2] (optional).

4.4 List of control transfer instructions

See page 24

Chapter 5

Description of instructions

Data move and conversion instructions

broad

format	opcode	operands
1.2 A	13	vector, g.p. register, and scalar
1.3 B	18	vector, g.p. register, and 8-bit signed constant
2.6	6	vector, g.p. register, and 32-bit signed or float constant

float v0 = broad(r1, v2)

float v0 = broad(r1, 2.5)

Broadcast a constant or the first element of a source vector into all elements of the destination vector with the length in bytes indicated by the first source operand (RS).

(This instruction is not called broadcast because that is a reserved keyword).

broadcast_max

format	opcode	operands
1.3 B	19	vector and 8-bit signed constant

float v0 = broadcast_max(1)

Broadcast a small constant to all elements of a vector with maximum length.

byte_reverse

format	opcode	operands
1.3 B	20	vectors

int32 v0 = byte_reverse(v1)

Reverse the order of bytes within each vector element. This is useful for converting big-endian file data.

clear

format	opcode	operands
Tiny	16	vectors

v0 = clear()

Clear vector register RD by setting the length to zero. All data are discarded. A cleared register is treated as unused.

compress

format	opcode	operands
1.2 A	4	vectors

double v0 = compress(r1, v2)

The elements of a vector are converted to half the element size. The length of the output vector will be half the length of the input vector. The OT field specifies the operand type of the input vector. Double precision floating point numbers are converted to single precision. Integer elements are converted to half the size by discarding the upper bits. Support for the following conversions are optional: single precision float to half precision, quadruple precision to double precision, 8-bit integer to 4-bit.

If the length of the input vector differs from the length specified by RS, then the length is converted to RS before compression.

compress_sparse

format	opcode	operands
1.2 A	9	vectors. Optional

Compress sparse vector elements indicated by mask bits into contiguous vector.

The length of the input vector is indicated by RS (bytes).

The length of the output vector is the number of true mask elements times the element size.

compress_ss

Compress signed with saturation.

See page 66.

compress_us

Compress unsigned with saturation.

See page 66.

concatenate

format	opcode	operands
2.2.6	0.1	vectors

float v0 = concatenate(v1, r2, v3)

A vector RU of length RS bytes and a vector RT of length RS bytes are concatenated into a vector RD of length 2·RS, with RT in the high end.

expand

format	opcode	operands
1.2 A	7	vectors

float v0 = expand(r1, v2)

This is the opposite of compress. The output vector has the specified length and the input vector has half this length. The OT field specifies the operand type of the output vector. Single precision floating point numbers are converted to double precision. Integers are converted to the double size by sign-extension. Support for the following conversions are optional: half precision float to single precision, double precision to quadruple precision, 4-bit integer to 8-bit.

If the length of the input vector differs from RS/2 then the length is converted before expansion. If the resulting length exceeds the maximum vector length for the specified operand type then the extra elements are lost.

expand_sparse

format	opcode	operands
1.2 A	10	vectors. Optional

This is the opposite of compress_sparse.

Expand contiguous vector into sparse vector with positions indicated by mask bits.

RS = length of output vector.

expand_u

format	opcode	operands
1.2 A	8	integer vectors

Same as expand, with zero extension.

Integers are expanded by zero-extension. Floating point operands cannot be used.

extract

format	opcode	operands
1.2 A	11	vectors

float v0 = extract(r1, v2)

Extract one element from vector RT, starting at offset RS·OS, with size OS into scalar in vector register RD. (OS = operand size).

An index out of range will produce zero. An operand size of 128 bits can be used, even if this size is not otherwise supported.

float2int

format	opcode	operands
1.3 B	12	vectors

double v0 = float2int(v1, 0)

Conversion of floating point to integer with the same operand size.

float32 is converted to int32. float64 is converted to int64.

The rounding mode is specified in bit 0-1 of IM1. See table 3.16 page 22.

get_len

format	opcode	operands
1.2 A	1	vectors

Get length in bytes of vector register RT into general purpose register RD.

get_num

format	opcode	operands
1.2 A	3	vectors

Get the number of elements in vector register RT into general purpose register RD. This is equal to the length divided by the operand size.

gp2vec

format	opcode	operands
1.3 B	0	g.p register in, vector register out

int64 v0 = gp2vec(r1)

Move integer value of general purpose register RS to scalar in vector register RD.

insert

format	opcode	operands
1.2 A	12	vectors

float v0 = insert(r1, v2)

Replace one element in vector RD, starting at offset RS·OS, with scalar RT. (OS = operand size).

An index out of range will leave the vector unchanged. An operand size of 128 bits can be used, even if this size is not otherwise supported.

insert_hi

format	opcode	operands
2.9	1	general purpose register, 32-bit immediate constant
2.6	1	vector register, 32-bit immediate constant

int64 r0 = insert_hi(r1, 2)

float v0 = insert_hi(v1, 2.1)

Insert 32-bit constant into the high part of a general purpose register, leaving the low part unchanged.
 $dest = (src1 \& 0xFFFFFFFF) | (IM2 \ll 32)$.

Make a vector of two elements. A constant is inserted into the second element, leaving the first element unchanged.

$dest[0] = src1[0], dest[1] = IM2$.

int2float

format	opcode	operands
1.3 B	13	vectors

int64 v0 = int2float(v1)

Conversion of integer to floating point with same operand size.

int32 is converted to float32. int64 is converted to float64.

interleave

format	opcode	operands
2.2.6	2.1	vectors. Optional

float v0 = interleave(v1, r2, v3)

Interleave the inputs from two vectors so that the even-numbered elements come from the first input vector and the odd-numbered elements come from the second input vector. The length in bytes of the destination vector is indicated by a general purpose register. The length of each input vector is half the indicated value.

load_hi

format	opcode	operands
2.5	0	vector. 32 bit immediate constant

float v0 = load_hi(1.2)

Make vector of two elements. $dest[0] = 0, dest[1] = IM2$.

move

format	opcode	operands
multi	2	all types
tiny	1	g.p. register = 4-bit zero-extended constant
tiny	8	two g.p. registers
tiny	14	g.p. register = memory operand with pointer
tiny	17	single prec. float scalar = 4 bit unsigned integer constant
tiny	18	double prec. float scalar = 4 bit unsigned integer constant
tiny	19	two vector registers
1.1 C	0	g.p. register = 16-bit sign extended constant
1.1 C	16	g.p. register = 8-bit sign extended constant with left shift
1.3 C	32	vector register 16-bit scalar = 16-bit constant. Optional
1.3 C	40	vector register 32-bit scalar = 8-bit sign extended constant with left shift. Optional
1.3 C	41	vector register 64-bit scalar = 8-bit sign extended constant with left shift. Optional
1.3 C	56	vector register single precision scalar = half precision immediate constant. Optional
1.3 C	57	vector register double precision scalar = half precision immediate constant. Optional

Copy A value from a register, memory operand or immediate constant to a register. If the destination is a vector register and the source is an immediate constant then the result will be a scalar. The value will not be broadcast because there is no other input operand that specifies the vector length. If a vector is desired then use the broadcast instruction instead.

The move instruction with an immediate operand is the preferred method for setting a register to zero.

move_u

format	opcode	operands
1.1 C	1	g.p. register, 16 bit unsigned constant

Copy 16-bit zero-extended constant to general purpose register.

This can be used as the first step of loading a 32-bit constant if double size instructions are not supported.

permute

format	opcode	operands
2.2.6	1.1	vectors
2.6	8	vectors and 32 bit immediate constant

float v0 = permute(v1, r2, v3)

float v0 = permute(r1, v2, 3)

This instruction permutes the elements of a vector. The vector is divided into blocks of size RS bytes each. The block size must be a power of 2 and a multiple of the operand size. Elements can be moved arbitrarily between positions within each block, but not between blocks. Each element of the output vector is a copy of an element in the input vector, selected by the corresponding index in an index vector. The indexes are relative to the start of the block they belong to, so that an index of zero will select the first element in the block of the input vector and insert it in the corresponding position of the output vector. The same element in the input vector can be copied to multiple elements in the output vector. An index out of range will produce a zero. The indexes are interpreted as integers regardless of the operand type.

The permute instruction has two versions. The first version specifies the indexes in a vector with the same length and element size as the input vector.

The second version specifies the indexes as a 32-bit immediate constant with 4 bits per element. This constant is split into a maximum of 8 elements with 4 bits in each. If the blocks have more than 8 elements each then the sequence of 8 elements is repeated to fill a block. The same pattern of indexes will be applied to all blocks in the second version of the permute instruction.

The maximum block size for the permute instruction is implementation-dependent and given by a special register. The reason for this limitation of block size is that the complexity of the hardware grows quadratically with the block size. A full permutation is possible if the vector length does not exceed the maximum block size. A trap is generated if RS is bigger than the maximum block size.

There are two ways to combine the outputs of multiple permute instructions. One method is to use indexes out of range to produce zeroes for unused outputs and then OR'ing the outputs. Another method is to use masks to combine the outputs.

Permute instructions are essential for a vector processor because it is often necessary to rearrange data to facilitate the vector processing. These instructions are useful for reordering data, for transposing a matrix, etc.

Permute instructions can also be used for parallel table lookup when the block size is big enough to contain the entire table.

Finally, permute instructions can be used for gathering and scattering data within an area not bigger than the vector length or the block size.

read_insert

format	opcode	operands
2.5 A	24	vectors. Optional

`int32 v0 = read_insert(v0, r1, [r2+0x8, scalar])`

Replace one element in vector RD, starting at offset RS-OS, with scalar memory operand [RT+IM2]. (OS = operand size).

repeat_block

format	opcode	operands
2.2.6	8.1	vectors. Optional

`float v0 = repeat_block(r1, v2, 8)`

Repeat a block of data to make a longer vector. This is the same as broadcast, but with a larger block of data. RT is an input vector containing a data block to repeat. IM2 is the length in bytes of the block to repeat. This must be a multiple of 4. RS is the length in bytes of the destination vector RD. This instruction is useful for matrix multiplication.

repeat_within_blocks

format	opcode	operands
2.2.6	9.1	vectors. Optional

`float v0 = repeat_within_blocks(r1, v2, 8)`

This divides a vector into blocks and broadcasts the first element of each block to the rest of the block. The block size is given by IM2. This must be a multiple of the operand size, and at least 4 bytes. There may be a maximum limit to the block size. RS is the length in bytes of the destination vector RD. This instruction is useful for matrix multiplication.

For example, if the input vector contains (0,1,2,3,4,5,6,7,8) and the block size is 3 times the operand size, then the result will be (0,0,0,3,3,3,6,6,6).

replace

format	opcode	operands
2.6	3	vectors and 32-bit immediate constant
3.1	8	vectors and 64-bit immediate constant. Optional

int32 v0 = replace(v1, 1), mask=v2, fallback=v3

double v0 = replace(v1, 2.3)

All elements of src1 are replaced by the integer or floating point constant src2.

When used without a mask, the constant is simply broadcast to make a vector of the same length as src1. This is useful for broadcasting a constant to all elements of a vector. Only the length of src1 (in bytes) is used, not its contents, when this instruction is used without a mask.

When used with a mask, the elements of src1 are selectively replaced. Elements that are not selected by the mask will be taken from a fallback register.

replace_even

format	opcode	operands
2.6	4	vectors and 32-bit immediate constant

Same as replace. Only even-numbered vector elements are replaced.

replace_odd

format	opcode	operands
2.6	5	vectors and 32-bit immediate constant

Same as replace. Only odd-numbered vector elements are replaced.

rotate_up

format	opcode	operands
1.2	22	vectors. Optional

int32 v0 = rotate_up(r1, v2)

Rotate vector up one element. Element number n is moved to position n+1, and the last element is moved to position zero. The length of the vector in bytes is indicated by general purpose register RS.

rotate_down

format	opcode	operands
1.2	23	vectors. Optional

int32 v0 = rotate_down(r1, v2)

Rotate vector down one element. Element number n is moved to position n-1, and element number zero is moved to the last position. The length of the vector in bytes is indicated by general purpose register RS.

set_len

format	opcode	operands
1.2	0	vectors

v1 = set_len(r2, v3)

Sets the length of a vector register to the number of bytes specified by a general purpose register. If the specified length is more than the maximum length for the specified operand type then the maximum length will be used.

If the output vector is longer than the input vector then the extra elements will be zero. If the output vector is shorter than the input vector then the extra elements will be discarded.

set_num

format	opcode	operands
1.2	2	vectors

The length of vector register RT is changed to the value of general purpose register RS. The length is indicated as number of elements. If the length is increased then the extra elements will be zero. If the length is decreased then the superfluous elements are lost.

This instruction differs from set_len by multiplying RS with the operand size.

shift_down

format	opcode	operands
1.2	21	vectors

int32 v0 = shift_down(r1, v2)

Shift elements of vector RT down by the number of elements indicated by general purpose register RS. The upper RS elements of RD will be zero, the lower RS elements of RT are lost. The length of the vector is not changed.

This instruction differs from shift_reduce by indicating the shift count as a number of elements rather than a number of bytes, and by not changing the length of the vector.

shift_expand

format	opcode	operands
1.2	18	vectors

int32 v0 = shift_expand(r1, v2)

The length of a vector is expanded by the specified number of bytes by adding zero-bytes at the low end and shifting all bytes up. If the resulting length is more than the maximum vector length for the specified operand type then the upper bytes are lost.

shift_reduce

format	opcode	operands
1.2	19	vectors

int32 v0 = shift_reduce(r1, v2)

The length of a vector is reduced by the specified number of bytes by removing bytes at the low end and shifting all bytes down. If the resulting length is less than zero then the result will be a zero-length vector. The specified operand type is ignored.

shift_up

format	opcode	operands
1.2	20	vectors

int32 v0 = shift_up(r1, v2)

Shift elements of vector RT up by the number of elements indicated by general purpose register RS. The lower RS elements of RD will be zero, the upper RS elements of RT are lost. The length of the vector is not changed.

This instruction differs from `shift_expand` by indicating the shift count as a number of elements rather than a number of bytes, and by not changing the length of the vector.

sign_extend

format	opcode	operands
multi	4	general purpose and integer scalar

`int8 r0 = sign_extend(r1) // result is 64 bits`

`int8 v0 = sign_extend(v1) // lower 8 bits of each 64-bit element is extended to 64 bits`

`int8 v0 = sign_extend([r1, scalar]) // memory operand is 8 bits, result is 64 bits scalar`

Sign-extend smaller integer to 64 bits.

The input can be an 8-bit, 16-bit or 32-bit integer. This integer is sign-extended to produce a 64-bit output in a general purpose register or a scalar in a vector register. If the input is a vector then only the first element in each 64-bit block of the input vector is used. Floating point types cannot be used.

sign_extend_add

format	opcode	operands
multi	5	general purpose registers. Optional

`int8 r0 = sign_extend_add(r1, r2)`

`int32 r0 = sign_extend_add(r1, [r2])`

`src2` is an integer of 8, 16, or 32 bits, usually a memory operand. This integer is sign-extended to produce a 64-bit integer. The value is added to the 64-bit integer in `src1` and the result is stored in the 64-bit destination register.

This instruction is useful for loading relative pointers from memory, where the reference point is in `src1`.

vec2gp

format	opcode	operands
1.3 B	1	vector register in, g.p. register out

`int64 r0 = vec2gp(v1)`

Copy value of first element of vector register RS to general purpose register RD. Integers are sign-extended. Single precision floating point values are zero-extended.

Data read and write instructions

add_cps

format	opcode	operands
tiny	28	vector and g.p. register

Add compressed size.

This instruction gets the size of the compressed image for vector RD and adds it to general purpose register RS. This is used for updating the stack pointer or any other pointer after restoring a vector from a compressed memory image.

See `save_cp` page 53 for details.

extract_store

format	opcode	operands
2.5 A	32	vector. Optional

int32 [r3+8, scalar] = extract_store(v1, r2)

Extract one element from vector RD, starting at offset RS·OS, with size OS into memory operand [RT+IM2].

(OS = operand size).

fence

format	opcode	operands
2.5 B	16	memory operand and immediate. Optional

int32 fence([r1], 2)

Memory fence at address [RS+IM2].

Options indicated by IM1:

IM1 value	meaning
1	read fence
2	write fence
3	read and write fence

move

The move instruction, described at page 48 can read a register from a memory operand.

prefetch

format	opcode	operands
multi	3	memory operand. Optional

Prefetch memory operand into cache for later read or write. Different variants (not yet defined) can be specified by IM3 for format with E template.

restore_cp

format	opcode	operands
tiny	30	vector and memory operand

Restore compressed image.

This will restore vector register RD from a compressed memory image pointed to by general purpose register RS. The memory image must have been written previously with the save_cp instruction. See save_cp for details.

save_cp

format	opcode	operands
tiny	31	memory operand and vector

Save compressed image.

This will save vector register RD to a compressed memory image pointed to by general purpose register RS.

This instruction is used for saving a vector register to the stack or to any other memory position in an efficient way that includes information about the vector length. When saving a vector register with variable length, we do not want to save the maximum length when only part of the register is used. Instead, we are using the `save_cp` and `restore_cp` instructions which are intended for saving and restoring a vector register without using more memory than necessary.

The format of the compressed memory image is implementation-dependent. Typically, it will contain an integer to indicate the vector length, followed by as many bytes of data as the length indicates. Any additional compression is allowed if it can be done sufficiently fast. It is recommended to include a checksum or hash of the data in order to check if the stack has been corrupted. The `restore_cp` instruction may generate a trap when attempting to restore a vector where the checksum does not match the data. The compressed image of an empty vector will normally use only a single stack entry (8 bytes).

The size of the compressed image must be a multiple of the stack word size if the pointer register is the stack pointer.

The software should never use the saved image for anything else than restoring a vector register on the same microprocessor model that saved it, because the image format is not compatible across microprocessors.

The size of the saved image can be added to a pointer with the `add_cps` instruction or subtracted from a pointer with the `sub_cps` instruction. `RS` indicates the pointer, which can be `r0-14` or `r31` (stack pointer).

A vector register `V` can be saved (pushed) on the stack with the following pair of tiny instructions:

```
sp = sub_cps(sp, V)
[sp] = save_cp(V)
```

A vector register `V` can be restored (popped) from the stack with the following pair of tiny instructions:

```
V = restore_cp([sp])
sp = add_cps(sp, V)
```

The same instructions can be used for saving vector registers during a task switch. Unused vector registers will only use very little space when saved in this way.

The `restore_cp` instruction is allowed to read more bytes than necessary, up to the maximum vector length plus 8 bytes, and discard any superfluous bytes afterwards when the actual length is known.

store

format	opcode	operands
multi	1	memory operand and g.p. or vector register
tiny	15	memory operand and g.p. register
2.5 B	8	memory operand and 32-bit constant. Optional

```
int32 [r0+r1*4] = r1
```

```
float [r0, length = r1] = v2
```

```
float [r0 + 0x10] = 2.5
```

Write the value of a register or constant to a memory operand.

The size of the memory operand is determined by the operand size `OS` when a scalar memory operand is specified, or by the vector length register in `RS` when a vector operand is specified.

An immediate constant cannot be bigger than 32 bits. A 64 bit integer constant can only be used if it fits into a 32-bit signed integer. A float64 constant can only be used if it can be represented as single precision without loss of precision.

The hardware must be able to handle memory operand sizes that are not powers of 2 without touching additional memory (read and rewrite beyond the memory operand is not allowed unless access from other threads is blocked during the operation and any access violation is suppressed). It is allowed for the hardware to write the operand in a piecemeal fashion.

Masked operation with a mask of zero will store a fallback value or zero. Masking cannot leave a memory element untouched.

sub_cps

format	opcode	operands
tiny	29	vector and g.p. register

Subtract compressed size.

This instruction gets the size of the compressed image for vector RD and subtracts it from general purpose register RS. This is used for adjusting the stack pointer or any other pointer before saving a vector to a compressed memory image.

See save_cp page 53 for details.

General arithmetic instructions

abs

format	opcode	operands
1.8 B	0	g.p. registers
1.3 B	16	vector registers

int32 r0 = abs(r1, 1)

Absolute value of signed number.

Signed integers can overflow when the input is the minimum value. The handling of overflow for signed integers is controlled by the constant IM1 as follows:

IM1	result when input is INT_MIN
0	INT_MIN (wrap around)
1	INT_MAX (saturation)
2	zero
4	generate trap (interrupt)

add

format	opcode	operands
multi	8	all types
tiny	2	g.p. register and 4-bit sign-extended constant
tiny	9	two g.p. registers
tiny	20	two vector registers, single precision float
tiny	21	two vector registers, double precision float
1.1 C	2	g.p. register and 16-bit sign-extended constant
1.1 C	7	g.p. register and 16-bit sign-extended constant shifted left by 16
1.1 C	17	g.p. register and 8-bit sign-extended constant shifted left by another constant
2.9	2	g.p. register and 32-bit zero-extended constant
2.9	4	g.p. register and 32-bit constant shifted left by 32
1.3 C	33	vector of 16-bit integer elements and broadcast 16 bit integer constant. Optional
1.3 C	42	vector of 32-bit integer elements and broadcast 8-bit sign-extended constant shifted left by another constant. Optional
1.3 C	43	vector of 64-bit integer elements and broadcast 8-bit sign-extended constant shifted left by another constant. Optional
1.3 C	58	single precision floating point vector and broadcast half precision floating point constant. Optional
1.3 C	59	double precision floating point vector and broadcast half precision floating point constant. Optional

$\text{int32 } r0 = r1 + r2$

$\text{int32 } r0 = r1 + 2$

$\text{int32+ } r0 += 4$

$\text{int32+ } r0++$

$\text{float } v0 = v1 + [r2 + 8, \text{length} = r5]$

Addition.

If you want to add a 64-bit constant to a general purpose register, and triple size instructions are not supported, then add the lower half first using the zero-extended version, and then add the upper half using the shifted version.

add_add

format	opcode	operands
multi	50	all types. Optional

This is two additions in one instruction:

$\text{dest} = \pm \text{src1} \pm \text{src2} \pm \text{src3}$

For optimal precision with floating point operands, the intermediate sum of the two numerically largest operands should preferably be calculated first with extended precision.

Only instruction formats that allow three operands are supported.

The signs of the operands can be inverted as indicated by the following 3-bit code

Table 5.3: Control bits for add_add

Format with E template	Other formats with mask	Meaning
IM3 bit 0	mask bit 10	change sign of src1

IM3 bit 1	mask bit 11	change sign of src2
IM3 bit 2	mask bit 12	change sign of src3

There is no sign change if there is no IM3 field and no mask.

This instruction should only be supported if it can be implemented so that it is faster than two consecutive add instructions. It may be supported for integer operands or floating point or both.

compare

format	opcode	operands
multi	7	all types

int8 r0 = r1 > r2

uint8 r0 = r1 > r2

float v0 = v1 >= 2.3

int32 r0 = compare(r1, 2), mask=r3, fallback=r4, options=0x11

The compare instruction compares two source operands and generates a boolean scalar or vector where bit 0 indicates the result, and the remaining bits are copied from the mask or numeric control register. This instruction can do different compare operations depending on the following 4-bit condition code:

Table 5.4: Condition codes for compare instruction

Bit 3-2-1-0	Meaning for integer	Meaning for float
_ 0 0 0	$a = b$	$a = b$
_ 0 0 1	$a \neq b$	$a \neq b$
_ 0 1 0	$a < b$	$a < b$
_ 0 1 1	$a \geq b$	$a \geq b$
_ 1 0 0	$a > b$	$a > b$
_ 1 0 1	$a \leq b$	$a \leq b$
_ 1 1 0	$\text{abs}(a) < \text{abs}(b)$	$\text{abs}(a) < \text{abs}(b)$
_ 1 1 1	$\text{abs}(a) \geq \text{abs}(b)$	$\text{abs}(a) \geq \text{abs}(b)$
0 _ _ _	compare as signed	unordered gives 0
1 _ _ _	compare as unsigned	unordered gives 1

Details: The absolute values of integers are corrected for overflow so that $\text{abs}(\text{INT_MIN}) > \text{abs}(\text{INT_MAX})$. If a and b are both infinity then $a = b$ will be true, but $a - b = 0$ will be unordered. You can check if a value is NAN by comparing it for unequal to itself with condition code 1001.

The condition codes are provided from different sources, depending on the instruction format:

Table 5.5: Source of condition codes

Formats	Condition code from	Remaining bits from
A template	mask bit 10-13	mask
B template	zero	numeric control register
E template	IM3 bit 0-3	mask

The condition code is zero (indicating compare for equal) if there is no mask and no IM3 field. The remaining bits are taken from the numeric control register if there is no mask.

Masking and fallback is possible. Alternative use of the fallback register (RU) as an extra boolean operand is supported with E template formats. This option is controlled by bits 4-5 of IM3:

Table 5.6: Alternative use of fallback register

bit 5 bit 4	Output
0 0	mask ? result : fallback
0 1	mask & result & fallback
1 0	mask & (result fallback)
1 1	mask & (result ^ fallback)

div

format	opcode	operands
multi	16	all types. Optional for integer vectors
1.1 C	6	g.p. register and 16-bit sign-extended constant

int32 r0 = r1 / r2

int32 r0 = div(r1, r2), options = 4

float v0 = v1 / [r2 + 8, length = r5]

Signed division.

This instruction has multiple rounding modes. The rounding mode for integer operands is controlled by IM3 in E template formats as follows:

Table 5.7: division instructions

IM3 bits 0-3	Meaning
0 0 0 0	Truncate towards zero
0 1 0 0	Nearest or even
0 1 0 1	Down
0 1 1 0	Up
0 1 1 1	Truncate towards zero
1 0 0 0	Rounding mode determined by bits 18-19 of mask or numeric control register
other values	Not allowed

Truncation is always used with integer operands in formats with no IM3 field.

The rounding mode for floating point operands is controlled by the mask or numeric control register. IM3 must be zero for floating point operands.

Division of floating point operands by zero gives $\pm\text{INF}$.

Division of integer operands by zero gives INT_MAX or INT_MIN.

Overflow occurs by division of INT_MIN by -1. The result will wrap around to give INT_MIN.

div_ex

format	opcode	operands
1.2 A	24	Integer vectors. Optional for more than one element

Divide vector of double-size signed integers RS by signed integers RT. RS has element size 2·OS. These are divided by the even numbered elements of RT with size OS. The truncated results are stored in the even-numbered elements of RD. The remainders are stored in the odd-numbered elements of RD. (OS = operand size).

div_ex_u

format	opcode	operands
1.2 A	25	Integer vectors. Optional for more than one element

Divide vector of double-size unsigned integers RS by unsigned integers RT. RS has element size 2·OS. These are divided by the even numbered elements of RT with size OS. The truncated results are stored in the even-numbered elements of RD. The remainders are stored in the odd-numbered elements of RD. (OS = operand size).

div_rev

format	opcode	operands
multi	18	all types. Optional for integer vectors

int32 r0 = 10 / r2

int32 v0 = div_rev(v1, v2), options = 4

Same as div, with the two source operands swapped.

The rounding mode is controlled in the same way as for the div instruction.

div_u

format	opcode	operands
multi	17	all integer types. Optional for integer vectors

uint32 r0 = r1 / r2

uint32 v0 = div_u(v1, v2), options=4

Unsigned integer division.

The rounding mode is controlled in the same way as for the div instruction.

Division by zero gives UINT_MAX.

max

format	opcode	operands
multi	24	all types

int32 r0 = max(r1, r2)

float v0 = max(v1, v2)

Get the maximum of two numbers:

max(src1,src2) = src1 > src2 ? src1 : src2

Integer operands are treated as signed.

The handling of floating point NAN operands is determined by bit 21 of the mask register or the numeric control register. If bit 21 is zero then the non-NAN operand is output when one of the inputs is NAN, in accordance with the IEEE Standard 754-2008. If bit 21 is one then the NAN input is propagated.

A NAN operand that is not propagated will generate a trap if flag bit 29 is set.

max_abs

format	opcode	operands
multi	25	all floating point types

float v0 = max_abs(v1, v2)

Gives the maximum of the absolute values of two floating point numbers.

max_abs(src1, src2) = max(abs(src1), abs(src2))

NAN values are treated in the same way as for the max instruction.

max_u

format	opcode	operands
multi	25	all integer types

uint32 r0 = max_u(r1, r2)

Gives the maximum of two unsigned integers.

max_u(src1,src2) = src1 > src2 ? src1 : src2

min

format	opcode	operands
multi	22	all types

int32 r0 = min(r1, r2)

float v0 = min(v1, v2)

Get the minimum of two numbers:

min(src1,src2) = src1 < src2 ? src1 : src2

Integer operands are treated as signed.

Floating point NAN values are treated in the same way as for the max instruction.

min_abs

format	opcode	operands
multi	23	all floating point types

float v0 = min_abs(v1, v2)

Gives the minimum of the absolute values of two floating point numbers.

min_abs(src1, src2) = min(abs(src1), abs(src2))

NAN values are treated in the same way as for the max instruction.

min_u

format	opcode	operands
multi	23	all integer types

uint32 r0 = min_u(r1, r2)

Gives the minimum of two unsigned integers.

min_u(src1,src2) = src1 < src2 ? src1 : src2

mul

format	opcode	operands
multi	11	all types
tiny	24	two single precision float vectors
tiny	25	two double precision float vectors
1.1 C	5	general purpose register and 16-bit sign-extended constant
1.3 C	60	single precision floating point vector and broadcast half-precision floating point constant. Optional
1.3 C	61	double precision floating point vector and broadcast half-precision floating point constant. Optional

int32 r0 = r1 * r2

float v0 *= 5.0

Multiplication.

The same instruction can be used for signed and unsigned integers.

mul_add, mul_add2

format	opcode	operands
multi	48	mul_add: dest = \pm src1 \pm (src2 \cdot src3). All types. Optional
multi	49	mul_add2: dest = \pm (src1 \cdot src2) \pm src3 . All types. Optional

Fused multiply and add.

The fused multiply-and-add instruction can often improve the performance of floating point code significantly. The intermediate product is calculated with extended precision according to the IEEE 754-2008 standard.

Only instruction formats that allow three operands are supported.

The signs of the operands can be inverted as indicated by the following 4-bit code

Table 5.8: Control bits for mul_add

E template formats	A template formats	Meaning
IM3 bit 0	mask bit 10	change sign of addend in even-numbered vector elements
IM3 bit 1	mask bit 11	change sign of addend in odd-numbered vector elements
IM3 bit 2	mask bit 12	change sign of product in even-numbered vector elements
IM3 bit 3	mask bit 13	change sign of product in odd-numbered vector elements

These option bits make it possible to do multiply-and-add, multiply-and-subtract, multiply-and-reverse-subtract, etc. It can also do multiply with alternating add and subtract, which is useful in calculations with complex numbers. There is no sign change if there is no IM3 field and no mask.

Support for integer operands is optional. Support for floating point operands is optional but desired.

mul_ex

format	opcode	operands
multi	14	integer vectors

int32 v0 = mul_ex(v1, v2)

Extended multiply, signed.

Multiply even-numbered signed integer vector elements to double size result. The result extends into the next odd-numbered vector element.

mul_ex_u

format	opcode	operands
multi	15	integer vectors

uint32 v0 = mul_ex_u(v1, v2)

Extended multiply, unsigned.

Multiply even-numbered unsigned integer vector elements to double size result. The result extends into the next odd-numbered vector element.

mul_hi

format	opcode	operands
multi	12	integer vectors

int32 r0 = mul_hi(r1, r2)

int32 v0 = mul_hi(v1, 2)

High part of signed integer product.

dest = (src1 · src2) >> OS

(Signed, OS = operand size in bits).

mul_hi_u

format	opcode	operands
multi	13	integer vectors

uint32 r0 = mul_hi_u(r1, r2)

High part of unsigned integer product.

dest = (src1 · src2) >> OS

(Unsigned, OS = operand size in bits).

rem

format	opcode	operands
multi	20	all types. Optional for vectors of more than one element

int32 r0 = r1 % r2

float v0 = rem(v1, v2)

Modulo.

Signed with integer operands or floating point operands.

A floating point number modulo zero gives NAN. An integer modulo zero gives zero.

rem_u

format	opcode	operands
multi	21	integers. Optional for vectors of more than one element

uint32 r0 = r1 % r2

Unsigned modulo or remainder.

An integer modulo zero gives zero.

round

format	opcode	operands
1.3 B	14	floating point vectors

float v0 = round(v1, 0)

Round floating point number to integer in floating point representation.

The rounding mode is specified in bit 0-1 of IM1. See table 3.16 page 22.

round_d2

format	opcode	operands
1.0 A	3	g.p. registers

Round unsigned integer down to nearest power of 2.

$dest = 1 \ll bitscan_r(src1)$.

The result is zero when `src1` is zero.

round_u2

format	opcode	operands
1.0 A	4	g.p. registers

Round unsigned integer up to nearest power of 2.

$dest = ((src1 \& (src1-1)) == 0) ? src1 : 1 \ll (bitscan_r(src1) + 1)$

The result is zero when `src1` is zero.

round2n

format	opcode	operands
1.3 B	15	vector registers. Optional

$float\ v0 = round2n(v1, -4)$

Round to nearest multiple of 2^n .

$dest = 2^n \cdot round(2^{-n} \cdot src1)$

`n` is a signed integer constant in IM1.

mul_2pow

format	opcode	operands
multi	32	all floating point types

Multiply by power of 2.

$dest = src1 * 2^{src2}$

`src1` and `dest` are floating point vectors, while `src2` is interpreted as a signed integer vector with the same element size as `src1` and `dest`.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

Overflow will produce infinity. The result will be zero rather than a subnormal number in case of underflow, regardless of bit 20 in the mask or numeric control register. The reason for this is that speed has priority here. This instruction will typically take a single clock cycle, while floating point multiplication by a power of 2 takes multiple clock cycles. This is useful for fast multiplication or division by a power of 2.

shift_left

format	opcode	operands
multi	32	all integer types

Shift integer left.

$dest = src1 \ll src2$

The result is zero if `src2` is out of range.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

shift_right_s

format	opcode	operands
multi	34	all integer types

Shift integer right with sign extension (arithmetic shift).

$dest = src1 \gg src2$

The result is 0 or -1 if $src2$ is out of range.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

shift_right_u

format	opcode	operands
multi	35	all integer types

Shift integer right with zero extension (logical shift).

$dest = src1 \gg src2$

The result is zero if $src2$ is out of range.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

shift_add

format	opcode	operands
1.8 B	1	g.p. registers

Shift left and add.

$dest = src1 + (src2 \ll src3)$.

$src1$ must use the same register as $dest$. $src3$ is an 8-bit unsigned immediate constant in IM1.

sqrt

format	opcode	operands
1.2 A	26	floating point vectors. Optional

Square root.

The square root of a negative number gives NAN.

sub

format	opcode	operands
multi	9	all types
tiny	3	g.p. register and 4-bit zero-extended constant
tiny	10	two g.p. registers
tiny	22	two vector registers, single precision float
tiny	23	two vector registers, double precision float
2.9	3	g.p. register and 32-bit zero-extended constant

int32 $r0 = r1 - r2$

int32 $r0 = r1 - 2$

int32+ $r0 -= 4$

int32+ r0--
float v0 = v1 - [r2 + 8, length = r5]

Subtraction.

sub_rev

format	opcode	operands
multi	10	all types

int32 r0 = 1 - r2
int32 v0 = - v2 + v1
float v0 = -v1 + [r2 + 8, length = r5]

Reverse subtraction.

dest = src2 - src1.

Arithmetic instructions with carry, overflow check, or saturation

abs

see page 55.

add_c

format	opcode	operands
1.2 A	28	integer vectors with two elements. Optional

Addition with carry.

The vector has two elements. The upper element of src1 is used as carry in. The upper element of dest is used as carry out. Only the lower element of src2 is used.

Longer vectors are not supported. See page 148 for an alternative for longer vectors.

add_oc

format	opcode	operands
1.2 A	38	vector registers. Optional

Integer addition with overflow check.

Instructions with overflow check use the even-numbered vector elements for arithmetic instructions. Each following odd-numbered vector element is used for overflow detection. If the first source operand is a scalar then the result operand will be a vector with two elements.

Overflow conditions are indicated with the following bits:

- bit 0. Unsigned integer overflow (carry).
- bit 1. Signed integer overflow.
- bit 2. Floating point overflow.
- bit 3. Floating point invalid operation.

The values are propagated so that the overflow result of the operation is OR'ed with the corresponding values of both input operands.

add_ss

format	opcode	operands
1.2 A	30	integer vectors. Optional

Add signed integers with saturation.

Overflow and underflow produces INT_MAX and INT_MIN.

add_us

format	opcode	operands
1.2 A	31	integer vectors. Optional

Add unsigned integers with saturation.
Overflow produces UINT_MAX.

compress_ss

format	opcode	operands
1.2 A	5	integer vectors. Optional

Compress, signed with saturation.

Same as compress (see page 45). Integers are treated as signed and compressed with saturation.
Floating point operands cannot be used.

compress_us

format	opcode	operands
1.2 A	6	integer vectors. Optional

Compress, unsigned with saturation.

Same as compress (see page 45). Integers are treated as unsigned and compressed with saturation.
Floating point operands cannot be used.

div_oc

format	opcode	operands
1.2 A	42	vector registers. Optional

Divide integers with overflow check.
See add_oc for options.

mul_oc

format	opcode	operands
1.2 A	41	vector registers. Optional

Multiply integers with overflow check.
See add_oc for options.

mul_ss

format	opcode	operands
1.2 A	34	integer vectors. Optional

Multiply signed integers with saturation.
Overflow and underflow produces INT_MAX and INT_MIN.

mul_us

format	opcode	operands
1.2 A	35	integer vectors. Optional

Multiply unsigned integers with saturation.
Overflow produces UINT_MAX.

shift_ss

format	opcode	operands
1.2 A	36	integer vectors. Optional

Shift signed integers left with saturation.

Overflow and underflow produces INT_MAX and INT_MIN.

shift_us

format	opcode	operands
1.2 A	37	integer vectors. Optional

Shift unsigned integers left with saturation.

Overflow produces UINT_MAX.

sub_b

format	opcode	operands
1.2 A	29	integer vectors with two elements. Optional

Subtraction with borrow.

The vector has two elements. The upper element of src1 is used as borrow in. The upper element of dest is used as borrow out. Only the lower element of src2 is used.

Longer vectors are not supported. See page 148 for an alternative for longer vectors.

sub_oc

format	opcode	operands
1.2 A	39	vector registers. Optional

Subtract integers with overflow check.

See add_oc for options.

sub_ss

format	opcode	operands
1.2 A	32	integer vectors. Optional

Subtract signed integers with saturation.

Overflow and underflow produces INT_MAX and INT_MIN.

sub_us

format	opcode	operands
1.2 A	33	integer vectors. Optional

Subtract unsigned integers with saturation.

Overflow and underflow produces UINT_MAX and 0.

Logic and bit manipulation instructions

and

format	opcode	operands
multi	28	all types
tiny	11	two g.p. registers
1.1 C	18	g.p. register and 8-bit signed constant shifted left by another constant
2.9	5	g.p. register and 32-bit constant shifted left by 32
1.3 C	34	vector of 16-bit integers, and broadcast 16-bit constant. Optional
1.3 C	44	vector of 32-bit integers, and broadcast sign-extended 8-bit constant shifted left by another constant. Optional
1.3 C	45	vector of 64-bit integers, and broadcast sign-extended 8-bit constant shifted left by another constant. Optional

`int32 r0 = r1 & r2`

`int32 v0 = v1 & 2`

Bitwise boolean and.

Floating point operands are treated as integers.

Do not use a floating point type with a constant operand unless you want the operand to be interpreted as floating point.

and_bit

format	opcode	operands
multi	39	all types

Clear all bits except one.

`dest = src1 & (1 << src2).`

Floating point operands are treated as integers.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

and_not

format	opcode	operands
multi	29	all types

Bitwise and not.

`dest = src1 & ~src2`

Floating point operands are treated as integers.

bit_reverse

format	opcode	operands
1.3 B	21	integer vectors. Optional

`int32 v0 = bit_reverse(v1)`

Reverse the order of bits in each element of a vector.

bits2bool

format	opcode	operands
1.2 A	14	integer vectors

int32 v0 = bits2bool(r1, v2)

Expand contiguous bits in a vector register to a boolean vector with each bit of the source going into bit 0 of each element of the destination. The remaining bits of each element are copied from the mask or numeric control register. The length in bytes of the result vector is specified by a general purpose register in RS.

bitscan_f

format	opcode	operands
1.0 A	1	general purpose registers
1.3 B	22	integer vectors. Optional for more than one element

int32 r0 = bitscan_f(r1)

int64 v0 = bitscan_f(v1)

Bit scan forward.

Find index to lowest set bit, i. e. highest X for which $((1 \ll X) - 1) \& \text{src1} == 0$.

The result is -1 when src1 is zero.

bitscan_r

format	opcode	operands
1.0 A	2	g.p. registers
1.3 B	23	integer vectors. Optional for more than one element

int32 r0 = bitscan_r(r1)

Bit scan reverse.

Find index to highest set bit, i. e. highest X for which $(1 \ll X) \leq \text{src1}$.

The result is -1 when src1 is zero.

bool_reduce

format	opcode	operands
1.2 A	16	integer vectors

int32 v0 = bool_reduce(r1, v2)

The boolean vector RT with length RS bytes is reduced by combining bit 0 of all elements. The output is a scalar integer where bit 0 is the AND combination of all the bits, and bit 1 is the OR combination of all the bits. The remaining bits are reserved for future use.

bool2bits

format	opcode	operands
1.2 A	15	integer vectors

int64 v0 = bool2bits(r1, v2)

The boolean vector RT with length RS bytes is packed into the lower n bits of RD, taking bit 0 of each element, where n is the number of elements ($n = RS / OS$). The length of RD is at least sufficient to contain n bits.

clear_bit

format	opcode	operands
multi	37	all types

Clear bit number src2 in src1.

$dest = src1 \& \sim(1 \ll src2).$

Floating point operands are treated as integers.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

compare

See page 57

fp_category

format	opcode	operands
1.3 B	17	floating point vectors

$float\ v0 = fp_category(v1, 1)$

The input is a floating point vector. The output is a boolean vector where bit 0 of each element indicates if the input RS belongs to any of the categories indicated by the bits in the immediate operand IM1. The remaining bits of the output are taken from the numeric control register. Any floating point value will belong to one, and only one, of these categories.

Table 5.9: Meaning of bits in fp_category

Bit number	Meaning
0	\pm NAN
1	\pm Zero
2	- Subnormal
3	+ Subnormal
4	- Normal
5	+ Normal
6	- Infinite
7	+ Infinite

make_mask

format	opcode	operands
2.6	2	integer vectors

$int32\ v0 = make_mask(v1, 2)$

Make a mask from the bits of the 32-bit integer constant src2. Each bit of src2 goes into bit 0 of one element of the output. The remaining bits of each element are taken from src1. The length of the output is the same as the length of src1. If there are more than 32 elements in the vector then the bit pattern of src2 is repeated.

make_sequence

format	opcode	operands
1.3 B	4	all vectors

$int32\ v0 = make_sequence(r1, 2)$

Makes a vector of length RS bytes. The number of elements is RS/OS, (OS = operand size).

The first element is equal to IM1, the next element is IM1+1, etc. Support for floating point is optional.

mask_length

format	opcode	operands
2.2.7	1.1	integer vectors

int64 v0 = mask_length(r1, v2, 0), options=2

Make a boolean vector to mask the first n bytes of a vector (RS = n).

The output vector RD will have the same length as the input vector RT. RS indicates the length of the part that is enabled by the mask.

IM3 contains the following option bits:

bit 0 = 0: bit 0 will be 1 in the first n bytes in the output and 0 in the rest.

bit 0 = 1: bit 0 will be 0 in the first n bytes in the output and 1 in the rest.

bit 1 = 1: copy remaining bits from input vector RT into each vector element.

bit 2 = 1: copy remaining bits from the numeric control register.

bit 4 = 1: broadcast remaining bits from IM2 into all 32-bit words of RD. Bit 1-7 of IM2 go to bit 1-7 of RD. Bit 8-11 of IM2 go to bit 18-21 of RD. Bit 12-15 of IM2 go to bit 26-29 of RD.

Output bits that are not set by any of these options will be zero. If multiple options are specified, the results will be OR'ed.

move_bits

format	opcode	operands
2.0.7	0.1	general purpose registers. Optional
2.2.7	0.1	integer vectors. Optional

int32 v0 = move_bits(v1, v2, 3, 4, 5)

Take one or more contiguous bits from one position in the second source operand and insert them into another position the first source operand. This can be used for all kinds of bitfield operations.

The position in src2 is the lower 8 bits of IM2. a = IM2 & 0xFF.

The position in src1 is the upper 8 bits of IM2. b = IM2 >> 0xFF.

The number of bits to move is c = IM3.

m = (1 << c) - 1

dest = (src1 & ~(m<<b)) | ((src2 >> a) & m) << b

or

format	opcode	operands
multi	30	all types
tiny	12	two g.p. registers
1.1 C	19	g.p. register and 8-bit signed constant shifted left by another constant
2.9	6	g.p. register and 32-bit constant shifted left by 32
1.3 C	35	vector of 16-bit integers, and broadcast 16-bit constant. Optional
1.3 C	46	vector of 32-bit integers, and broadcast sign-extended 8-bit constant shifted left by another constant. Optional
1.3 C	47	vector of 64-bit integers, and broadcast sign-extended 8-bit constant shifted left by another constant. Optional

int32 r0 = r1 | r2

int32 v0 = v1 | 2

Bitwise boolean or.

Floating point operands are treated as integers.

Do not use a floating point type with a constant operand unless you want the operand to be interpreted as floating point.

popcount

format	opcode	operands
1.3 B	24	integer vectors. Optional for more than one element

`int32 v0 = popcount(v1)`

The popcount instruction counts the number of 1-bits in an integer. It can also be used for parity generation.

replace_bits

format	opcode	operands
2.9	9	general purpose registers. Optional
2.6	9	integer vectors. Optional

`int32 v0 = replace_bits(v1, 2, 3, 4)`

Replace one or more contiguous bits in the first source operand with a constant.

The second source operand is the replacement bits, with a maximum size of 16 bits. (Any additional bits will be zero)

The third source operand is the position of the bits to replace.

The fourth source operand is the number of bits to replace.

Algorithm:

`val = IM2 & 0xFFFF`

`pos = (IM2 >> 16) & 0xFF`

`num = (IM2 >> 24) & 0xFF`

`mask = (1 << num) - 1`

`dest = (src1 & ~(mask<<pos)) | ((val & mask) << pos)`

rotate

format	opcode	operands
multi	33	all integer types

Rotate the bits of src1 left if src2 is positive, or right if src2 is negative.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

set_bit

format	opcode	operands
multi	36	all integer types

Set bit number src2 in src1 to one.

`dest = src1 | (1 << src2)`

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

test_bit

format	opcode	operands
multi	40	all integer types

Test the value of bit number src2 in src1, and make it the least significant bit of the output (to use as a boolean). The result is zero if src2 is out of range.

result = (src1 >> src2) & 1.

The remaining bits of the output are taken from a mask or from the numeric control register if there is no mask.

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

Masking and fallback is possible. Alternative use of the mask register and the fallback register as extra boolean operands is supported with E template formats with no memory operand. These options are controlled by bit 0-1 of the constant IM3 in the same way as for the compare instruction (see page 58). Bit 2 inverts the result, bit 3 inverts the fallback, and bit 4 inverts the mask. These options are summarized in the following table, giving the value of bit 0 of the destination register. These options cannot be used if there is a memory operand.

Table 5.10: Alternative use of mask and fallback register controlled by IM3

bit 4	bit 3	bit 2	bit 1	bit 0	Output
0	0	0	0	0	mask ? result : fallback
0	0	1	0	0	mask ? !result : fallback
0	1	0	0	0	mask ? result : !fallback
0	1	1	0	0	mask ? !result : !fallback
1	0	0	0	0	!mask ? result : fallback
1	0	1	0	0	!mask ? !result : fallback
1	1	0	0	0	!mask ? result : !fallback
1	1	1	0	0	!mask ? !result : !fallback
0	0	0	0	1	mask & result & fallback
0	0	1	0	1	mask & !result & fallback
0	1	0	0	1	mask & result & !fallback
0	1	1	0	1	mask & !result & !fallback
1	0	0	0	1	!mask & result & fallback
1	0	1	0	1	!mask & !result & fallback
1	1	0	0	1	!mask & result & !fallback
1	1	1	0	1	!mask & !result & !fallback
0	0	0	1	0	mask & (result fallback)
0	0	1	1	0	mask & (!result fallback)
0	1	0	1	0	mask & (result !fallback)
0	1	1	1	0	mask & (!result !fallback)
1	0	0	1	0	!mask & (result fallback)
1	0	1	1	0	!mask & (!result fallback)
1	1	0	1	0	!mask & (result !fallback)
1	1	1	1	0	!mask & (!result !fallback)
0	0	0	1	1	mask & (result ^ fallback)
0	0	1	1	1	mask & (!result ^ fallback)
0	1	0	1	1	mask & (result ^ !fallback)
0	1	1	1	1	mask & (!result ^ !fallback)
1	0	0	1	1	!mask & (result ^ fallback)
1	0	1	1	1	!mask & (!result ^ fallback)

1	1	0	1	1	!mask & (result ^ !fallback)
1	1	1	1	1	!mask & (!result ^ !fallback)

test_bits

format	opcode	operands
multi	41	all integer types

Test if at least one of the indicated bits is 1. The boolean result is placed in the least significant bit of the output. The remaining bits of the output are taken from a mask or from the numeric control register if there is no mask.

result = ((src1 & src2) != 0)

Masking and fallback is possible. Alternative use of the mask register and the fallback register as extra boolean operands is supported with E template formats with no memory operand. These options are controlled by the bits of IM3 in the same way as for test_bit, as indicated in table 5.10.

test_bits_all1

format	opcode	operands
multi	42	all integer types

Test if all the indicated bits are 1. The boolean result is placed in the least significant bit of the output. The remaining bits of the output are taken from a mask or from the numeric control register if there is no mask.

result = ((src1 & src2) == src2)

Masking and fallback is possible. Alternative use of the mask register and the fallback register as extra boolean operands is supported with E template formats with no memory operand. These options are controlled by the bits of IM3 in the same way as for test_bit, as indicated in table 5.10.

toggle_bit

format	opcode	operands
multi	38	all types

Change the value of bit number src2 in src1 to its opposite.

dest = src1 ^ (1 << src2)

This instruction can have a memory operand and an immediate operand, using an E template format and IM3 as the immediate operand. It is not possible to have both a register operand and a memory operand with these formats.

truth_tab2

format	opcode	operands
1.3 B	25	integer vectors

int32 v0 = truth_tab2(v0, v1, 8)

This instruction can make an arbitrary boolean function of two boolean vector input variables, expressed by a 4-bit truth table in an immediate constant IM1. The result in bit 0 of each vector element is the arbitrary boolean function of bit 0 of the corresponding elements of each of the input operands. Bit 0 of the result is a bit from the truth table selected by the combined input bits. The remaining bits of the output vector are copied from the first input operand.

The first source operand must be the same as the destination operand.

This instruction can be used as a universal instruction for manipulating and combining boolean vectors and masks.

$$\text{result} = ((\text{IM1} \gg ((\text{src1} \& 1) | (\text{src2} \& 1) \ll 1)) \& 1) | (\text{src1} \& \sim 1)$$

truth_tab3

format	opcode	operands
2.2.7	2.1	integer vectors. Optional

int32 v0 = truth_tab3(v1, v2, v3, 9), options=0

This instruction can make an arbitrary boolean function of three boolean vector input variables, expressed by an 8-bit truth table in an immediate constant IM2. The result in bit 0 of each vector element is the arbitrary boolean function of bit 0 of the corresponding elements of each of the input operands. Bit 0 of the result is a bit from the truth table selected by the combined input bits. The remaining bits of the output vector are copied from the mask, or from the first input operand if there is no mask.

This instruction can be used as a universal instruction for manipulating and combining boolean vectors and masks.

Algorithm:

$$\text{select} = (\text{src1} \& 1) | (\text{src2} \& 1) \ll 1 | (\text{src3} \& 1) \ll 2$$

$$\text{dest} = (\text{IM2} \gg \text{select}) \& 1$$

The remaining bits of dest are determined by bit 1-4 of the option bits as described below.

IM3 contains the following option bits:

bit 0 = 0: Normal use of mask with src1 as fallback.

bit 0 = 1: The normal function of the mask is disabled. Instead, bit 0 of the result is AND'ed with the corresponding mask bit.

bit 1 = 1: copy remaining bits from input vector src1.

bit 2 = 1: copy remaining bits from the numeric control register.

bit 3 = 1: copy remaining bits from mask register.

bit 4 = 1: broadcast remaining bits from IM2 into all 32-bit words of RD. Bit 8-11 of IM2 go to bit 18-21 of dest. Bit 12-15 of IM2 go to bit 26-29 of RD.

Output bits that are not set by any of these options will be zero. If multiple options are specified, the results will be OR'ed.

xor

format	opcode	operands
multi	31	all types
tiny	13	two g.p. registers
1.1 C	20	g.p. register and 8-bit signed constant shifted left by another constant
2.9	7	g.p. register and 32-bit constant shifted left by 32
1.3 C	36	vector of 16-bit integers, and broadcast 16-bit constant. Optional
1.3 C	48	vector of 32-bit integers, and broadcast sign-extended 8-bit constant shifted left by another constant. Optional
1.3 C	49	vector of 64-bit integers, and broadcast sign-extended 8-bit constant shifted left by another constant. Optional

int32 r0 = r1 ^ r2

int32 v0 = v1 ^ 2

Bitwise boolean exclusive or.

Floating point operands are treated as integers.

Do not use a floating point type with a constant operand unless you want the operand to be interpreted as floating point.

Combined ALU and branch instructions

These instructions are doing an arithmetic or logic operation and a conditional jump depending on the result. Each instruction can be coded in a number of different formats described on page 23.

The instructions are listed below in pairs, where the second instruction has the branch condition inverted.

These instructions cannot have a mask, and they do not depend on the numeric control register. Numeric overflow cannot generate a trap in these instructions.

Addition and subtraction instructions here do not support floating point operands because the longer latency will not fit into the pipeline structure.

add/jump_zero

format	opcode	instruction	operands
all	16	add/jump_zero	integer
all	17	add/jump_nzero	integer

Add two integer operands and jump if the result is zero.

add/jump_neg

format	opcode	instruction	operands
all	18	add/jump_neg	integer
all	19	add/jump_nneg	integer

Add two integer operands and jump if the signed result is negative.

The result will wrap around in the case of overflow and jump if the result has the sign bit set.

add/jump_pos

format	opcode	instruction	operands
all	20	add/jump_pos	integer
all	21	add/jump_npos	integer

Add two integer operands and jump if the signed result is positive.

The result will wrap around in the case of overflow and jump if the result is not zero and does not have the sign bit set.

add/jump_overfl

format	opcode	instruction	operands
all	22	add/jump_overfl	integer
all	23	add/jump_noverfl	integer

Add two signed integer operands and jump if the result overflows.

add/jump_carry

format	opcode	instruction	operands
all	24	add/jump_carry	integer
all	25	add/jump_ncarry	integer

Add two unsigned integer operands and jump if the operation produces a carry.

add/jump

format	opcode	instruction	operands
all	55	add/jump	integer

Add two integer operands and jump unconditionally.

increment_compare/jump_above

format	opcode	instruction	operands
all	50	increment_compare/jump_above	integer
all	51	increment_compare/jump_beloweq	integer

Add 1 to the first source operand and jump if the signed result is bigger than the second source operand. The result is saved in the destination operand. This is useful for implementing a simple “for” loop.

The result will wrap around from INT_MAX to INT_MIN in case of overflow, and jump only if the second source operand is also INT_MAX.

The hardware implementation can use the following algorithm to avoid problems with overflow and to minimize the latency:

dest = src1 + 1

jump if src1 ≥ src2

Note that a “for” loop using this instruction will stop at n, where n = src2. If it is desired to stop at n-1 then subtract 1 from n first. (An implementation that stops at n-1 would be more difficult to implement with a short latency).

sub_maxlen/jump_pos

format	opcode	instruction	operands
1.5 2.5.2	52	sub_maxlen/jump_pos	int64 in g. p. register
1.5 2.5.2	53	sub_maxlen/jump_npos	int64 in g. p. register

Subtract the maximum vector length (in bytes) from a general purpose register and jump if the result is positive. The 8-bit immediate operand indicates the operand type for which the maximum vector length is obtained.

The register operand must be a 64-bit general purpose register.

This instruction makes it easy to implement the type of vector loop described on on page 8.

sub/jump_zero

format	opcode	instruction	operands
1.4, 2.5.0	0	sub/jump_zero	integer
1.4, 2.5.0	1	sub/jump_nzero	integer

Subtract two integer operands and jump if the result is zero.

Immediate constants are not allowed, and the corresponding format is used for other purposes. If a constant operand is needed then change the sign of the operand and use add/jump_zero.

sub/jump_neg

format	opcode	instruction	operands
1.4, 2.5.0	2	sub/jump_neg	integer
1.4, 2.5.0	3	sub/jump_nneg	integer

Subtract two integer operands and jump if the signed result is negative.

The result will wrap around in the case of overflow and jump if the result has the sign bit set.

Immediate constants are not allowed, and the corresponding format is used for other purposes. If a constant operand is needed then change the sign of the operand and use add/jump_neg.

sub/jump_pos

format	opcode	instruction	operands
1.4, 2.5.0	4	sub/jump_pos	integer
1.4, 2.5.0	5	sub/jump_npos	integer

Subtract two integer operands and jump if the signed result is positive. The result will wrap around in the case of overflow and jump if the result is not zero and does not have the sign bit set.

Immediate constants are not allowed, and the corresponding format is used for other purposes. If a constant operand is needed then change the sign of the operand and use add/jump_pos.

sub/jump_overfl

format	opcode	instruction	operands
1.4, 2.5.0	6	sub/jump_overfl	integer
1.4, 2.5.0	7	sub/jump_noverfl	integer

Subtract two signed integer operands and jump if the result overflows.

Immediate constants are not allowed, and the corresponding format is used for other purposes. If a constant operand is needed then change the sign of the operand and use add/jump_overfl.

sub/jump_borrow

format	opcode	instruction	operands
1.4, 2.5.0	8	sub/jump_borrow	integer
1.4, 2.5.0	9	sub/jump_nborrow	integer

Subtract two unsigned integer operands and jump if the operation produces a borrow.

Immediate constants are not allowed, and the corresponding format is used for other purposes. If a constant operand is needed then change the sign of the operand and use add/jump_nborrow.

sub/jump

format	opcode	instruction	operands
all	54	sub/jump	integer

Subtract two integer operands and jump unconditionally.

compare/jump_equal

format	opcode	instruction	operands
all	32	compare/jump_equal	integer or floating point
all	33	compare/jump_nequal	integer or floating point

Compare two integer or floating point operands and jump if they are equal. The destination operand (if any) will be unchanged.

Two floating point NAN operands will be treated as not equal, even if they are identical. 0.0 and -0.0 are treated as equal.

compare/jump_sbelow

format	opcode	instruction	operands
all	34	compare/jump_sbelow	integer or floating point
all	35	compare/jump_saboveeq	integer or floating point

Compare two signed integer or floating point operands and jump if the first source operand is less than the second source operand.

The destination operand (if any) will be unchanged.

Overflow cannot occur.

0.0 and -0.0 are treated as equal.

compare/jump_sbelow will jump and compare/jump_saboveeq will not jump if one or both operands are NAN.

compare/jump_sabove

format	opcode	instruction	operands
all	36	compare/jump_sabove	integer or floating point
all	37	compare/jump_sbeloweq	integer or floating point

Compare two signed integer or floating point operands and jump if the first source operand is bigger than the second source operand.

The destination operand (if any) will be unchanged.

Overflow cannot occur.

0.0 and -0.0 are treated as equal.

compare/jump_sabove will jump and compare/jump_sbeloweq will not jump if one or both operands are NAN.

compare/jump_ubelow

format	opcode	instruction	operands
all	38	compare/jump_ubelow	integer
all	39	compare/jump_uaboveeq	integer

Compare two unsigned integer operands and jump if the first source operand is less than the second source operand.

The destination operand (if any) will be unchanged.

Overflow cannot occur.

compare/jump_uabove

format	opcode	instruction	operands
all	40	compare/jump_uabove	integer
all	41	compare/jump_ubeloweq	integer

Compare two unsigned integer operands and jump if the first source operand is bigger than the second source operand.

The destination operand (if any) will be unchanged.

Overflow cannot occur.

compare/jump_absbelow

format	opcode	instruction	operands
all	38	compare/jump_absbelow	floating point
all	39	compare/jump_absaboveeq	floating point

Compare two floating point operands and jump if the absolute value of the first source operand is less than the the absolute value of the second source operand.

The destination operand (if any) will be unchanged.

compare/jump_absbelow will jump and compare/jump_absaboveeq will not jump if one or both operands are NAN.

compare/jump_absabove

format	opcode	instruction	operands
all	38	compare/jump_absabove	floating point
all	39	compare/jump_absbeloweq	floating point

Compare two floating point operands and jump if the absolute value of the first source operand is bigger than the the absolute value of the second source operand.

The destination operand (if any) will be unchanged.

compare/jump_absbelow will jump and compare/jump_absaboveeq will not jump if one or both operands are NAN.

compare/jump_nfinite

format	opcode	instruction	operands
all	24	compare/jump_nfinite	floating point
all	25	compare/jump_finite	floating point

Jump if not finite, i.e. if at least one of the source operands is \pm infinity or NAN.

The destination operand (if any) will be unchanged.

test_bit/jump_zero

format	opcode	instruction	operands
all	42	test_bit/jump_zero	all
all	43	test_bit/jump_nzero	all

Test a single bit in the first source operand as indicated by the an index in the second source operand and jump if the indicated bit is 0.

All operands are treated as unsigned integers.

Floating point operands are treated as unsigned integer scalars in vector registers.

jump if $((src1 \gg src2) \& 1) == 0$

The destination operand (if any) will be unchanged.

test/jump_all1

format	opcode	instruction	operands
all	44	test/jump_all1	all
all	45	test/jump_nall1	all

Bitwise test. Jump if the indicated bits are all 1.

jump if $(src1 \& src2) == src2$

All operands are treated as integers. Floating point operands are treated as unsigned integer scalars in vector registers.

The destination operand (if any) will be unchanged.

and/jump_zero

format	opcode	instruction	operands
all	26	and/jump_zero	all
all	27	and/jump_nzero	all

Bitwise and. Jump if zero.

dest = src1 & src2
 jump if dest == 0

All operands are treated as integers. Floating point operands are treated as unsigned integer scalars in vector registers.

or/jump_zero

format	opcode	instruction	operands
all	28	or/jump_zero	all
all	29	or/jump_nzero	all

Bitwise or. Jump if zero.

dest = src1 | src2
 jump if dest == 0

All operands are treated as integers. Floating point operands are treated as unsigned integer scalars in vector registers.

xor/jump_zero

format	opcode	instruction	operands
all	30	xor/jump_zero	all
all	31	xor/jump_nzero	all

Bitwise exclusive or. Jump if zero.

dest = src1 ^ src2
 jump if dest == 0

All operands are treated as integers. Floating point operands are treated as unsigned integer scalars in vector registers.

shift_left/jump_zero

format	opcode	instruction	operands
not 1.5	10	shift_left/jump_zero	all
not 1.5	11	shift_left/jump_nzero	all

Shift the first source operand left by the count of the second source operand. Jump if the result is zero.

dest = src1 << src2
 Jump if dest == 0

Floating point operands are treated as scalar integer operands in vector registers.

shift_right_u/jump_zero

format	opcode	instruction	operands
not 1.5	12	shift_right_u/jump_zero	all
not 1.5	13	shift_right_u/jump_nzero	all

Shift the first source operand right by the count of the second source operand with zero extension (logical shift). Jump if the result is zero.

dest = src1 << src2

Jump if dest == 0

Floating point operands are treated as scalar integer operands in vector registers.

rotate/jump_carry

format	opcode	instruction	operands
not 1.5	14	rotate/jump_carry	all
not 1.5	15	rotate/jump_ncarry	all

Rotate the bits of the first source operand by the count of the second source operand, left if positive or right if negative. Jump if the carry bit is 1.

The carry bit is the least significant bit of the result if the shift count is positive or zero, and the most significant bit of the result if the shift count is negative.

Floating point operands are treated as scalar integer operands in vector registers.

This instruction is useful for generating an arbitrary sequence of jump and not jump.

Unconditional and indirect jump, call, and return instructions

Control transfer instructions are available in a number of different formats, described on page 23.

Direct jump

format	opcode	operands
1.5 C	58	jump with 16 bit relative address
1.5 D	0	jump with 24 bit relative address
2.5.2 C	58	jump with 32 bit relative address
3.1.0 B	58	jump with 64 bit absolute address

Unconditional jump.

Direct function call

format	opcode	operands
1.5 C	59	call with 16 bit relative address
1.5 D	8	call with 24 bit relative address
2.5.2 C	59	call with 32 bit relative address
3.1.0 B	59	call with 64 bit absolute address

Function call.

The return address is stored on the call stack. The calling conventions are described in chapter 12.4.

Indirect jump

format	opcode	operands
1.4 B	58	64 bit absolute address in memory operand with 8 bit offset
2.5.0 B	58	64 bit absolute address in memory operand with 32 bit offset
1.5 C	60	64 bit absolute address in register
1.4 A	60	Multi-way jump with table of relative addresses (see below)

Indirect call

format	opcode	operands
1.4 B	59	64 bit absolute address in memory operand with 8 bit offset
2.5.0 B	59	64 bit absolute address in memory operand with 32 bit offset
1.5 C	61	64 bit absolute address in register
1.4 A	61	Multi-way call with table of relative addresses (see below)

Multi-way jump and call

format	opcode	operands
1.4 A	60	Jump with table of relative addresses. Has reference point, base and scaled index
2.5.0 B	60	Jump with table of relative addresses. Has reference point, base and offset
1.4 A	61	Call with table of relative addresses. Has reference point, base and scaled index
2.5.0 B	61	Call with table of relative addresses. Has reference point, base and offset

The table-based indirect jump and call instructions are intended to facilitate multiway branches (switch/case statements), function tables in code interpreters, and virtual function tables in object oriented languages with polymorphism.

The table of jump or call addresses is stored in memory as signed offsets relative to an arbitrary reference point, which may be the table address, the code base, or any other reference point. The operand type specifies the size of the table entries. The use of relative addresses makes the table more compact than if 64-bit absolute addresses were used. Masks are not supported.

This instruction works as follows. Calculate the address of a table entry as the base pointer (RT) plus the offset (unscaled) or the index (RS) multiplied by the operand size. Read a signed value from this address, sign-extend to 64 bits, and scale by 4. Then add the reference point (RD). Jump or call to the calculated address. The array index (RS) is scaled by the operand size, while the table entries are scaled by the instruction word size (4). The reference point must be aligned by 4.

The table used by the table-based jump and call instructions is preferably placed in the constant data section (CONST). This makes it possible to use the table base as reference point. This also improves security by giving read-only access to the table.

return

format	opcode	operands
1.4 B	62	

Return from function call. The return address is taken from the call stack.

Return instructions do not need a stack offset when the calling conventions specified in chapter 12.4 are used.

System call, system return, and traps

See page 88.

Miscellaneous instructions

address

format	opcode	operands
2.9	32	g.p. registers

Calculate an address relative to a pointer by adding a 32-bit sign-extended constant to a special pointer register. The pointer register can be THREADP (28), DATAP (29), IP (30) or SP(31).

compare_swap

format	opcode	operands
2.5 B	18	g. p. registers and memory operand with 32 bit offset. Optional

int32 compare_swap(r1, r2, [r3+0x100])

Atomic compare and swap instruction, used for thread synchronization and for lock-free data sharing between threads. src1 and src2 are register operands, src3 is a memory operand, which must be aligned to a natural address. All operands are treated as integers, regardless of the specified operand type. The operation is:

```
temp = src3;
if (temp == src1) src3 = src2;
return temp;
```

Further atomic instructions can be implemented if needed, preferably with the same format and consecutive values of OP1.

filler

format	opcode	operands
1.5 C	63	

This instruction is used for filling unused code memory. It will generate a trap (interrupt) if executed.

All fields are filled with ones. The complete instruction code word is 0x6FFFFFFF.

nop

format	opcode	operands
multi	0	
tiny	0	
3.0	0	

No operation. Used as a filler to replace removed code or to align code entries.

The processor is allowed to skip NOPs as fast as it can at an early stage in the pipeline. A pair of tiny instructions where the second instruction is a NOP can be treated as a single instruction.

These NOPs cannot be used as timing delays, only as fillers.

undef

format	opcode	operands
multi	63	

Undefined code. Guaranteed to generate trap (interrupt) in all future implementations

userdef

format	opcode	operands
multi	55-62	any types

Reserved for user-defined instructions.

System instructions

input

format	opcode	operands
1.2 A	62	vector registers

int64 v0 = input(r1, r2)

Read from input port. Privileged instruction.

RD = destination (vector register), RS = vector length in bytes, RT = port address.

output

format	opcode	operands
1.2 A	63	vector registers

int64 output(v0, r1, r2)

Write to output port. Privileged instruction.

RD = source (vector register), RS = vector length in bytes, RT = port address.

read_capabilities

format	opcode	operands
1.8 B	34	g.p. register, capabilities register

Read processor capabilities register. These registers are used for indicating capabilities of the processor, such as support for optional instructions and limitations to vector lengths. The size is 64 bits. These registers are initialized with their default values at program start.

The immediate constant in IM1 determines details of the operation:

Table 5.11: Meaning of immediate constant in read_capabilities and write_capabilities instructions

Bit number	Meaning
0	0: read/write the capabilities for the operand type specified in bit 5-7. 1: read the typical capabilities for all operand types / write the capabilities for all relevant operand types.
1	0: read the current value of the register, which may have been modified. 1: read the real capabilities of the hardware (cannot write.)
5-7	Operand type for capabilities.

Table 5.12: List of capabilities registers

Capabilities register number	Meaning
capab0	Maximum vector length for general instructions.
capab1	Maximum vector length for permute instructions.
capab2	Maximum block size for permute instructions.
capab3	Maximum vector length for compress_sparse and expand_sparse.
capab8	Support for optional instructions in general purpose registers. Each bit indicates a specific instruction.
capab9	Support for optional instructions on scalars in vector registers. Each bit indicates a specific instruction.
capab10	Support for optional instructions on vectors. Each bit indicates a specific instruction.

Changing the values of the maximum vector length has the following effects. If the maximum length is reduced below the physical capability then any attempt to make a longer vector will result in the reduced length. The behavior of vector registers that already had a longer length before the maximum length was reduced, is implementation dependent. If the maximum vector length is set to a higher value than the physical capability then any attempt to make a vector longer than the physical capability will cause a trap to facilitate emulation. Capabilities registers 0-3 can be increased for the purpose of emulation. The value of capabilities registers 0-3 must be powers of 2.

Capabilities registers 8-9 can be modified for test purposes or to tell the software not to use a specific instruction. The same value will be returned when reading the register. Attempts to execute an instruction that is not supported will cause a trap, regardless of the value of the capabilities register.

read_memory_map

format	opcode	operands
1.2 A	60	memory map, vector register

int64 v0 = read_memory_map(r2, r3)

Read memory map and save it to a vector register. Privileged instruction.

RD = destination vector register, RT-RS = internal address.

read_call_stack

format	opcode	operands
1.2 A	58	internal call stack, vector register

int64 v0 = read_call_stack(r1, r2)

Read the internal call stack into a vector register. This instruction is used for saving the internal call stack to system memory in case of overflow. Privileged instruction.

RD = destination vector register, RT-RS = internal address.

read_perf

format	opcode	operands
1.8 B	36	g.p. register, performance counter register

Read performance counter register. A number of internal registers are used for counting performance related events. The details are not defined yet.

This instruction has an integer constant as second source operand. This is reserved for future purposes and must be 0.

read_perfs

format	opcode	operands
1.8 B	37	g.p. register, performance counter register

This is the same as the read_perf instruction, but serializing. This means that no instruction can execute out of order with read_perfs.

This instruction has an integer constant as second source operand. This is reserved for future purposes and must be 0.

read_spec

format	opcode	operands
1.8 B	32	g.p. register, special register

int64 r0 = read_spec(spec1, 2)

Read a special system register. The following special registers are currently defined. The size is 64 bits. These registers are initialized with their default values at program start.

This instruction has an integer constant as second source operand. This is reserved for future purposes and must be 0.

Table 5.13: List of special registers

Special register number	Meaning
spec0	Numeric control register (NUMCONTR)
spec1	Microprocessor brand ID
spec2	Microprocessor version number
spec28	Thread environment block pointer (THREADP)
spec29	Data section pointer (DATAP)

read_spev

format	opcode	operands
1.3 B	2	vector register, special register

int64 v0 = read_spev(r1, 2)

Read special register IM1 into vector register RD with length RS bytes. The value is broadcast if the vector register is longer than the special register.

The following special registers are currently defined:

Table 5.14: Special registers that can be read into vectors

Special register number	Meaning
spec0	Numeric control register (NUMCONTR). The value is broadcast into all elements of the destination register with the indicated operand size and length.
spec1	Name of processor. The output is a zero-terminated UTF-8 string containing the brandname and model name of the microprocessor.

read_sys

format	opcode	operands
1.8 B	38	g.p. register, system register

Read system register. Details are not defined yet. This instruction is privileged.

This instruction has an integer constant as second source operand. This is reserved for future purposes and must be 0.

sys_call

System calls use ID numbers rather than addresses to identify system functions. The ID is the combination of a module ID identifying a particular system module or device driver and a function ID identifying a particular function within this module. The module ID and the function ID are both 16 or 32 bits, so that the combined system call ID is up to 64 bits. The sys_call instruction has the following variants:

Table 5.15: Variants of system call instruction

Format	Operand type	Function ID	Module ID
1.4 A	32 bit	RT bit 0-15	RT bit 16-31
1.4 A	64 bit	RT bit 0-31	RT bit 32-63
2.5.1 B	32 bit	IM2 bit 0-15	IM2 bit 16-31
2.5.4 C	64 bit	IM1,IM2 bit 0-15	IM3 bit 0-31
3.1.0 B	64 bit	IM2 bit 0-31	IM3 bit 0-31

The sys_call instruction can indicate a block of memory to be shared with the system function. The address of the memory block is pointed to by the register specified in RD and the length is in register RS. This memory block, which the caller must have access rights to, is shared with the system function. The system function will get the same access rights to this block as the calling thread has, i.e. read access and/or write access. This is useful for fast transfer of data between the caller and the system function. No other memory is accessible to both the caller and the called function. If the RD and RS fields are both zero (i.e. indicating register r0) then no memory block is shared. The sys_call instruction in format 2.5.4 cannot have a shared memory block. System calls cannot have a mask.

Parameters for system functions are transferred in registers, following the same calling conventions as normal functions. The registers used for function parameters are usually different from the registers in the RD, RS and RT fields. Function parameters that do not fit into registers must reside in the shared memory block.

sys_return

format	opcode	operands
1.5 C	62	

Return from system call.

trap

format	opcode	instruction	immediate operand
1.5 C	63	trap	0-254
1.5 C	63	filler	255

Traps work like interrupts. The unconditional trap has an 8-bit interrupt number in IM1. This is an index into the interrupt vector table, which initially starts at absolute address zero. The unconditional trap instruction may use IM2 for additional information.

A trap instruction with all 1's in all fields (opcode 0x6FFFFFFF) can be used as filler in unused parts of code memory.

conditional trap

format	opcode	instruction	immediate operand
2.5.3 C	63	compare, trap_uabove	40
2.5.3 C	63	conditional trap	IM1 = interrupt number, IM2 = OPJ

The conditional trap generates a trap if the specified condition is true.

IM1 contains the interrupt number.

IM2 contains the condition code OPJ, specified in table 3.18.

Currently, only the condition code 40 is supported, corresponding to compare unsigned and jump if above. This will generate a trap if $RD > IM3$. This is useful for checking if an array index exceeds the upper bound. The lower bound does not have to be checked because we use unsigned compare.

write_capabilities

format	opcode	operands
1.8 B	35	g.p. register, capabilities register

Write processor capabilities register. See the read_capabilities instruction, page 85, for details.

write_memory_map

format	opcode	operands
1.2 A	61	memory map, vector register

int64 write_memory_map(v1, r2, r3)

Write a vector register to memory map. RD = vector register source. RT-RS = internal address. Privileged instruction.

write_call_stack

format	opcode	operands
1.2 A	59	memory map, vector register

int64 write_call_stack(v1, r2, r3)

Write a vector register to the internal call stack. This instruction is used for restoring the internal call stack. Privileged instruction.

write_spec

format	opcode	operands
1.8 B	33	g.p. register, special register

Write special register. See read_spec instruction page 87 for details.

write_sys

format	opcode	operands
1.8 B	39	g.p. register, system register

Write system register. Details are not defined yet. This instruction is privileged.

5.1 Common operations that have no dedicated instruction

This section discusses some common operations that are not implemented as single instructions, and how to code these operations in software.

Change sign

For integer operands, do a reverse subtract from zero. For floating point operands, use the `toggle_bit` instruction on the sign bit.

Not

To invert all bits in an integer, do an XOR with -1. To invert a Boolean, do an XOR with 1.

Rotate through carry

Rotates through carry are rarely used, and common implementations can be very inefficient. A left rotate through carry can be replaced by an `add_c` with the same register in both source operands.

Push and pop registers

There are no push and pop instructions, but pseudo-instructions named `push` and `pop` are available for saving and restoring registers on the stack. For example, pushing and popping a general purpose register `r1` is implemented as follows.

```
// push (r1)
int64 sp -= 8
int64 [sp] = r1
// pop (r1)
int64 r1 = [sp]
int64 sp += 8
```

Vector registers can be saved with the `sub_cps` and `save_cp` instructions and restored with the `restore_cp` and `add_cps` instructions. For example, pushing and popping a vector register `v1` is implemented as follows.

```
// push (v1)
int64 sp = sub_cps(sp, v1)
[sp] = save_cp(v1)
// pop (v1)
v1 = restore_cp([sp])
int64 sp = add_cps(sp, v1)
```

See page 53 for details about these instructions

Save and restore all registers

The registers need to be saved and restored at task switches. There are no instructions to save all registers because we do not want complex instructions. The saving and restoring of registers is described under `push` and `pop` on page 90.

Horizontal vector add

See example 14.1.

5.2 Unused instructions

Unused instructions and opcodes can be divided into three types:

1. The opcode is reserved for future use. Attempts to execute it will trigger a trap (synchronous interrupt) which can be used for generating an error message or for emulating instructions that are not supported.
2. The opcode is guaranteed to generate a trap, not only in the present version, but also in all future versions. This can be used as a filler in unused parts of the memory or for indicating unrecoverable errors. It can also be used for emulating user-specific instructions.
3. The error is ignored and does not trigger a trap. It can be used for future extensions that improve performance or functionality, but which can be safely ignored when not supported.

All three types are implemented, where type 1 is the most common.

Nop instructions with nonzero values in unused fields are type 3. These instructions are ignored.

Prefetch and fence instructions with no memory operand, with nonzero values in unused fields, or with undefined values in IM3 are type 3. These instructions are ignored.

Unused bits in masks and numeric control register are type 3. These bits are ignored.

Trap instructions and conditional trap instructions with nonzero values in unused fields or undefined values in any field are type 2. These instructions are guaranteed to generate a trap. A special version of the trap instruction is intended as filler in unused or inaccessible parts of code memory.

The undef instruction is type 2. It is guaranteed to generate a trap in all systems. It can be used for testing purposes and emulation.

The userdef_ instructions are type 1. These instructions are reserved for user-defined and application-specific purposes.

Instructions with erroneous coding should preferably behave as type 1. This includes instruction codes with nonzero values in unused fields, operand types not supported, or any other bit pattern with no defined meaning in any field. Type 3 behavior may alternatively be allowed in these cases. If so, the instruction should behave as if it were coded correctly.

All other opcodes not explicitly defined are type 1. These may be used for future instructions.

Small systems with no operating system and no trap support should define alternative behavior.

Chapter 6

Other implementation details

6.1 Endianness

The memory organization is little endian. Instructions for byte swapping are provided for reading and writing big endian binary data files.

Rationale

The storage of vectors in memory would depend on the element size if the organization was big endian. Assume, for example, that we have a 128 bit vector register containing four 32-bit integers, named A, B, C, D. With little endian organization, they are stored in memory in the order:

A0, A1, A2, A3, B0, B1, B2, B3, C0, C1, C2, C3, D0, D1, D2, D3,

where A0 is the least significant byte of A and D3 is the most significant byte of D. With big endian organization we would have:

A3, A2, A1, A0, B3, B2, B1, B0, C3, C2, C1, C0, D3, D2, D1, D0.

This order would change if the same vector register is organized, for example, as eight integers of 16 bits each or two integers of 64 bits each. In other words, we would need different read and write instructions for different vector organizations.

Little endian organization is more common for a number of reasons that have been discussed many times elsewhere.

6.2 Implementation of call stack

There are various methods for saving the return addresses for function calls: a link register, a separate call stack, or a unified stack for return addresses and local data. Here, we will discuss the pro's and con's of each of these methods.

Link register

Some systems use a link register to hold the return address. The advantage of a link register is that a leaf function can be called without storing anything on the stack. This saves cache bandwidth in programs with many leaf function calls. The disadvantage is that every non-leaf function needs to save the link register on a stack before calling another function, and restore the leaf register before returning.

If we decide to have a link register then it should be a special register, not one of the general purpose registers. A link register does not need to support all the things that a general purpose register can do. If the link register is included as one of the general purpose registers then it will be tempting for a programmer to save it to another register rather than on the stack, and then end the function by

jumping to that other register. This will work, of course, but it will interfere with the way returns are predicted. The branch predictor uses a special mechanism for predicting returns, which is different from the mechanism used for predicting other jumps and branches. This mechanism, which is called a return stack buffer, is a small rolling cache that remembers the addresses of the last calls. If a function returns by a jump to another register than the link register then it will use the wrong prediction mechanism, and this will cause severe delays due to misprediction of the subsequent series of returns. The return stack buffer will also be messed up if the link register is used for indirect jumps or other purposes.

The only instructions that are needed for the link register other than call and return, are push and pop. We can reduce the number of instructions in non-leaf functions by making a combined instruction for “push link register and then call a function” which can be used for the first function call in a non-leaf function, and another instruction for “pop link register and then return” to end a non-leaf function. However, this will violate the principle that we want to avoid complex instructions in order to simplify the pipeline design.

The only performance gain we get from using a link register is that it saves cache bandwidth by not saving the return address on leaf function calls. It will not affect performance in applications where cache bandwidth is not a bottleneck. The performance of the return instruction is not influenced by cache bandwidth because it can rely on the prediction in the return stack buffer.

The disadvantage of using a link register is that the compiler has to treat leaf functions and non-leaf functions differently, and that non-leaf functions need extra instructions for saving and restoring the leaf register on the stack.

Therefore, we will not use a link register in the ForwardCom architecture.

Separate call stack

We may have two stacks: a call stack for return addresses and a data stack for the local data of each function. A program without recursive functions will usually have a quite limited call depth so that the entire call stack, or at least the “hot” part of it, can be stored on the chip. This will improve the performance because no memory or cache operations are needed for call and return operations – at least not in the critical innermost loops of the program. It will also simplify prediction of return addresses because the on-chip rolling stack and the return stack buffer will be one and the same structure.

The call stack can be implemented as a rolling register stack on the chip. The call stack is spilled to memory if it overflows. A return instruction after such a spilling event will use the on-chip value rather than the value in memory as long as the on-chip value has not been overwritten by new calls. Therefore, the spilling event is unlikely to occur more than once in the innermost part of a program.

The pointer for the call stack should not be a general purpose register because the programmer will rarely need to access it directly. Direct manipulation of the call stack is only needed in a stack unroll event (after an exception or long jump) or a task switch.

A function does not have easy access to the return address that it was called from. Information about the caller may be supplied explicitly as a function parameter in the rare case that it is needed. There is a security advantage in hiding the return address inside the chip. This prevents overwriting return addresses in case of program errors or malicious buffer overflow attacks.

The disadvantage of having a separate call stack is that it makes memory management more complicated because there are two stacks that can potentially overflow. The size of the call stack can be predicted accurately for programs without recursive functions by using the method described on page 124.

A separate call stack may be implemented with the ForwardCom architecture. The size of the on-chip stack buffer and other details will be implementation-dependent.

Unified stack for return addresses and local data

Many current systems use the same stack for return addresses and local data. This method may be used with the ForwardCom architecture because it is simple to implement.

Conclusion for ForwardCom

A ForwardCom system may use a separate call stack or a unified stack, but not a link register. The hardware implementation of call and return instructions depends on whether there is one or two stacks. The dual stack system will be used for large processors where performance or security is important, while the unified stack system may be used in small processors where simplicity is preferred. A ForwardCom microprocessor does not have to support both systems, but the software does. The calling conventions defined in chapter 12.4 will make the software compatible with both single stack and dual stack processors. Tail calls can be implemented efficiently with a simple jump instruction regardless of the stack type.

6.3 Floating point errors and exceptions

Exceptions for floating point errors are disabled by default, but can be enabled with bits 26-29 in the numeric control register or a mask register. Enabled exceptions are caught as traps (synchronous interrupts).

It is a problem that an exception caused by a single element in a vector will interrupt the processing of the whole vector. The behavior of a program using floating point vectors will depend on the vector length in case of traps caused by a single vector element. We can rely on the generation and propagation of NAN and INF values instead of traps if we want consistent results on different processors with different vector lengths.

NAN values will be propagated through the sequence of floating point calculations. A NAN can contain a bit pattern of diagnostic information called the payload, and this bit pattern is propagated to the result. A problem arises when two different NANs are combined, for example $\text{NAN}_1 + \text{NAN}_2$. The IEEE standard (754-2008) specifies that only one of the two NAN operands is propagated to the result. This violates the fundamental principle that addition is commutative. The result can be inconsistent when a compiler swaps the two operands. Another problem with the IEEE standard is that NAN values are not propagated through the max and min instructions according to this standard.

Here, it is proposed to deviate from this unfortunate standard and output the OR combination of the input NAN payloads when multiple NAN operands are combined. This will make the propagation of NANs more useful and consistent. Different bits in the NAN payload can be used for indicating different error conditions. If multiple different error conditions have arisen in a sequence of calculations then all these conditions can be traced in the final result. This better propagation of NAN values is enabled by setting bit 21 in the numeric control register or in a mask register.

The implementation will use only one bit in the NAN payload for each error condition. A quiet NAN has bit number -1 of the significand set, while the remaining bits are available for any payload information. The ForwardCom processor puts diagnostic information in the payload if better NAN propagation is enabled by bit 21 in the numeric control register or a mask register. Bit number -2 in the significand indicates invalid arithmetic operations such as $0/0$, $0 \cdot \infty$, $\infty - \infty$, etc. Bit number -3 indicates a square root of a negative number, and other complex number results. The remaining payload bits are available for other purposes such as function libraries.

Other methods for generating error messages in function libraries are discussed on page 112.

6.4 Detecting integer overflow

There is no common standard method for detecting overflow in integer calculations. The detection of overflow in signed integer operations is a real nightmare in some programming languages like C++ (see e. g. stackoverflow.com/questions/199333/how-to-detect-integer-overflow-in-c-c).

It would be nice to have a reliable way of detecting integer overflow and perhaps to propagate it through a series of calculations, analogous to the NAN propagation for floating point calculations, so that errors can be checked at the end of a series of calculations rather than after each operation. Compilers could support this method by offering overflow detection with a try/catch block. It is more likely that compilers will support integer overflow detection if the hardware offers a reasonable method.

The following methods have been proposed:

1. Use a few vacant bits in the mask registers for detecting and propagating overflow and other errors. This method has a number of problems that will impede out-of-order execution. The mask register will be used not only for input to each instruction but also output. Each instruction will then have two outputs rather than one. This will make the out-of-order scheduler much more complicated, and it will cause undesired dependencies when the same mask register is used for multiple instructions that otherwise would be independent.
2. Use the even-numbered elements in a vector register for normal calculation on integers and use the following odd-numbered elements for the overflow information. The overflow information is propagated together with the calculated values. This method will be efficient for scalar integer calculations, but wasteful for vectors because half the vector elements are used only for this purpose.
3. Use one element of a vector for the overflow bits of all the other elements. This method may be tempting because it does not waste as much register space as the previous method, but it will have inferior performance because of the transport delay when moving overflow bits to a distant part of a long vector.
4. Add extra bits in the vector registers for overflow information. All vector registers will have one extra overflow bit for each 32 bits of normal data. These overflow bits are preserved when a vector register is saved and restored with the `save_cp` and `restore_cp` instructions, but they are lost when the vector is saved as normal data. The behavior of the overflow bits is controlled by bits in the numeric control register or a mask register to enable the detection and propagation of signed and unsigned integer overflow. An extra instruction must be provided for extracting the overflow bits from a vector register.
5. Generate a trap in case of integer overflow. Use a mask register or the numeric control register to control this behavior. Bit 6 enables a trap for unsigned integer overflow and bit 7 enables a trap for signed integer overflow. This method requires little extra code, but it is subject to the problem that the behavior of vector code depends on the vector length in case of traps, as explained in the previous chapter for floating point errors.

Method 2 is tentatively supported here with the optional instructions `add_oc`, etc., described on page 65.

Support for method 4 may be considered, since it would be more efficient and useful. The cost of implementing method 4 is that we will need 3% more bits in the vector registers; the `save_cp` and `restore_cp` instructions will be more complicated; and the compiler has to check for overflow before saving vectors to memory in the normal way.

Method 5 should be supported. It is useful for integer code in general purpose registers and it is useful for verifying that overflow does not occur in vector registers.

These methods should not detect overflow in saturated arithmetic instructions and shift instructions.

6.5 Multithreading

The ForwardCom design makes it possible to implement very large vector registers to process large data sets. However, there are practical limits to how much you can speed up the performance by using larger vectors. First, the actual data structures and algorithms often limit the vector length that can be used. And second, large vectors mean longer physical distances on the semiconductor chip and longer transport delays.

Additional parallelism can be obtained by running multiple threads in each their CPU core. The design should allow multiple CPU chips or multiple CPU cores on the same physical chip.

Communication and synchronization between threads can be a performance problem. The system should have efficient means for these purposes, including speculative synchronization.

It is probably not worthwhile to allow multiple threads to share the same CPU core and level-1 cache simultaneously (this is what Intel calls hyper-threading) because this could allow a low priority thread to steal resources from a high priority thread, and it is difficult for the operating system to determine which threads might be competing for the same execution resources if they are run in the same CPU core.

6.6 Security features

Security is included in the fundamental design of both hardware and software. This includes the following features.

- A flexible and efficient memory protection mechanism.
- Separation of call stack and data stack so that return addresses cannot be compromised by buffer overflow.
- Each thread has its own protected memory space, except where compatibility with legacy software requires a shared memory space for all threads in an application.
- Device drivers and system functions have carefully controlled access rights. These functions do not have general access to application memory, but only to a specific block of memory that an application may share with a system function when calling it. A device driver has only access to a specific range of input/output ports and system registers as specified in the executable file header and controlled by the system core.
- A fault in a device driver should not generate a “blue screen of death”, but generate an error message and close the application that called it and free its resources.
- Application programs have only access to specific resources as specified in the executable file header and controlled by the system.
- Array bounds checking is simple and efficient, using an addressing mode with built-in bounds checking or a conditional trap.
- Various optional methods for checking integer overflow.
- There is no “undefined” behavior. There is always a limited set of permissible responses to an error condition.

How to improve the security of applications and systems

Several methods for improving security are listed below. These methods may be useful in ForwardCom applications and operating systems where security is important.

Protect against buffer overflow

Input buffers must be protected against overflow. If a software-based protection is not sufficient then you may allocate an isolated block of memory for the input buffer. See page 105.

Protect arrays

Array bounds should be checked.

Protect against integer overflow

Use one of the methods for detecting integer overflow mentioned on page 95.

Protect thread memory

Each thread in an application should have its own protected memory space. See page 105.

Protect code pointers

Function pointers and other pointers to code are vulnerable to control flow hijack attacks. These include:

Return addresses. Return addresses on the stack are particularly vulnerable to buffer overflow attacks. Use a dual stack design to isolate the return stack from other data.

Jump tables. Switch/case multiway branches are often implemented as tables of jump addresses. These should use the jump table instruction with the table placed in a read-only section. See page 83.

Virtual function tables. Programming languages with object polymorphism, such as C++, use tables of pointers to virtual functions. These should use the call table instruction with the table placed in the a read-only section. See page 83.

Procedure linkage tables. Procedure linkage tables, import tables and symbol interposition are not used in ForwardCom. See page 123.

Callback function pointers. If a function receives a pointer to a callback function as parameter, then keep this pointer in a register rather than saving it to memory.

State machines. If a state machine or similar algorithm is implemented with function pointers then place these function pointers in a constant array, use a state variable as index into this array and check the index for overflow. The compiler should have support for defining an array of relative function pointers in a read-only section and access them with the call table instruction.

Other function pointers. Most uses of function pointers can be covered by the methods described above. Other uses of function pointers should be avoided in high security applications, or the pointers should be placed in protected memory areas or with unpredictable addresses. (See Code-Pointer Integrity link).

Control access rights of application programs

The executable file header of an application program should include information about which kinds of operations the application needs permission to. This may include permission to various network activities, access to particular sensitive files, permission to write executable files and scripts, permission to install drivers, permission to spawn other processes, permission to inter-process communication, etc. The user should have a simple way of checking if these access rights are acceptable. We may implement a system for controlling the access rights of scripts as well. Web page scripts should run in a sandbox.

Control access rights of device drivers

Many operating systems are giving very extensive rights to device drivers. Rather than having a bureaucratic centralized system for approval of device drivers, we should have a more careful control of the access rights of each device driver. The system call instruction in ForwardCom gives a device driver access to only a limited area of application memory (see page 88). The executable file header of a device driver should have information about which ports and system registers the device driver has access to. The user should have a simple way of checking if these access rights are acceptable.

Standardized installation procedure

Malware protection should be an integral part of the operating system, not a third-party add on with possible compatibility problems. The operating system should provide a standardized way of installing and uninstalling applications. The system should refuse to run any program, script or driver that has not been installed through this procedure. This will make it possible for the user to review the access requirements of all installed programs and to remove any malware or other unwanted software through the normal uninstallation procedure.

Chapter 7

Programmable application-specific instructions

Rather than implementing a lot of special instructions for specific applications, we may provide a means for generating user-defined instructions which can be coded in a hardware description language, e. g. VHDL or Verilog.

The microprocessor can have an optional FPGA or similar programmable hardware. This structure can be used for making application-specific instructions or functions, e. g. for coding, encryption, data compression, signal processing, text processing, etc.

If the processor has multiple CPU cores then each core may have its own FPGA. The hardware definition code is stored in its own cache for each core. The operating system should prevent, as far as possible, that the same core is used for different tasks that require different hardware codes. There may be features for allowing an application to monopolize an FPGA or part of it.

If it cannot be avoided that multiple applications use the same FPGA in the same CPU core, then the code, as well as the contents of any memory cells in the FPGA, must be saved on each task switch. This saving may be implemented as lazy, i. e. the contents is only swapped when the second task needs the FPGA structure that contains code for the first task.

There must be instructions for accessing the user-defined functions, including means for input and output, and for adapting to the latency of the user-defined functions.

Chapter 8

Microarchitecture and pipeline design

The ForwardCom instruction set is intended to facilitate a consistent and efficient design of the pipeline of a superscalar microprocessor. Instructions can have no more than one destination operand, up to three source operands, a mask register, a fallback register, and a register specifying vector length. The last source operand can be a register, a memory operand or an immediate constant. All other operands are registers, except for memory write instructions. The total number of input registers to an instruction, including source operands, mask, fallback, memory base pointer, index, and vector length specifier cannot exceed five, and may be limited to four.

No instruction can have more than one memory operand. Only few instructions can have both a memory source operand and an immediate operand. Any extra immediate operand field can be used for option bits, memory offset, or memory limit.

A high performance pipeline may be designed as superscalar with the following stages.

- Fetch. Fetching blocks of code from the instruction cache, one cache line at a time, or as determined by the branch prediction machinery.
- Instruction length decode. Determine the length of each instruction and identify tiny instructions. Distribute the first P instructions into each their pipeline lane, where P is the number of parallel lanes implemented in the pipeline. Excess instructions may be queued for the next clock cycle. The length of an instruction is determined by two bits of the first code word in order to simplify this process.
- Instruction decode. Identify and classify all operands, opcode and option bits. Determine input and output dependencies. A consistent template system simplifies this step.
- Register allocation and renaming.
- Instruction queue.
- Put instructions into reservation station. Schedule for address calculator.
- Calculate address and length of memory operand. Check access rights. Do calculations that depend on immediate operand only.
- Read memory operand. Schedule for execution units.
- Execution units.
- Retire or branch.

It is not necessary to split instructions into micro-operations if the reading of memory operands is done in a separate pipeline stage and instructions are allowed to stay in the reservation station until the memory operand has been read.

Each stage in the pipeline should ideally require only one clock cycle. Instructions waiting for an operand should stay in the reservation station. Most instructions will use only one clock cycle in the execution

unit. Multiplication and floating point addition need a pipelined execution unit with several stages. Division and square root may use a separate state machine.

Jump, branch, call, and return instructions also fit into this pipeline design.

The reservation station has to consider all the input and output dependencies of each instruction. Each instruction can have up to four or five input dependencies and one output dependency.

There may be multiple execution units so that it is possible to run multiple instructions in the same clock cycle if their operands are independent.

An efficient out-of-order processing requires renaming of the general purpose registers and vector registers, but not necessarily the special registers.

Complex instructions and microcode should generally be avoided. We do not have an instruction for saving or restoring all registers during a task switch. Instead, the necessary instructions for saving and restoring registers are implemented as tiny instructions to reduce the size of an instruction sequence that saves all registers.

The following instructions are moderately complex: call, return, div, rem, sqrt, cmp_swap, save_cp, restore_cp. These instructions may be implemented as dedicated state machines. The same applies to traps, Interrupts and system calls.

Some current CPUs have a “stack engine” in order to predict the value of the stack pointer for a push, pop or call instruction when preceding stack operations are delayed due to operands that are not available yet. Such a system is not needed if we have a dual stack design (see page 93). Even with a single stack design, there is little need for a stack engine because push and pop operations will be rare in critical parts of the code if the function calling conventions in this document are followed (chapter 12.4).

Branch prediction is important for the performance. We may implement four different branch prediction algorithms: one for ordinary branches, one for loops, one for indirect jumps, and one for function returns. The long form of branch instructions have an option bit for indicating loop behavior. The short form of branch instructions does not have space for such a bit. The initial guess may be to assume loop behavior if the branch goes backwards and ordinary branch behavior if the branch goes forwards. This assumption may be corrected later, if necessary, by the branch prediction machinery.

The code following a branch is executed speculatively until it is determined whether the prediction was right. We may implement features for running both sides of a branch speculatively at the same time.

The ForwardCom design allows large microprocessors with very long vector registers. This requires special design considerations. The chip layout of vector processors is typically divided into “data lanes” so that the vertical transfer of data from a vector element to the corresponding vector element in another vector (i. e. same lane) is faster than the horizontal transfer of data from one vector element to another element at another position of the same vector (i. e. different lane). This means that instructions that transfer data horizontally across a vector, such as broadcast and permute instructions, may have longer latencies than other vector instructions. The scheduler needs to know the instruction latency, and this can be a problem if the latency depends on the distance of data transfer on very long vectors. This problem is addressed by indicating the vector length or the distance of data transfer for such instructions in a separate operand, which always uses the RS register field. This information may be redundant because the vector length is stored in the vector register operands, but the scheduler needs this information as early as possible. The other register operands are typically not ready until the clock cycle where they go to the execution unit, while the vector length is typically known earlier. The microprocessor can read the RS register at the address calculation stage in the pipeline, where it also reads any pointer, index register and vector length for memory operands. This allows the scheduler to predict the latency a few clock cycles in advance. The instruction set provides the extra information about vector length or data transfer length in RS for all instructions that involve horizontal data transfer, including memory broadcast, permute, insert, extract and shift instructions, but not

broadcasting of immediate constants.

The data path to the data cache and memory should be quite wide, possibly matching the maximum vector length, because cache access and memory access are typical bottlenecks.

8.1 Proposal for reducing branch misprediction delay

Modern superscalar processors often have quite long pipelines. This gives a long branch misprediction delay. The branch misprediction delay is normally equal to the number of pipeline stages from a branch instruction is fetched until it is executed.

It may be possible to reduce this delay by executing branch instructions as early as possible in the pipeline. Any preceding instructions that a branch instruction depends on may also be executed early. It is proposed to execute all control transfer instructions (branch, jump, call, and return) in the front end of the pipeline, before the register renaming and scheduler. Any preceding instruction that a branch depends on, could likewise be executed in the front end. This idea is somewhat similar to the principle described in the article:

Sheikh, Rami, James Tuck, and Eric Rotenberg. "Control-Flow Decoupling." In Proceedings of the 2012 45th Annual IEEE/ACM International Symposium on Microarchitecture, 329–340. IEEE Computer Society, 2012. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.952.6936&rep=rep1&type=pdf>

The vision is this: There are two sets of ALU's, one in the front end used for control flow and other simple instructions, and one in the back end used for all other calculations. The front end does not use register renaming, but relies on the permanent register file.

In most cases, the control flow depends only on simple integer operations such as incrementing a loop counter. The front end ALU's should be able to handle simple integer operations on general purpose registers. The front end may handle all simple integer instructions on general purpose registers if the operands are available in the permanent register file. If the capacity of the front end ALU's is insufficient then it may prioritize instructions on those registers that have recently been used as loop counters or branch conditions, but it may be cheaper to implement a few extra ALU's to increase the front end capacity than to implement such a mechanism for prioritization.

Instructions that are executed in the front end use the permanent register file for both input and output. The remaining instructions are sent to the out-of-order back end. All general purpose register operands are replaced by their value if the value is available in the permanent register file before the instruction is sent to the back end. Instructions that execute in the front end may write their result not only to the permanent register file but also to any pending instruction that needs it.

The result of this mechanism is that a loop that is controlled only by simple instructions on general purpose registers will be unrolled in the front end. The instructions in the loop body may be passed on to the back end with the loop counter replaced by its value. The back end has register renaming so that it can execute the body of the second iteration before it has finished the body of the first iteration, etc.

All other control flow, such as branches, jumps, calls, returns, etc., will be unrolled in the front end as well so that control flow is effectively resolved before execution as far as it does not depend on delayed data.

The front end may support a minimum of out-of-order execution in the sense that an instruction that is waiting for an operand should not delay any independent subsequent instructions. But the front end does not need a complicated structure with a long queue and scheduler.

The front end should have access to the permanent register file, while the back end uses temporary registers with register renaming. The renamed registers will be written to the permanent register file when they retire. We have to keep track of whether the permanent registers are up to date. When the decoder sees an instruction that writes to a register, it will mark this register as invalid in the per-

manent register file and remember which instruction it belongs to. When this instruction is executed (whether in the front end or the back end), the register entry is marked as valid, unless another instruction writing to the same register has been seen in the meantime.

The front end may fetch and decode speculatively, but not execute speculatively. It may even fetch both sides of a branch that it expects to have poor prediction. Fetching both sides of a branch becomes cheaper here than with a traditional superscalar design because it only needs to duplicate two pipeline stages (fetch and decode) in order to minimize branch misprediction bubbles.

The fetch and decode stages should have higher throughput than the rest of the pipeline so that it can catch up and fill the pipeline after a branch bubble. The fetch stage should be able to fetch a large block of code in one clock cycle, possibly crossing a cache line boundary. The decoder should be able to decode the first one or two instructions in a fetched code block in the first clock cycle, and determine the starting points of the remaining instructions. This will allow it to decode maybe four or more instructions in the second clock cycle after fetching a block of code. The aim of this is to keep the branch misprediction delay as low as possible. This delay will be 3 clock cycles if there are three stages in the front end path: fetch, decode, and execute. This requires that the two decode stages are joined into one.

Multi-way branches with a jump table are particularly critical here because they have a memory operand that may be delayed due to cache miss. Implementing memory access in the front end may be too complicated. Therefore, such instructions probably have to go to the back end. The out-of-order mechanism may give high priority to these instructions and execute them as soon as the register operand is available. Jump tables are normally placed in the read-only memory block (see page 104). The speed of access to jump tables may be increased by having a separate cache for read-only memory.

The demand for advanced branch prediction is somewhat relaxed if we can drastically reduce the branch misprediction delay with the measures described here. The branch prediction machinery in modern processors is very complicated and requires a lot of silicon space for branch target buffers and pattern/history tables. We can cut down on this and use a simpler branch prediction scheme if branches are resolved in the front end. A few extra ALU's in the front end will use less silicon space than an advanced branch prediction mechanism.

An extra advantage of this design is that the registers used in memory operands for pointer, index, and vector length are likely to be available early so that memory operands can be fetched before they are needed.

Chapter 9

Memory model

The address space is using unsigned 64-bit addresses and 64-bit pointers. Future extension to 128-bit addresses is possible, but this will probably not be relevant in a foreseeable future.

Absolute addresses are rarely used. Most data objects, functions and jump targets are addressed with signed offsets of 32 bits or less relative to some reference point contained in a 64-bit pointer. This pointer can be the instruction pointer (IP), the data section pointer (DATAP), the thread data pointer (THREADDP), the stack pointer (SP), or a general purpose register.

An application can have access to the following sections of data:

- Program code (CODE). This memory block is executable with or without read access, but without write access. The CODE section can be shared between multiple processes running the same program.
- Read-only program data (CONST). This contains constants and tables used by the program without write access. It may be shared between multiple processes.
- Static read/write program data sections, which can be initialized (DATA) and uninitialized (BSS). This is used for global data and for static data inside functions. Multiple instances are needed if multiple processes are running the same code.
- Stack data (STACK). This is used for non-static data inside functions. Each process or thread has its own stack, addressed relative to the stack pointer. The stack grows downward from high to low addresses when data are added to the stack.
- Program heap (HEAP). Used for dynamic memory allocation by an application program.
- Thread data (THREADD). Allocated when a thread is created and used for thread-local static data and thread environment block. This has one instance for each thread.

References within the CODE section use 8-bit, 16-bit, 24-bit and 32-bit signed references relative to the instruction pointer, scaled by the code word size which is 4 bytes.

The CONST section is preferably placed immediately before the CODE section. Data in the CONST section are mostly addressed relative to the instruction pointer with no scale factor. (In case of a pure Harvard architecture, the CONST section may be placed in readable program memory to be addressed relative to the instruction pointer, or it may be placed in data memory and addressed relative to DATAP).

The DATA and BSS sections are addressed relative to the data section pointer (DATAP) which is a special register that points to some reference point in these sections. The preferred reference point is where DATA ends and BSS begins. Multiple running instances of the same program will have different values of the data section pointer. The CODE and CONST sections contain no absolute references to DATA or BSS, only references relative to the data section pointer. This makes it possible for multiple processes to share the same CODE and CONST sections, but have each their private DATA

and BSS sections without the need for virtual address translation. The DATA and BSS sections can be placed anywhere in the address space independently of where CONST and CODE are placed.

STACK data are addressed relative to the stack pointer (SP). Heap data are addressed through pointers provided by the heap allocation function.

Thread data are addressed relative to the thread data pointer (THREADP), which is separate for each thread in the process. The thread data section may be allocated on the stack when a new thread is created.

The STACK, DATA, BSS, and HEAP data sections are preferably kept together in one contiguous block in order to optimize caching and memory management.

This model allows the program to access up to 8 GB of CODE, 2 GB of CONST, 2 GB of DATA, 2 GB of BSS, 2 GB of THREADD, an almost unlimited size of STACK with 2 GB frames, and an almost unlimited amount of HEAP data. A pointer to the CONST section is provided in the thread environment block in order to access CONST data in the rare case that the distance between code and data exceeds 2 GB or in order to avoid address relocation.

The end of the combined data memory block must have an unused space of the same size as the maximum vector length. This will enable the `restore_cp` instruction to read more than necessary when restoring a vector of unknown length. It will also allow a function that searches for the end of a zero-terminated string to read one vector-length piece of the string at a time without causing access violation by reading into unavailable memory space.

Most microprocessor systems have the stack growing backward. The ForwardCom system has the same, but mainly for a different reason. When a vector register is saved on the stack, it is stored as the length followed by the amount of data indicated by the length. When the vector register is restored (using the `restore_cp` instruction), it is necessary to read the length followed by the data. The stack pointer must point to the low end where the length is stored, otherwise it would be impossible to find where the length is stored.

9.1 Thread memory protection

Each thread must have its own stack. The thread data (THREADD) may be placed on this stack. The ForwardCom system allows inter-thread memory protection. The stack data of the main thread of a program is accessible to all its child threads, but all other threads in the program can have private data which is not accessible to any other threads, not even to the main thread. Any communication and synchronization between threads must use static memory or memory belonging to the main thread.

It is recommended to use this inter-thread memory protection in all cases except where legacy software requires one memory space shared by all threads.

Isolated memory blocks

It is possible to make a system function that allocates an isolated memory block surrounded by inaccessible memory on both sides. Such a memory block, which will be accessible only to a specific thread, can be used for example for an input buffer in cases where security requirements are high. Each thread can have only a limited number of such protected memory blocks because of the limited size of the memory map.

9.2 Memory management

It is a design goal to minimize memory fragmentation and to minimize the need for virtual address translation. Current designs often have very complicated memory management systems with multi-level address translation, large translation-lookaside-buffers (TLB), and huge page tables. We want to

replace the TLB, which has a large number of fixed-size memory blocks, by a memory map with a few memory blocks of variable size. In most cases, the main thread of an application will only need three blocks of memory: CONST (read only), CODE (execute only), and the combined STACK+DATA+BSS+HEAP (read-write). A child thread needs one more entry for its private stack. Similar blocks are defined for system code.

A memory map with such a limited number of entries can easily be implemented on the chip in a very efficient way and it can easily be changed on task switches. Each process and each thread must have its own memory map. The memory is not organized into fixed-size pages.

The memory map supports virtual address translation in the form of a constant offset that defines the distance between the virtual address and the physical address for each map entry. The hardware should not waste time and power on virtual address translation when it is not used.

A limited number of extra entries are provided in the memory map to deal with cases where the memory becomes fragmented, but memory fragmentation can be avoided in most cases. The following techniques are provided to simplify memory management and avoid memory fragmentation:

- There is only one type of function libraries which can be used for both static and dynamic linking. These are linked with a mechanism that keeps the CONST, CODE and DATA sections contiguous with the similar sections of the main program in most cases. This technique is described on page 123 below.
- The required stack size is calculated by the compiler and the linker so that stack overflow can be avoided in most cases. This technique is described on page 124.
- The operating system can keep statistical records of the heap use of each program in order to predict the required heap size. The same technique can be used for predicting stack use in cases where the required stack size cannot be predicted exactly (e. g. recursive function calls).

The memory space may become fragmented despite the use of these techniques. Problems that can result in memory fragmentation are listed below.

- Recursive functions can use unlimited stack space. We may require that the programmer specifies a maximum recursion level in a pragma.
- Allocation of variable-size arrays on the stack using the `alloca` function in C. We may require that the programmer specifies a maximum size.
- Runtime linking. The program can reserve space for loading and linking function libraries at runtime (see page 123). The memory may become fragmented if the memory space reserved for this purpose turns out to be insufficient.
- Script languages and byte code languages. It is difficult to predict the required size of stack and heap when running interpreted or emulated code. It is recommended to use a just-in-time compiler instead. Self-modifying scripts cannot be compiled. The same problem can occur with large user-defined macros.
- Unpredictable number of threads without protection. The required stack size for a thread may be computed in advance, but in some cases it may be difficult to predict the number of threads that a program will generate. Multiple threads will mostly share the same code sections, but they need separate stacks. The stack of a thread can be placed anywhere in memory without problems if inter-thread memory protection is used. But if memory is shared between threads and the number of threads is unpredictable then the shared memory space may become fragmented.
- Unpredictable heap size. Programs that process large amounts of data, e. g. multimedia processing, may need a large heap. A heap can use discontinuous memory, but this will require extra entries in the memory map.

- Lazy loading and code overlay. A large program may have certain code units that are rarely used and loaded only when needed. Lazy loading can be useful to save memory, but it may require virtual memory translation and it may cause memory fragmentation. A straightforward solution is to implement such code units as separate executable programs.
- Hot patching, i. e. updating of code while it is running.
- Shared memory for inter-process communication. This requires extra entries in the memory map as explained below.
- Many programs running. The memory can become fragmented when many programs of different sizes are loaded and unloaded randomly or swapped to memory.

A possible remedy against overflow of stack and heap is to place the STACK, DATA, BSS and HEAP data together (in this order) in an address range with large unused virtual address spaces below and above, so that the stack can grow downwards and the heap can grow upwards into the vacant spaces. This method can avoid fragmentation of the virtual address space, but not the physical address space. Fragmentation of the physical address space can be remedied by moving data from a memory block of insufficient size to another block that is larger. This method has the cost of a time delay when the data are moved.

If runtime linking runs into memory problems and lack of memory map entries then it is allowed to mix CONST and CODE sections together in a common section with both read and execute access. If a library function contains constant data that originate from an untrusted source, while the code is trusted, then it is preferred to put the untrusted data into the DATA section rather than the CONST section in order to prevent execution of malicious code placed in the CONST section.

An application that needs many independent memory blocks can make a separate thread to service each memory block. Each thread has its own memory map with an entry for its private memory block.

Shared memory can be used when there is a need to transfer large amounts of data between two processes. One process shares a part of its memory with another process. The receiving process needs an extra entry in its memory map to indicate read and/or write access rights to the shared memory block. The process that owns the shared memory block does not need any extra entry in its memory map. There is a limit to how many shared memory blocks an application can receive access to, because we want to keep the memory map small. If one program needs to communicate with a large number of other programs then we can use one of these solutions: (1) let the program that needs many connections own the shared memory and give each of its clients access to one part of it, (2) run multiple threads in (or multiple instances of) the program that needs many connections so that each thread has access to only one shared memory block, (3) let multiple communication channels use the same shared memory block or parts of it, (4) communicate through function calls, (5) communicate through network sockets, or (6) communicate through files.

Executable memory cannot be shared between different applications. The mechanism of interprocess calls must be used if one application needs to call a function in another application. This is described on page 112.

We can probably keep memory fragmentation so low, by using the principles discussed here, that a relatively small memory map for each thread will be sufficient to cover normal cases. This will be much more efficient than the large TLB and multilevel address translation of current designs. It will save silicon space and power; we can avoid the cost of TLB misses and page faults, and it will make task switches very fast. The actual size of the memory map will depend on the hardware implementation.

This system puts the priority on performance-critical applications. The programmer will be able to get top performance by observing some discipline to avoid memory fragmentation, for example by recycling allocated memory. However, the system should also be able to run less well-behaved applications. Applications that cause heavy memory fragmentation are probably built with less regard for

performance. The system should have methods to allow such applications to work, perhaps even software based methods to deal with memory fragmentation, but we can regard it as acceptable to make a system that prioritizes the performance of well-behaved applications at the cost of inferior performance for applications that cause heavy memory fragmentation.

Chapter 10

System programming

The system instructions have not been fully defined yet. There is more work to do making an efficient system design. However, the first experimental implementations of ForwardCom will be without operating system so the system design does not have to be fixed yet. It is preferred to spend more time on optimizing the system design rather than to define a complete standard at this early stage of development.

There should be at least three different levels of privilege:

- The system core has the highest privilege level. Memory management and thread scheduling takes place here. This is the only part that can modify memory maps and control access rights at the lower levels.
- Device drivers and system plugin modules have carefully controlled access rights. A structure similar to the memory map (see page 105) gives a device driver access to the particular range of input/output ports and system registers that it needs. A user application can give a device driver read and write access to a specific range of the data memory it owns. This is done through the system call instruction. A device driver has no access to the code memory of the application that calls it. This means that callback function pointers cannot be used with system calls.
- An application program has access to only the memory that is allocated to it or shared with it. Memory belonging to a thread is usually not shared with other threads in the same process. Application programs have access to a few system registers and no input/output ports.

Transitions between these levels are managed by the system call and system return instructions and by traps and interrupts.

There are various system registers for control purposes. In addition, there are two sets of registers used for temporary storage, one set for the device driver level and one for the system core level. The temporary registers for the device driver level are cleared for security reasons every time a device driver is called. These registers are used mainly for temporary saving of the general purpose registers.

10.1 Memory map

There are three kinds of memory access: read, write, and execute access. These kinds of access are separate, but can be combined. For example, execute access does not imply read access. Write access and execute access should not normally be combined, because self-modifying code is not allowed.

The memory map is stored in the CPU chip. Each entry has three fields: A virtual address (up to 64 bits), access rights (3 bits), and an addend for address translation (up to 64 bits). There is no memory paging. Instead, the memory blocks have variable sizes.

The entries in the memory map must be kept sorted at all times so that each memory block ends where the next block begins. The addresses must be divisible by 8. Each thread has its own memory map. A typical memory map for an application thread may look like this.

Table 10.1: Example of memory map

Start address	Access	Addend	Comment
0x10000	Read	0	CONST section
0x10100	Execute	0	CODE section
0x10800	None	0	Belongs to other processes
0x20000	Read, Write	0	Main STACK, DATA, BSS, and HEAP sections
0x24000	None	0	Belongs to other processes
0x30000	Read, Write	0	Thread STACK, thread environment block, and thread static data
0x32000	None	0	The rest belongs to other processes

There may be a few further entries for memory blocks shared between processes and for secure isolated memory blocks. A virtual memory block may have multiple entries in case the memory becomes fragmented. The addends are used for keeping the virtual addresses of the block contiguous while the physical addresses are noncontiguous. The start addresses are virtual memory addresses.

The size of the memory map is variable. The maximum size is implementation dependent. There are at least three memory maps on the chip, one for each privilege level. This makes transitions between the levels fast. The chip space used for memory maps may be reconfigurable so that the memory maps of multiple processes can remain on the chip in case the memory maps are small. This makes task switching faster.

The memory maps are controlled at the system core level. The instructions `read_memory_map` and `write_memory_map` use the vector loop mechanism for fast manipulation of memory maps.

The methods described on page 105 for avoiding memory fragmentation are important for keeping the memory maps small.

Task switches will be very fast because we have replaced the large page tables and translation-lookaside-buffer (TLB) of traditional systems with a small on-chip memory map. This makes the system suitable for real-time operating systems.

10.2 Call stack

It is possible to have either a unified stack for function data and return addresses or two separate stacks. See page 92. ForwardCom currently supports both systems. The two-stack system is safer and more efficient, while the single-stack system may be used for small processors where the simpler single-stack system is preferred.

The two-stack system has the call stack stored inside the CPU rather than in RAM memory. A method is required for saving this stack to memory when it is full. This method may use vector-size memory access. It should be possible to manipulate the call stack for task switches and for stack unrolling in the exception handler.

10.3 System calls and system functions

Calls to system functions are made with a system call instruction (`sys_call`). The system call instruction does not use addresses, but ID numbers. Each ID number consists of a function ID in the lower half and a module ID in the upper half. The module ID identifies a system module or device driver. The system core has ID = 0. Each part of the ID can be either 16 bits or 32 bits so that the combined ID is 32, 48, or 64 bits.

System add-on modules and device drivers do not necessarily have fixed ID numbers because this would require some central authority to assign these ID numbers. Instead, the program will have to

translate the module name to an ID number before the first call to a module. This translation is done by a system function with a fixed ID number. The functions within a module can have fixed or variable ID numbers.

The ID number of a system function can be put into the program in three ways:

1. The most important system functions have fixed ID numbers which can be inserted at compile time.
2. The ID number can be found at load time in the same way as load-time linking works. This is described on page 123. The loader will find the ID number and insert it in the code before running the program.
3. The ID number is found at run time before the first call to the desired function.

The calling convention for system functions is the same as for other functions, using registers for parameters and for return value. The registers used for parameters are determined by the general calling convention. The calling conventions are described in chapter 12.4. The parameter registers should not be confused with the operands for the system call instruction.

The system call instruction has three operands. The first operand is the combined ID, contained in a register (RT) or an immediate constant. The second operand (RD) is a pointer to a memory block that may be used for transferring data between the calling program and the system function. The third parameter (RS) is the size of this memory block. The last two parameters must be divisible by 8.

The calling thread must have access rights to the memory block that it shares with the system function. This can be read access or write access or both. These access rights are transferred to the system function. The system function has no access rights to any other part of the application's memory.

It is not possible to use callback function pointers with a system call because executable memory cannot be shared with a system function. Instead, the system function can call an exported function provided by the application, using the method for inter-process calls, described below.

Device driver functions should preferably have separate stacks. The system call goes first to the system core which assigns a stack to the device driver function and makes a memory map for it before dispatching the call to the desired function. Preferably, no stack is used during this dispatching. The two registers identifying a shared memory block are copied to special registers which are accessible to the called system function. The system function runs in the same thread as the application that called it, but not with the same stack.

The old values of instruction pointer, stack pointer, and DATAP are saved in system registers, to be restored when returning to the application code.

System functions, device drivers and interrupt handlers are allowed to use all general purpose registers and vector registers if they are saved and restored according to the normal calling conventions. Interrupt handlers must save and restore all registers they use.

A method is provided to get information about the register use of system functions so that it is possible to call them using the register usage conventions of either method 1 or method 2, described on page 121. The stack use of system functions is irrelevant for the caller because they do not use the stack of the calling application program.

Some important system functions must be standardized and must be available in all operating systems. This will make it possible, for example, to make a third-party function library that works in all operating systems, even if this library needs to call system functions. It will also make it easier to adapt a program for different operating systems. The list of system functions that might be standardized includes functions for thread creation, thread synchronization, setting thread priority, memory allocation, time measurement, system information, access to environment variables, etc.

There should be a selection of system libraries providing the most common user interface forms, such as graphical user interface, console mode, and server mode. These user interface system libraries should be provided for each operating system that the architecture can run on, so that the same executable program can run in different operating systems simply by linking with the appropriate user interface library when the program is installed. Such user interface libraries may be based on existing platform-independent GUI libraries such as, e. g., wxWidgets or QT. All user interface libraries must support the `error_message` function mentioned below.

10.4 Inter-process calls

Inter-process calls are mediated by a system function. This works in the following way. An application program can export a function with an entry in its executable file header. Another application can get access to this exported function by calling a system function that checks for permission and switches the memory map, the `DATAP` and `THREADP` registers and the stack pointer before calling the exported function, and switches back before returning to the caller. The call will appear as a separate thread to the called program. The general purpose registers and vector registers can be used for parameters and return value in the same way as for normal functions. This mechanism does not generate any shared memory between caller and callee. Therefore, the exported function must use only simple types that fit into registers for its parameters and return type. A block of memory can be shared between the two processes as described on page 107.

10.5 Error message handling

There is a need for a standardized way of reporting errors that occur in a program. Many current systems fail to satisfy this need, or they use methods that are not portable or thread-safe. In particular, the following situations would benefit from such a standard.

1. A function library detects an error, for example an invalid parameter, and needs to report the error to the calling program. The calling program will decide whether to recover from the error or terminate.
2. A trap is generated because of a numerical error. The program fails to catch it as an exception, or the programming language has no support for structured exception handling. The operating system must make an informative error message.
3. A program can run in different environments that require different forms of error reporting.
4. A function library in source code form, a class library, or any other piece of code needs to report an error without knowing which user interface paradigm is used (e. g. console mode or graphical user interface or server mode). It needs a standardized way of reporting the error to the operating system or to the user interface framework, which must present an error message to the user in the way that is appropriate for the user interface (e. g. pop up a message box, print to `stderr`, print to a log file, or send a message to an administrator).

It is proposed to define a standard library function named `error_message` for this purpose. All user interface frameworks must define this function. It is possible to choose between different versions of this function when the program is installed by linking with the appropriate library. The main program may override this function by defining its own function with the same name.

The `error_message` function must have the following parameters: a numerical error code, a string pointer giving an error message, and another string pointer giving the name of the function where the error occurred. These strings are coded as zero-terminated UTF-8 strings. The error message is in the English language by default. It is not reasonable to require support for many different languages (see this link for a discussion of problems with internationalization). Instead, a manual in the desired language can contain a list of error codes.

The error message string may include numerical values and diagnostic information, such as the value of a parameter that is out of range.

The `error_message` function may or may not return. If it returns then the function that called it must return in a graceful way. The `error_message` function may alternatively terminate the application or it may raise an exception or trap which is handled by the operating system in case the exception is not caught by the program.

Chapter 11

Support for multiple instruction sets

A microprocessor may support multiple instruction sets. It will be useful for compatibility with legacy software that a microprocessor with support for ForwardCom also supports one or more older instruction sets. Such a microprocessor will have different modes, one for each instruction set. Instructions of different instruction sets cannot be mixed freely. Virtualization support might be useful so that we can have one virtual machine for each instruction set.

The transitions between ForwardCom and other instruction sets can be implemented with instructions dedicated to this purpose or with software interrupts or system calls.

If mode transitions are allowed only in system code then it is not necessary to save any registers across the mode switch. But if mode transitions are possible also in application code then the following principles should be applied:

- The instruction that makes the mode switch must be aligned to an address divisible by 4.
- The hardware or operating system must make sure that all general purpose registers are preserved across a mode switch between ForwardCom and another instruction set if similar registers exist in the other instruction set.
- The software must set up stack pointers and other special registers immediately before or after the mode switch.
- The software must take care of differences in function calling conventions between the two modes.
- The hardware or operating system must make sure that the contents of at least the first eight vector registers is preserved across the mode switch. The remaining vector registers must be cleared if they are not preserved.
- If the target instruction set has a lower maximum vector length than the actual length of a vector register before conversion, then the length is truncated to the maximum length for the target instruction set.
- The vector length information is lost in the transition from ForwardCom to another instruction set if the other instruction set has no similar way of representing vector length.
- The length of each vector register is set to a value that is sufficient to contain the nonzero part of the register, or longer, when converting from another instruction set to ForwardCom,

Details that are specific to a particular other instruction set are discussed in the following sections.

11.1 Transitions between ForwardCom and x86-64

Transitions between ForwardCom and the x86-64 instruction set (with the AVX512 extension) involve the following registers:

Table 11.1: ForwardCom and x86-64 registers

x86-64	ForwardCom	Comments
rax	r0	
rcx	r1	
rdx	r2	
rbx	r3	
rsp	r4	Stack pointer. Must be set before or after conversion to x86-64.
rbp	r5	
rsi	r6	
rdi	r7	
r8 - r15	r8 - r15	
	r16-r30	Not converted
	r31	Stack pointer. Must be set after conversion to ForwardCom.
flags		Flags register. Not converted.
k0 - k7		Mask registers. Not converted.
zmm0 - zmm7	v0 - v7	Vector registers. Converted.
zmm8 - zmm31	v8 - v31	Vector registers. Converted or cleared.

It is possible to make multi-mode functions that can be called from either ForwardCom or x86-64 mode in the following way. The first four bytes of the multi-mode function consist of a short x86-64 jump instruction, which is two bytes long, followed by two bytes of zero. The jump leads to an x86-64 implementation of the code. The four bytes will be interpreted as a NOP (no operation) if the processor is in ForwardCom mode. After this follows a ForwardCom implementation of the function.

11.2 Transitions between ForwardCom and ARM

Transitions between ForwardCom and the ARM (AArch64) instruction set involve the following registers:

Table 11.2: ForwardCom and ARM registers

ARM	ForwardCom	Comments
r0 - r30	r0 - r30	
r31	r31	Stack pointer in both instruction sets
v0 - v7	v0 - v7	Vector registers. Converted
v8 - v31	v8 - v31	Vector registers. Converted or cleared
p0 - p15		Predicate registers. Not converted

The Scalable Vector Extensions (SVE) is a future extension to the ARM architecture that allows the length of vector registers to vary from 128 to 2048 bits in increments of 128 bits. The vector length in SVE is apparently controlled through predicate masks. The vector length information cannot be converted to ForwardCom because there is no unambiguous connection between each vector register and the predicate register that contains the length information.

11.3 Transitions between ForwardCom and RISC-V

Transitions between ForwardCom and the RISC-V instruction set involve the following registers:

Table 11.3: ForwardCom and RISC-V registers

RISC-V	ForwardCom	Comments
x0		Always zero.
x1 - x31	r1 - r31	General purpose registers. x1 = link register, x14 = stack pointer, x15 = thread pointer.
f0 - f7	v0 - v7	Floating point and vector registers. Converted.
f8 - f31	v8 - v31	Floating point and vector registers. Converted or cleared.

The SIMD / Vector Extensions to RISC-V are not fully developed yet (January 2017). It is therefore too early to tell whether the vector length information can be converted in a useful way during transitions between ForwardCom and RISC-V.

Chapter 12

Standardization of ABI and software ecosystem

The goal of the ForwardCom project is a vertical redesign that defines new standards not only for the instruction set, but also for the software that uses it. This will have the following advantages.

- Different compilers will be compatible. The same function libraries can be used with different compilers.
- Different programming languages will be compatible. It will be possible to compile different parts of a program in different programming languages. It will be possible to compile a function library in a programming language different from the program that uses it.
- Debuggers, profilers and other development tools will be compatible.
- Different operating systems will be compatible. It will be possible to use the same function libraries in different operating systems, except if they use system-specific functions.

The previous chapter described standardization of system calls, system functions, and error messaging. The present chapter discusses standardization of the following aspects of the software ecosystem.

- Compiler support.
- Binary data representation.
- Function calling conventions.
- Register usage conventions.
- Name mangling for function overloading
- Binary format for object files and executable files.
- Format and link methods for function libraries.
- Exception handling and stack unrolling.
- Debug information.
- Assembly language syntax.

12.1 Compiler support

Compilers can have three different levels of support for variable-length vector registers.

Level 1

The compiler will not use variable-length vectors. The compiler can call a vector function in a function library with a scalar parameter if the function is not available in a scalar version.

Level 2

The compiler can call vector functions, but not generate such functions. The compiler can vectorize a loop automatically and call a vector library function from such a loop.

Level 3

Full support. The compiler supports data types for variable-length vectors. These data types can be used for variables, function parameters, and function returns. Variable-length vectors can not be included in structures, classes, or unions because such composite types must have known sizes. Support for variable-length vectors in static and global variables is optional. General operations on variable-length vectors can be specified explicitly, including options for applying boolean vector masks and fall-back values.

Other compiler features

The compiler may support pointer arithmetic on function pointers in order to write compact call tables with relative addresses explicitly. The difference between two function pointers should be scaled by the code word size, which is 4. Without this feature, the function pointers have to be type cast to integer pointers and back again.

The compiler may have support for detecting integer and floating point overflow and other numerical errors in try-catch blocks using one of the methods discussed on page 95.

The compiler may support array bounds checking using the indexed addressing mode with bounds, or the conditional trap instruction, or simple conditional jumps to an error function.

12.2 Binary data representation

Data are stored in little-endian form in RAM memory. See page 92 for the rationale.

Integer variables are represented with 8, 16, 32, 64, and optionally 128 bits, signed and unsigned. Signed integers use 2's complement representation. Integer overflow wraps around, except in saturated arithmetic instructions.

Floating point numbers are coded with single precision (32 bits) and double precision (64 bits). There is limited support for half precision (16 bits) and optional support for quadruple precision (128 bits). All follow the IEEE Standard 754-2008.

Floating point variables with NAN values can contain and propagate diagnostic information about the cause of errors as discussed on page 94.

Boolean variables are stored as integers of at least 8 bits with the values 0 and 1 for FALSE and TRUE. Only bit 0 of the boolean variable is used, while the other bits are ignored. This rule makes it possible to use boolean variables as masks and to implement boolean functions such as AND, OR, XOR, and NOT in an efficient way with simple bitwise instructions, rather than the method used in many current systems that have a branch for each variable to check if the whole integer is nonzero. A branch instruction is needed in the compilation of expressions like $(A \ \&\& \ B)$ and $(A \ || \ B)$ only if the evaluation of B has side effects.

All variables not bigger than 8 bytes must be kept at their natural alignment.

Arrays not smaller than 8 bytes must be aligned to addresses divisible by 8. It may be recommended to align large arrays by the cache line size.

Multidimensional arrays are stored in row-major order, except where the programming language makes this impossible.

Text strings may be stored in language-dependent forms, but a standardized form is needed for system functions and for functions that are intended to be compatible with all programming languages. The proposed standard uses UTF-8 encoding. The length of the string may be determined by a terminating zero or a length specifier, or both. The rationale is this. The CPU processing time is insignificant for text strings of a length suitable for human reading. The priority is therefore on compactness. Compactness matters if the string is stored in a file or transmitted over a network. UTF-8 is more compact than UTF-16 in most cases, though less compact for some Asian languages. UTF-8 is the most common encoding used on the Internet.

12.3 Further conventions for object-oriented languages

Object oriented languages require further standards for the binary representation of special features such as virtual function tables, runtime type identification, member pointers, etc.

These details must be standardized within each programming language for the sake of compatibility between different compilers, and if possible also between different programming languages that have compatible features.

Member pointers should be implemented in a way that prioritizes good performance in the general case where only a simple offset (to data) or a pointer (to a function) is required, while additional information for contrived cases of multiple inheritance is added only when needed.

12.4 Function calling convention

Function calls will use registers for parameters as much as possible. Integers of up to 64 bits, pointers, references, and boolean scalars are transferred in general purpose registers. Vector parameters can have variable length. Floating point scalars, vectors of any type with a fixed length of up to 16 bytes, and vectors of variable length are transferred in vector registers.

The first 16 parameters to a function that fit into a general purpose register are transferred in register $r0 - r15$. The first 16 parameters that fit into a vector register are transferred in $v0 - v15$. The length of a variable-length vector parameter is contained in the same vector register that contains the data.

Composite types are transferred in vector registers if they can be considered "simple tuples" no bigger than 16 bytes. A simple tuple is a structure or class or encapsulated array for which all non-static elements have the same type, which is not a pointer. A union is treated as a structure according to its first element.

Parameters that do not fit into a single register are transferred by a pointer to a memory object allocated by the caller. This applies to: structures and classes with elements of different types, or bigger than 16 bytes. It also applies to objects that require special handling such as a non-standard copy constructor or destructor, and objects that require extra implicit storage such as tables of virtual member functions. It is the responsibility of the caller to call any copy constructor and destructor.

If there are not enough registers for all parameters then the additional parameters are provided in a list, which can be stored anywhere in memory. A pointer to this parameter list is transferred in a general purpose register. Such a list is also used if there is a variable argument list. There can be no more than one parameter list, as the same list is used for all purposes.

The rules for a parameter list are as follows. A parameter list is used if there are more than 16 parameters that fit into a general purpose register, if there are more than 16 parameters that fit into a vector register, or if there is a variable argument list. If there are less than 16 general purpose parameters then these parameters are put in general purpose registers, and the next vacant general purpose

register is used as pointer to the list. If there are 16 or more general purpose parameters, and a parameter list is needed for any reason, then the first 15 general purpose parameters are put in r0-r14, the list pointer is in r15, and the remaining general purpose parameters are put in the list. If there are more than 16 vector parameters then the first 16 vector parameters are put in v0-v15 and the remaining vector parameters are put in the list. All parameters in the list are placed in the order that they appear in the function definition, regardless of type. Variable arguments are placed last in the list because they always appear last in a function definition.

The list consists of entries of 8 bytes each. A general purpose parameter uses one entry. A vector parameter with a constant size of 8 bytes or less uses one entry. A vector parameter with a constant size of more than 8 bytes or a variable size uses two entries in the list. The first entry is the length (in bytes) and the second entry is a pointer to an array containing the vector. A parameter that would not fit into a register, if one was vacant, is transferred by a pointer in the list according to the same rules as if the pointer was in a register.

The parameter list belongs to the called function in the sense that it is allowed to modify parameters in the list if they are not declared as constant parameters. The same applies to arrays and objects with a pointer in the list. The caller can rely on parameters in the list being unchanged only if they are declared constant. The caller must put the list in a place where it cannot be modified by other threads.

Function return values follow the following rules: A single return value is returned in r0 or v0, using the same rules as for function parameters.

Multiple return values of the same type are treated as a tuple if possible and returned in v0 if the total size is no more than 16 bytes.

A function with two return values will use two registers for return, using two of the registers r0, r1, v0, v1 as appropriate, if each of the two values will fit into a single register according to the above rules. For example, a function can return a result in v0 and an error code in r0. Or a function can return two vectors of variable length.

In all cases where the return value or set of return values does not fit into at most two registers according to the above rules, the return value is placed in a space allocated by the caller through a pointer transferred by the caller in r0 and returned in r0. Any constructor is called by the callee.

A “this” pointer for a class member function is transferred in r0, except if r0 is used for a return object. In this case the “this” pointer is transferred in r1.

Rationale

It is much more efficient to transfer parameters in registers than on the stack. The present proposal allows up to 32 parameters, including variable length vectors, to be transferred in registers, leaving 15 general purpose registers and 16 vector registers for the function to use for other purposes while handing the parameters. This will cover almost all practical cases, so that parameters only rarely need to be stored in memory.

Nevertheless, we must have precise rules for covering an unlimited number of parameters if the programming language has no limit to the number of parameters. We are putting any extra parameters in a list rather than on the stack as most other systems do. The main reason for this is to make the software independent of whether there is a separate call stack or the same stack is used for return addresses and local variables. The addresses of parameters on the stack would depend on whether there is a return address on the same stack. The list method has further advantages. There will be no disagreement over the order of parameters on the stack and whether the stack should be cleaned up by the caller or the callee. The list can be reused by the caller for multiple calls if the parameters are constant, and the called function can reuse a variable argument list by forwarding it to another function. The function is guaranteed to return properly without messing up the stack even if caller and callee disagree on the number of parameters. Tail calls are possible in all cases regardless of the

number and types of parameters.

12.5 Register usage convention

Most systems have rules that certain registers have callee-save status. This means that a function must save these registers and restore them before it returns, if they are used. The caller can then rely on these registers being unchanged after the function call.

Current systems have a problem with assigning callee-save status to vector registers. Future CPU versions may make the vector registers longer, and the instructions for saving the longer registers have not been defined yet. Some systems now have callee-save status on part of a vector register because of poor foresight. It is impossible in current systems to save a vector register in a way that will be compatible with future extensions.

This problem is solved by the ForwardCom design with variable vector length. It is possible to save and restore a vector register of any length, even if this length was not supported at the time the code was compiled. It is also possible to know how much of a long vector register is actually used, because the length of a vector is saved in the register itself, so that we only need to save the part of the register that is actually used. The `save_cp` and `restore_cp` instructions are designed for this purpose (see page 53). Unused vector registers will use only little space for saving.

It still takes a lot of cache space to save the vector registers if they are long. Therefore, we want to minimize the need for saving registers. It is proposed to have two different methods to choose between. These methods are explained here.

Method 1

This is the default method which can be used in all cases, but not the most efficient method.

The rule is simply that registers `r16 – r31` and `v16 – v31` have callee-save status.

A function can use registers `r0 – r15` and `v0 – v15` freely. Sixteen registers of each type will be sufficient for most functions. If the function needs additional registers, it must save them.

All system registers and special registers have callee-save status, except in functions that are intended for manipulating these registers.

Method 2

It will be more efficient if we actually know which registers are used by each function. If function A calls function B, and A knows which registers are used by B, then A can simply choose some registers that are not used by B for any data that it needs to save across the call to B. Even a long chain of nested function calls can avoid the need to save any registers as long as there are enough registers.

If function A and B are compiled together in the same process then the compiler can easily manage this information. But if A and B are compiled separately, then we need to store the necessary information about which registers are used. This is possible with the object file format described on page 122. The information about register use must be saved in the compiled object file or library file, not in some other file that could possibly come out of sync.

Function B is preferably compiled first into an object file. This object file must contain information about which registers are modified by function B. The necessary information is simply a 64-bit number with one bit for each register that is modified (bit 0-31 for `r0-r31`, and bit 32-63 for `v0-v31`). Any registers used for parameters and return value are also marked if they are modified by the function.

When function A is compiled next, the compiler will look in the object file for B to see which registers it modifies. The compiler will choose some registers not modified by B for data that need to be saved across the call to B. Registers that are modified by B can advantageously be used in A for temporary variables that do not need to be saved across the call to B. Likewise, it will be advantageous

to use the same register for multiple temporary variables if their live ranges do not overlap, in order to modify as few registers as possible. The object file for A will contain a list of registers modified by A, including all registers modified by B and by any other functions that A may call. The object file for A contains a reference to function B. This reference must contain information about which registers A expects B to modify. If B is later recompiled, and the new version of B modifies more registers, then the linker will detect the discrepancy and prompt for a recompilation of A.

If, for some reason, A is compiled before B or no information is available about B when A is compiled, then the compiler will have to make assumptions about the register use of B. The default assumption is as specified in method 1. Function A may later be recompiled if B violates these assumptions, or simply to improve efficiency.

If two functions A and B are mutually calling each other then the easiest solution is to rely on method 1. The functions should still include the information about register use in their object files.

The compiler should preferentially allocate the lower registers first in order to minimize the problem that different library functions use different registers. It may skip `r6` and `v6` for the caller to use for masks.

The main program function is allowed to use method 2 and to modify all registers if it includes the necessary information in its object file.

Object files that are contained in a function library must include the information about register use.

System functions and device drivers cannot be accessed in the same way as normal library functions (see page 110). System functions must obey the rules for method 1, but the system should provide a method for getting information about the register use of each system function. This can be useful for just-in-time compilers.

12.6 Name mangling for function overloading

Programming languages that support function overloading use internal names with prefixes and suffixes on the function names in order to distinguish between functions with the same name but different parameters or different classes or namespaces. Many different name mangling schemes are in use, and some are undocumented. It is necessary to standardize the name mangling scheme in order to make it possible to mix different compilers or different programming languages.

The most common name mangling schemes are Microsoft and Gnu. The Microsoft scheme uses characters that cannot occur in function names (`?@$`). This prevents name clashes, but makes it impossible to call the mangled name directly or to translate e. g. C++ to C. The Gnu scheme generates mangled names that look unwieldy, but contain no special characters that prevent calling the mangled name directly. Therefore, the proposal is to use the Gnu mangling scheme (version 4 or later) with necessary additions for variable-length vectors, etc.

Functions with mangled names may optionally supplement the mangled name with the simple (non-mangled) name as a weak public alias in the object file. This makes it easier to call the function from other programming languages without name mangling. The weak linking of the alias prevents the linker from making error messages for duplicate names.

The compiler must prefix an underscore to all symbol names for languages that do not use name mangling, such as C, in order to avoid name clashes with reserved names in the assembler.

12.7 Binary format for object files and executable files

The executable file format must be standardized. The most flexible and well-structured format in common use is probably ELF. It is proposed to use ELF format for object files, function libraries, and executable files.

The details of an ELF format for ForwardCom are specified in a file named `elf_forwardcom.h`. This specification includes details for section types, symbol types, relocation types, etc. Additional information about register use (see page 121) and stack use (see page 124) is added to the file format.

File names must have extensions that indicate their type. It is proposed to use the following extensions. Assembly code: `.as`, object file: `.ob`, library file: `.li`, executable file: `.ex`.

12.8 Function libraries and link methods

Dynamic link libraries (DLLs) and shared objects (SOs) are not used in the ForwardCom system. Instead, we will use only one type of function libraries that can be used in four different ways:

1. **Static linking.** The linker finds the required functions in the library and copies them into the executable file. Only the parts of the library that are actually needed by the specific main program are included. This is the normal way that static libraries are used in current systems (`.lib` files in Windows, `.a` files in Unix-like systems such as Linux, BSD, and Mac OS).
2. **Installation-time linking.** The library is linked or re-linked into the executable during the installation process. There may be a selection of different libraries for different platforms. The library may be updated at any time if needed.
3. **Load-time linking.** The library may be distributed separately from the executable file. The required parts of the library are loaded into memory together with the executable file, and all links between the main executable and the library functions are resolved by the loader in the same way as for static linking.
4. **Run-time linking.** The required function is extracted from the library and loaded into memory, preferably at a memory space reserved for this purpose by the main program. Any reference from the newly loaded function to other functions, whether already loaded or not, can be resolved in the same way as for static linking.

These methods will improve the performance and remedy many of the problems that we encounter with the traditional DLLs and SOs. A typical program in Windows and Unix systems will require several DLLs or SOs when it is loaded. These dynamic libraries will all be loaded into each their memory block, using an integral number of memory pages each, and possibly scattered over the memory space. This leads to a waste of memory space and poor caching. A further performance disadvantage with shared objects is that they use procedure linkage tables (PLT) and global offset tables (GOT) for all accesses to functions and variables in order to support the rarely used feature of symbol interposition. This requires a lookup in the PLT or GOT for every access to a function or variable in the library, including internal references to globally visible symbols.

The ForwardCom system replaces the traditional dynamic linking with method 2 or 3 above, which will make the code just as efficient as with static linking because the library sections are contiguous with the main program sections, and all access is immediate with no intermediate tables. The time required to load the library will be similar to the time required for dynamic linking because the bottleneck will be disk access, not calculation of function addresses.

A DLL or SO can share its code section (but not its data section) between multiple running programs that use the same library. A ForwardCom library can share its code section between multiple running instances of the same program, but not between different programs. The amount of memory that is wasted by possibly loading multiple instances of the same library code is more than compensated for by the fact that we are loading only the part of the library that is actually needed and that the library does not require its own memory pages. It is not uncommon in Windows and Unix systems to load a dynamic library of one megabyte and use only one kilobyte of it.

These linking methods are efficient in the ForwardCom system because of the way relative addresses are used. The main program typically contains a `CONST` section immediately followed by a `CODE`

section. The CONST section is addressed relative to the instruction pointer so that these two sections can be placed anywhere in memory as long as they have the same position relative to each other. Now, we can place the CONST section of the library function before the CONST section of the main program, and the CODE section of the library function after the CODE section of the main program. We don't have to change any cross-references in the main program. Only cross references between the main program and the library function and between the CODE and CONST sections of the library function have to be calculated by the linking or re-linking process and inserted in the code.

A library function does not necessarily have any DATA and BSS sections. In fact, a thread-safe function has little use of static data. However, if the library function has any DATA and BSS sections, then these sections can be placed anywhere within the $\pm 2\text{GB}$ range of the DATAP pointer. The references in the library function to its static data have to be calculated relative to the point that DATAP points to; but no references to data in the main program have to be modified when a library is added as long as DATAP still points to the border between the DATA and BSS sections of the main program.

The combined main program and library file can now be loaded into any vacant spaces in memory. It will need only three entries in the memory map: (1) the combined CONST sections of library and main program, (2) the combined CODE sections of main program and library functions, and (3) the combined STACK, DATA, BSS, and HEAP of the main program and the library functions.

Run-time linking works slightly differently. The reference from the main program to the library function goes through a function pointer that is provided when the library is loaded. Any references the other way – from the library function to functions or global data in the main program – can be resolved in the same way as for static linking or through pointer parameters to the function. The main program should preferably reserve space for the CONST, CODE and DATA/BSS sections of any libraries that it will load at run time. The sizes of these reserved spaces are provided in the header of the executable file. The loader has considerable freedom to place these sections anywhere it can in the event that the reserved spaces are insufficient. The only requirements are that the CONST section of the library function is within a range of $\pm 2\text{GB}$ of the CODE section of the library, and the DATA and BSS sections of the library are within $\pm 2\text{GB}$ of DATAP. The library function may be compiled with a compiler option that tells it not to use DATAP. The function will load the absolute address of its DATA section into a general purpose register and access its data with this register as pointer.

12.9 Predicting the stack size

In most cases, it is possible to calculate exactly how much stack space an application needs. The compiler knows how much stack space it has allocated in each function. We only have to make the compiler save this information. This can be accomplished in the following way. If a function A calls a function B then we want the compiler to save information about the difference between the value of the stack pointer when A is called and the stack pointer when A calls B. These values can then be summed up for the whole chain of nested function calls. If function A can call both function B and function C then each branch of the call tree is analyzed and the value for the branch that uses most stack space is used. If a function is compiled separately into its own object file, then the information must be stored in the object file.

A function can use any amount of memory space below the address pointed to by the stack pointer (a so-called red zone) if this is included in the stack size reported in the object file, provided that the system has a separate system stack.

The amount of stack space that a function uses will depend on the maximum vector length if full vectors are saved on the stack. All values for required stack space are linear functions of the vector length: $\text{Stack_frame_size} = \text{Constant} + \text{Factor} \cdot \text{Max_vector_length}$. Thus, there are two values to save for each function and branch: Constant and Factor. We need separate calculations for each thread and possibly also information about the number of threads. If there are two stacks then we

need to save separate values for the call stack and the data stack. The size of the call stack does not depend on the maximum vector length.

The linker will add up all this information and store it in the header of the executable file. The maximum vector length is known when the program is loaded, so that the loader can finish the calculations and allocate a stack of the calculated size before the program is loaded. This will prevent stack overflow and fragmentation of the stack memory. Some programs will use as many threads as there are CPU cores, for optimal performance. It is not essential, though, to know how many threads will be created because each stack can be placed anywhere in memory if thread memory protection is used (see page 105).

In theory, it is possible to avoid the need for virtual address translation if the following four conditions are met:

- The required stack size can be predicted and sufficient stack space is allocated when a program is loaded and when additional threads are created.
- Static variables are addressed relative to the data section pointer. Multiple running instances of the same program have different values in the data section pointer.
- The heap manager can handle fragmented physical memory in case of heap overflow.
- There is sufficient memory so that no application needs to be swapped to a hard disk.

A possible alternative to calculating the stack space is to measure the actual stack use the first time a program is run, and then rely on statistics to predict the stack use in subsequent runs. The same method can be used for heap space. This method is simpler, but less reliable. The calculation of stack requirements based on the compiler is sure to cover all branches of a program, while a statistical method will only include branches that have actually been used.

We may implement a hardware register that measures the stack use. This stack-measurement register is updated every time the stack grows. We can reset the stack-measurement register when a program starts and read it when the program finishes. This method can be useful if the program contains recursive function calls. We don't need a hardware register to measure heap size. This information can be retrieved from the heap manager.

These proposals can eliminate or reduce memory fragmentation in many cases so that we only need a small memory map which can be stored on the CPU chip. Each process and each thread will have its own memory map. However, we cannot completely eliminate memory fragmentation and the need for virtual memory translation because of the complications discussed on page 105.

12.10 Exception handling, stack unrolling and debug information

Executable files must contain information about the stack frame of each function for the sake of exception handling and stack unrolling for programming languages that support structured exception handling. It should also be used for programming languages that do not support structured exception handling in order to facilitate stack tracing by a debugger.

This system should be standardized, and both single stack and dual stack systems should be supported. It is recommended to use a table-based method that does not require a stack frame register.

Debuggers need information about line numbers, variable names, etc. This information should be included in object files when requested. The debug information may be copied into the executable file or saved in a separate file which is stored together with the executable file. It is yet to be decided which system to use.

12.11 Assembly language syntax

The definition of a new instruction set should include the definition of a standardized assembly language syntax. The syntax should be suitable for human processing, not only for machine processing. We must avoid a situation similar to the x86 environment where many different syntaxes are in use, with different instruction names and different orders of the operands.

ForwardCom will use a syntax similar to C and Java in order to make it more intelligible to high-level language programmers and to avoid any confusion over what is source and destination. The details are described on page 129.

Chapter 13

Binary tools

All the basic development tools for ForwardCom except compilers are combined into a single executable file named **forw** or **forw.exe**.

The tools can be run using the following command lines:

assemble	forw -ass assemblyfile objectfile
disassemble	forw -dis objectfile assemblyfile
link	forw -link exefile objectfiles
make library	forw -lib libraryfile objectfiles
emulate	forw -emu exefile
dump	forw -dump objectfile
help	forw -help

General options:

The following options can be used with most or all of the commands listed above:

@file	read additional command line options or file names from file.
-list=file	alternative instruction list file.
-wdNNN	disable Warning NNN.
-weNNN	treat Warning NNN as Error. -wex: treat all warnings as errors.
-edNNN	disable Error number NNN.
-ewNNN	treat Error number NNN as Warning.

The list of instructions is stored separately in a comma-separated file named **instruction_list.csv**. This file must be present when running the assembler, disassembler, or emulator. The format for the instruction list is defined in table 4.1.

13.1 Assembler

Introduction

The assembler is using a syntax similar to C and Java in order to make the code intelligible to high-level language programmers. The basic syntax of an instruction looks like this:

```
datatype destination_operand = instruction(source_operands)
```

This syntax leaves no doubt about which operands are source and destination. You can also use common operators, such as + - * / etc. instead of instructions that correspond to these operators.

Branches and loops can use conditional jump instructions or the high-level language keywords: if, else, for, do, while, break, continue.

Before defining the syntax details we will look at a few examples.

The following example shows a function that calculates the factorial of an integer:

Example 13.1.

```
code section execute          // define executable code section

// factorial function calculates n!
// input: r0, output: r0
factorial function public
if (uint64 r0 <= 20) {      // check for overflow, 64 bit unsigned
    uint64 r1 = 1           // start with 1
    while (uint64 r0 > 1) { // loop through r0 values
        uint64 r1 *= r0     // multiply all values
        uint64 r0--        // count down to 1
    }
    uint64 r0 = r1         // put result in r0
    return                 // normal return from function
}
int64 r0 = -1              // overflow. return max unsigned value
return                     // error return
factorial end              // end of function

code end                    // end of code section
```

The next example illustrates the use of the efficient vector loop described on page 8. It calculates the polynomial $y = 0.5x^2 - 4x + 1$ for all elements of an array x and stores the results in an array y .

Example 13.2.

```
data section read write datap    // define data section
% arraysize = 100                 // define constant
float x[arraysize], y[arraysize] // define arrays
data end                          // end of data section

code section execute              // define code section

// This function calculates a polynomial on all elements of an
// array x and stores the results in an array y
polyn function public
int64 r1 = address([x+arraysize*4]) // end of array x
int64 r2 = address([y+arraysize*4]) // end of array y
int64 r0 = arraysize*4             // array size in bytes = 400

for (float v0 in [r1-r0]) {       // vector loop
    float v0 = [r1-r0, length=r0] // read x vector
    float v1 = v0 * 0.5           // 0.5 * x
    float v1 -= 4                  // 0.5 * x - 4
    float v0 = v0 * v1 + 1        // (0.5 * x - 4) * x + 1
    float [r2-r0, length=r0] = v0 // save y vector
}
return                             // return from function
polyn end                           // end of function

code end                            // end of code section
```

While this looks very much like high-level language code, you have to explicitly specify which register to use for each variable, and you cannot put more code on one line than fits into a single instruction. In example 13.2, the line `float v0 = v0 * v1 + 1` is allowed because it fits the `mul_add2` instruction, but we cannot write `float v1 = v0 * 0.5 - 4` because an instruction cannot have two immediate constants. The line `int64 r1 = address([x+arraysize*4])` also fits a single instruction because `arraysize*4` can be calculated at assembly time (it involves only constants) and the result can be added to the relative address of `x` by the linker. The only thing the `address` instruction has to do at runtime is to add a constant to the datapointer.

More code examples are given in chapter 14

Command line

The command line for the assembler has the following format:

```
forw -ass assemblyfile objectfile [options]
```

The following options are supported:

-list=name	Make output list file. This is very useful for checking the generated code.
-O0	Optimization level 0: The assembler finds the smallest possible instruction that fits the specified operands. Two consecutive tiny instructions are joined together if possible.
-O1	Optimization level 1 (default): An instruction may be replaced by another instruction that does the same thing more efficiently. For example, <code>float v1 *= 4</code> can be replaced by <code>float v1 = mul_2pow(v1, 2)</code> . A conditional jump statement is merged with a preceding arithmetic instruction when possible. An <code>if</code> statement immediately followed by an unconditional jump is converted to a single conditional jump.
-O3	Optimization level 3. Enable optimizations that ignore subnormal numbers. For example, <code>float v1 *= 0.25</code> can be replaced by <code>float v1 = mul_2pow(v1, -2)</code> .
-maxerrors=n	Specify maximum number of error messages from the assembler.
-datasize=n	Specify maximum combined size of writeable static data sections. Static data can be accessed with 16 bit relative addresses if <code>datasize ≤ 32000</code> .
-codesize=n	Specify maximum combined size of code and read-only sections. Inter-module jumps and call tables can use 16 bit relative addresses when <code>codesize ≤ 131000</code> , and 24 bit relative addresses when <code>codesize ≤ 33000000</code> . Read-only static data can be accessed with 16 bit relative addresses if <code>codesize ≤ 32000</code> .
-ilist=name	Specify alternative instruction list name.

The preferred extension for file names in ForwardCom are `.as` for assembly files and `.ob` for object files.

File format

An assembly file can be in ASCII or UTF-8 format. An UTF-8 byte order mark is allowed at the beginning of the file, but not required.

Whitespace can be spaces or tabs. The use of tabs is discouraged because different editors may have different tabstops.

Linefeeds can be UNIX style (`\n`), Mac style (`\r`), or Windows style (`\r\n`). There is no limit to the line length.

Comments are C style: `/* */` or `//`

Nested comments are allowed. This makes it possible to comment out a piece of code that already contains comments.

Names of symbols

The names of data symbols, code labels, functions, etc. are case sensitive. A name can consist of letters a-z, A-Z, numbers 0-9, the special characters _ \$ @, and unicode letters. A name cannot begin with a number. There is no limit to the length of a name.

The names of keywords and instructions are not case sensitive. The following are reserved keywords:

```
align
break broadcast
capab0-capab31 case communal constant continue
datap do double
else end execute extern
fallback float float16 float32 float64 float128 for function
if in int8 int16 int32 int64 int128 ip
length limit
mask
options
perf0-perf31 pop public push
r0-r31 read
scalar section sp spec0-spec31 string switch sys0-sys31
threadp
uint8 uint16 uint32 uint64 uint128 uninitialized
v0-v31
weak while write
```

Sections

A section containing code or data is defined as follows:

```
name section options
...
name end
```

The following options can be defined for a section. Multiple options are separated by space or comma.

read	Readable data section.
write	Writeable data section.
execute	Executable code section. This does not imply read access. Write access is not allowed for executable sections.
ip	Addressed relative to the instruction pointer. This is the default for executable and read-only sections.
datap	Addressed relative to the data pointer. This is the default for writeable data sections.
threadp	Addressed relative to the thread data pointer. The section will have one instance for each thread.
communal	Communal section. Identical sections in different modules can be merged together. Unreferenced sections can be removed.
uninitialized	Data section containing only zeroes. The data of this section does not take up space in object files.
align=n	Align the beginning of the section to an address divisible by n, which must be a power of 2. The default alignment for executable sections is 4. The default alignment for a data section is the size of the biggest data type in the section. A higher alignment can be specified with the align directive.

Sections with the same name are placed consecutively in the executable file. They must have the same attributes, except for alignment.

Sections with different names but the same attributes (except for alignment) are placed in alphabetical order in the executable file.

Constant expressions

Integer constants can be expressed in the following ways:

decimal numbers	contains only digits 0-9 Numbers beginning with 0 are interpreted as decimal, not octal.
binary numbers	0b followed by digits 0-1
octal numbers	0O followed by digits 0-7
hexadecimal numbers	0x followed by digits 0-9, a-f
character constants	1-8 ASCII characters enclosed in single quotes ' '. The first character will be contained in the lowest byte of a 64 bit integer. For example 'AB' = 0x4241
imported constant	extern name: constant
difference between addresses	symbol1 - symbol2. The two symbols must have the same base pointer, either ip, datap, threadp, or none.

Integer constants can be combined with all common operators:

+ - * / % & | ^ ~ && || ! << >> >>> < <= > >= == !=
?:

The operators have the same order of precedence as in C. The result is calculated as signed 64-bit integers, except for >>> which is an unsigned (logical) shift right.

Imported constants and differences between addresses cannot be allowed in general calculations. The only operations allowed for these are addition of a local constant, and division by a power of 2. These calculations are done by the linker except for a difference between two local symbols in the same section.

Floating point constants must contain a dot or an E and at least one digit, for example 1.23E-4

Floating point expressions are calculated with double precision. The following operators can be used:

+ - * / < <= > >= == != ?:

String constants are sequences of ASCII or UTF-8 characters enclosed in " ".

The following escape sequences are recognized: \\ \n \r \t \"

String constants can be concatenated with the + operator like in Java.

Data types

The following data types are used in data definitions and instructions:

int8	8 bit signed integer
uint8	8 bit unsigned integer
int16	16 bit signed integer
uint16	16 bit unsigned integer
int32	32 bit signed integer
uint32	32 bit unsigned integer
int64	64 bit signed integer
uint64	64 bit unsigned integer
int128	128 bit signed integer (optional)
uint128	128 bit unsigned integer (optional)
float16	half precision floating point (limited support)
float	single precision floating point
float32	single precision floating point
double	double precision floating point
float64	double precision floating point
float128	quadruple precision floating point (optional)

A '+' after an integer type indicates that the assembler can use a bigger type if this makes the instruction smaller or more efficient, for example int16+.

Data definitions

Static data can be defined inside a section. Several different forms are allowed:

Assembly style data definition: label : datatype value, value, ...
C style data definition: datatype name = value, name = value, ...
C style array definition, uninitialized: datatype name[number]
C style array definition, initialized: datatype name[number] = {value, value, ...}
Assembly style string definition: label : int8 "string", "string", ...
C style string definition: int8 name = "string"

A terminating zero after a string must be added explicitly if needed.

Examples:

```
mydata section read write datap
alpha: float 0.1, 0.2, 0.3
int16 beta = 4, gamma = 5
int8 delta[25]
align (8)
int8 epsilon[] = {6, 7, 8, 9}
zeta: int8 "Dear reader", 0
int8 eta = "Nice to meet you"
mydata end
```

Function definitions and labels

A function can be defined inside an executable section. It is defined like this:

```
name function attributes
...
name end
```

Attributes can be 'public' and 'weak'. The function will be local if no attributes are specified and the name does not appear in a public declaration.

Example:

```
mycode section execute
// this function calculates the square of a double
square function
double v0 *= v0
return
square end
mycode end
```

The calling conventions for functions are defined in chapter 12.4.

Instructions

It is convenient to have only one instruction per line. Multiple instructions on the same line must be separated by semicolon.

Instructions can be defined inside an executable section. The general form looks like this:

```
label : datatype destination = instruction(source_operands), options
```

The destination is a register in most cases. A few instructions allow a memory operand as destination. The source operands can be registers, memory operands, or immediate constants. No instruction can have more than one memory operand. Only a few instructions can have multiple immediate constants or both a memory operand and an immediate constant.

The following options are possible:

```
options = integer constant
mask = register
fallback = register
fallback = 0
```

Only certain instructions can have options=constant.

All multi-format instructions can have a mask. Single-format instructions with A or E templates can also have a mask. Jump instructions cannot have a mask.

The mask register can be r0-r6 or v0-v6. The fallback register can be r0-r30 or v0-v30. The mask and fallback registers must be the same type of register as the destination (general purpose or vector register). The fallback specifies the result when bit 0 of the mask is zero.

The default fallback register is the same as the first source register. A different fallback register is possible only if there is a vacant register field in the code template. The fallback register cannot be different from the first source register in the following cases:

- the instruction has three source operands
- the instruction needs 64 bits for an immediate constant
- the instruction has a memory operand with index
- the instruction has vector registers and a memory operand

Instructions can be written simply with an operator instead of the instruction name if the instruction does the same as the operator. The general form is:

```
label : datatype destination = operand1 operator operand2
```

In this case, you can specify mask and fallback with the ?: operator:

```
datatype destination = mask ? operand1 operator operand2 : fallback
```

A memory operand must be enclosed in square brackets. Memory references are always relative to a base pointer. A memory operand can contain:

A base pointer. This is a general purpose register or a special pointer (ip, datap, threadp, sp).
A scaled index. This is a general purpose register multiplied by a scale factor. The scale factor is the same as the operand size in most cases. Instructions with a general purpose register destination can also have a scale factor of 1. Instructions with a vector register destination can also have a scale factor of -1.
An offset. This is an integer constant.
The name of a data definition. The base pointer is implicit when there is a data name.
A limit for the index: limit=constant.
Memory operands in vector instructions must have one of the following options:
scalar
length=register
broadcast=register

Examples:

```
LABEL1 :
int32 r1 = add(r2, r3)
int32 r1 = r2 + r3           // same
int32 r1 = r1 + 1
int32 r1 += 1                // same
int32 r1 ++                  // same
int32 r1 = add(r2, r3), mask = r4, fallback = r5
int32 r1 = r4 ? r2 + r3 : r5 // same
int32 r1 = r2 + [r3 + 0x10]
int32 r1 = r2 + [r3 + 4*r4, limit = 100]
int32 v1 = v2 + [r3 - r4, length = r4]
float v1 = v2 + [alpha, scalar]
float v1 = v2 + [alpha, broadcast=r3]
float v1 = -v2
double v1 = mul_add(v2, v3, 1.5), options=0xC
double v1 = v2 - v3 * 1.5    // same
```

The details for each instruction are described in chapter 5.

Unconditional and indirect jumps, calls, and returns

Direct jumps, calls, and returns:

```
jump target
call target
return
```

Indirect jumps and calls to a register or memory operand containing an absolute 64-bit address:

```
jump (register)
call (register)
jump ([base_register+offset])
call ([base_register+offset])
```

Indirect jumps and calls to a pointer relative to an arbitrary reference point:

```
datatype jump (reference_register, [base_register+offset])
datatype call (reference_register, [base_register+offset])
datatype jump (reference_register, [base_register+index_register*scale])
datatype call (reference_register, [base_register+index_register*scale])
```

The datatype specifies the size of the relative pointer stored in memory. This must be a signed integer type. The reference point must be contained in a 64-bit register.

Examples:

```
extern function1: function, function2: function

rodata section read ip // read-only data section, ip based
// table of addresses relative to the table itself:
align (4)
calltable:
int16 (function1-calltable) / 4 // relative address of function1
int16 (function2-calltable) / 4 // relative address of function2
rodata end

code section execute ip
function0 function
    call function1 // calls function1
    int64 r1 = address([function1])
    call (r1) // calls function1
    int64 r1 = address([calltable])
    int16 call (r1, [r1+2]) // calls function2
    int16 jump (r1, [r1+r0*2]) // jumps to function selected by r0
    return
function0 end
code end
```

Conditional jumps and loops

Conditional jumps always involve an arithmetic or logic operation and a jump conditional upon the result:

```
datatype destination = instruction(source_operands), jump_condition target
```

Examples:

```
int32+ r1 = add(r1, r2), jump_pos L1
uint64    compare(r2, 5), jump_pos L1
float     test_bit(v1, 31), jump_nzero L1
L1:
```

If there are two source registers then the first source register must be the same as the destination register. Vector registers can be used only when the data type does not fit into a general purpose register. Floating point addition and subtraction cannot be used with conditional jumps.

The most common conditional jumps can be coded more conveniently using the high-level language keywords: if, else, for, do, while, break, continue.

An 'if' statement can have the form:

```
if (datatype register operator register_or_constant) {...} else {...}
```

Line breaks are allowed before and after { and }. The curly brackets cannot be omitted.

Comparing a register value with another register or a constant is done with one of the operators:

< <= > >= == !=

It is also possible to test a single bit in a register with the & operator and a constant with only one bit set (in other words, a power of 2).

Examples:

```
if (int32+ r1 >= r2) {
    nop // r1 >= r2
} else { // r1 < r2
    if (int64 !(r3 & (1 << 20))) {
        nop // bit 20 in r3 is not set
    }
}
```

The conditions cannot be combined with && || etc. because the statement must fit into a single instruction.

The 'while', 'do-while', and 'for' loops are written in the same way as if statements:

```
while (datatype condition) {...}
do {...} while (datatype condition)
for (datatype initialization; condition; increment) {...}
```

Examples:

```
for (int32+ r1 = 1; r1 <= 100; r1++) {
    int64 r2 = 4 // executed 100 times
    while (int64 r2 > 0) {
        int64 r2-- // executed 400 times
    }
}
do {
    int32 r2 = r1 - 1 // executed once
} while (int32 r2 > r1) // never true
```

The initialization and increment instructions in the 'for' loop can be anything that fits into a single instruction with the specified data type.

Vector loops are written in the following way:

```
for (datatype vector_register in [end_pointer - index_register]) {...}
```

Example:

```
int64 r1 = address([my_array]) // address of my_array
int64 r2 = my_arraysize * 8 // size of my_array in bytes
int64 r1 += r2 // point to end of my_array
for (float v1 in [r1 - r2]) { // loop through my_array
    float v1 = [r1 - r2] // read vector of maximum length
    float v1 *= 2 // multiply values by 2
    float [r1 - r2] = v1 // store results in my_array
}
```

```
}
```

This loop will subtract the maximum vector length from r2 for each iteration and continue as long as r2 (interpreted as a signed 64 bit integer) is positive. The maximum vector length may depend on the data type. The loop will use the maximum vector length for the data type specified in the for statement (float in this case). It is possible to use more vector registers and more end-of-array pointers than the ones specified in the 'for' statement.

It is recommended to use optimization level 1 or higher when assembling branches and loops.

Imports and exports

A function, data symbol, or constant defined in one assembly module can be accessed from another module if it is exported from the first module and imported to the second module. The necessary cross references are calculated and inserted by the linker.

Symbols are imported as follows:

```
extern symbolname : attributes, symbolname : attributes, ...
```

The following attributes can be specified:

function	the symbol is an executable function
ip	the symbol is a data object addressed relative to the ip pointer
datap	the symbol is a data object addressed relative to the datap pointer
threadp	the symbol is a data object addressed relative to the threadp pointer
constant	the symbol is a constant with no address
read	the symbol is in a readable data section
write	the symbol is in a writeable data section
execute	the symbol is executable code
data type	the data type for the data symbol
weak	weak linking: the symbol will be resolved only if it exists

An external symbol must have one, and only one, of the following attributes: function, ip, datap, threadp, constant. The other attributes are optional. Multiple attributes are separated by space.

A weak external symbol will only be linked if it exists. The linker will not search function libraries to find the weak symbol, but the symbol will be resolved if it exists in a module that is linked in for other reasons. A weak external constant or data symbol that is not resolved will be zero. A weak function that has not been resolved will return zero.

Symbols are exported as follows:

```
public symbolname : attributes, symbolname : attributes, ...
```

The possible attributes are the same as for external symbols. It is not necessary to specify the attributes because the attributes of locally defined symbols are already known.

It is convenient to place extern and public declarations in the beginning of the assembly file. Forward references are allowed.

Weak public symbols work as follows. If more than one module contains a weak public symbol with the same name then the linker will not issue an error message but link the first symbol. If there is one occurrence of a non-weak symbol with this name then the non-weak symbol will be chosen. If there is more than one non-weak public symbol with the same name then the linker will issue an error message.

Other directives

The 'align' directive can align data or code:

```
align (number)
```

The number must be a power of 2. The next data or code will be aligned to an address divisible by the specified number. The beginning of the section will automatically be aligned to at least the same size.

13.2 Metaprogramming

Metaprogramming means variables and instructions that do not form part of the final executable code but are useful for controlling the assembly process.

Traditional assemblers use macros for this purpose. The syntax for this is often confusing, especially when macros are nested.

Instead, the ForwardCom assembler will implement metaprogramming involving the following features:

- variables that are used only during the assembly process
- include files
- assembly-time branches and loops
- assembly-time functions
- generate a text string and emit this string as assembly code

Metaprogramming commands are indicated by a percent sign at the beginning of the line. At present, only variables are implemented.

Metaprogramming variables

Metaprogramming variables are defined and assigned on a separate line in the following way:

```
% name = value
```

Meta-variables are weakly typed. They can have one of the following types:

integer	evaluated as signed 64-bit integer
floating point	evaluated as double precision
string	ASCII or UTF-8 text string
register	the variable can be an alias for any register
memory operand	the variable can be an alias for any memory operand

It makes no difference whether this meta-code is inside a section or not.

Meta-variables can be redefined or reassigned at any time. Meta-code is sequential so that the same variable can have different values at different places in the code.

An integer meta-variable can be exported as a constant with a public declaration if it has only one value. This is the only case where a forward reference to a meta-symbol is allowed.

While general assembly code allows forward references to data and code labels, meta-code cannot in general have forward references.

Example:

```
% A = 5           // meta-variable integer A = 5
% R = r1         // alias for register r1
int32 R = A      // r1 = 5
% A++           // change value of A to 6
int32 R = A      // r1 = 6
```

13.3 Disassembler

The disassembler can convert object files and executable files back to assembly language. The command line for the disassembler has the following format:

```
forw -dis inputfile outputfile [options]
```

The following options are supported:

-ilist=name	specify alternative instruction list name.
-------------	--

The disassembler produces output lines that may look like this example:

```
float v1 = add(v2, 2.5) // 002C _ 227_E 08.0 5 01.02.02.02 _ 4100 00
```

The comment is interpreted as follows:

002C	hexadecimal address relative to the beginning of the current section
227.E	il, mode, and mode2 in the E2 format template. The last digit can indicate various kinds of subdivision of the code format
08.0	operation code OP1 and OP2
5	operand type. 5 means float
01.02.02.02	register fields RD, RS, RT, RU
-	the mask field is unused in this case
4100	the 16-bit data field IM2 contains the value 2.5 in half precision. If the value had been 2.6 we would need full precision and another code format
00	the IM3 field is unused in this case

The disassembler is used internally to generate a list output for the assembler. The only difference between a disassembly and an assembly listing is that the names of local labels are preserved in the assembly listing while these names may have been lost in a disassembly.

13.4 Linker

The linker joins object files and library files into an executable file.

The command line for the linker has the following format:

```
forw -link executablefile objectfiles [options]
```

The preferred extension for file names in ForwardCom are .ob for object files and .ex for executable files.

The linker is not implemented yet.

13.5 Library manager

A function library is a collection of object files each defining one or more functions that can be called from other modules. The library manager can make a function library and add or remove modules from it.

The command line for the library manager has the following format:

```
forw -lib libraryfile [options] objectfiles
```

The following options are supported:

-a	(Default). Add the following object files to the library. Any existing object file with the same name will be replaced.
-d	Delete the following object files from the library.
-x	Extract the following object files from the library. The library itself is not modified by this command.
-l	List object files and their public symbols.

The preferred extension for file names in ForwardCom are .ob for object files and .li for library files.

The library manager is not implemented yet.

13.6 Emulator

The emulator can execute a ForwardCom executable file on another platform. It is useful for testing and debugging.

The command line for the emulator has the following format:

```
forw -emu executable_file [options]
```

The emulator is not implemented yet.

13.7 Dump utility

The dump utility can show metadata from object files and executable files.

The command line for the dump utility has the following format:

```
forw -dump-options object_file
```

The following options are supported:

f	file header.
h	section headers.
s	symbol table.
n	string table.
r	relocation records.

13.8 Compiling the forw tools

These tools can be compiled for Windows, Linux, MacOS, and other platforms.

Compiling for Windows with MS Visual Studio: Use the project file forw.vcxproj.

Compiling with Gnu C++ compiler: Use the makefile.

Other compilers: Make a project containing all the .cpp files. Compile for console mode, preferably 64 bits. The platform must have little endian memory organization.

Chapter 14

Code examples

This chapter contains examples of assembly code to illustrate the features of the ForwardCom instruction set. The syntax for assembly language is described in chapter 13.1. The function calling conventions are described in chapter 12.4.

Horizontal vector add

ForwardCom has no instruction for adding all elements of a vector because this would be a complex instruction with variable latency depending on the vector length.

The sum of all elements in a vector can be calculated by repeatedly adding the lower half and the upper half of the vector. This method is illustrated by the following example, finding the horizontal sum of a vector of single precision floats.

Example 14.1.

```
v0 = my_vector           // we want the horizontal sum of this
int64 r0 = get_len(v0)   // length of vector in bytes
int64 r0 = round_u2(r0)  // round up to nearest power of 2
float v0 = set_len(r0,v0) // adjust vector length
while (uint64 r0 > 4) {  // loop to calculate horizontal sum
    uint64 r0 >>= 1      // the vector length is halved
    float v1 = shift_reduce(r0,v0) // get upper half of vector
    // the result vector has the length of the first operand:
    float v0 = v1 + v0   // Add upper half and lower half
}
// The sum is now a scalar in v0
```

Horizontal vector minimum

The same method can be used for other horizontal operations. It may cause problems that the `set_len` instruction inserts elements of zero if the vector length is not a power of 2. Special care is needed if the operation does not allow extra elements of zero, for example if the operation involves multiplication or finding the minimum element. A possible solution is to mask off the unused elements in the first iteration. The following example finds the smallest element in a vector of floating point numbers:

Example 14.2.

```
v0 = my_vector           // find the smallest element in this
r0 = get_len(v0)         // length of vector in bytes
int64 r1 = round_u2(r0)  // round up to nearest power of 2
```

```

uint64 r1 >>= 1          // half length
v1 = shift_reduce(r1,v0) // upper part of vector
int64 r2 = r0 - r1       // length of v1
float v0 = set_len(r1,v0) // reduce length of v0
// make mask for length of v1 because the two operands may
// have different length
int64 v2 = mask_length(r2, v0, 0), options=4
// Get minimum. Elements of v0 fall through where v1 is empty
float v0 = min(v0, v1, mask=v2, fallback=v0) // minimum
// loop to calculate the rest. vector length is now a power of 2
while (uint64 r1 > 4) {
    // Half vector length
    uint64 r1 >>= 1
    // Get upper half of vector
    float v1 = shift_reduce(r1, v0)
    // Get minimum of upper half and lower half
    float v0 = min(v1, v0) // has the length of the first operand
}
// The minimum is now a scalar in v0

```

Switch-case statement

A switch-case multiway branch can be implemented efficiently with a table of relative addresses. This table is placed in read-only memory for security reasons.

Example 14.3.

```

/* C code:
int i, x;
switch (i) {
case 1:
    x = 10; break;
case 2:
    x = 12; break;
case 5:
    x = 20; break;
default:
    x = 99; break;
}
*/
// ForwardCom code:

rodata section read ip
// table of relative addresses, using DEFAULT as reference point
align (4)
jumptable: int16 0          // case 0
int16 (CASE1 - DEFAULT) / 4 // case 1
int16 (CASE2 - DEFAULT) / 4 // case 2
int16 0                    // case 3
int16 0                    // case 4
int16 (CASE5 - DEFAULT) / 4 // case 5
rodata end

```

```

code section execute ip
// r0 = i
// r1 = x
// test if i is outside of the range
// use unsigned test to avoid testing for r0 < 0
if (uint32 r0 > 5) {
    jump DEFAULT
}
int64 r2 = address([jumptable])
int64 r3 = address([DEFAULT])
// relative jump with r3 = DEFAULT as reference point,
// r2 as table base and r0 as index
int16 jump (r3, [r2 + r0 * 2])

CASE1:
    int32+ r1 = 10
    jump FINISH
CASE2:
    int32+ r1 = 12
    jump FINISH
CASE5:
    int32+ r1 = 20
    jump FINISH
DEFAULT:
    int32+ r1 = 99
FINISH:

code end

```

Boolean operations

Boolean combinations of conditions can be implemented with branches as shown in this example.

Example 14.4.

```

/* C code:
float condfunc (float a, float b) {
    if (a >= 0 && (a < 20 || a == b)) {
        a = sqrt(a);
    }
    return a;
}
*/

// ForwardCom code:

code section execute ip
extern _sqrtf : function

// v0 = a, v1 = b
_condfunc function public

```

```

if (float v0 >= 0) {
    if (float v0 < 20) {jump L1}
    if (float v0 == v1) {
        L1:
        call _sqrtf
    }
}
return // return value is in v0
_condfunc end

code end

```

Branches can be quite slow, especially if they are poorly predictable. It is often faster to generate boolean variables for each condition and use bit operations to combine them. This corresponds to replacing `&&` with `&` and `||` with `|`. The code below shows the same example where three conditional jumps are reduced to one conditional jump and two bit operations. Note that this transformation is not valid if the evaluation of unused conditions has side effects.

Example 14.5.

```

_condfunc2 function public
float v2 = v0 >= 0 // boolean variable for a >= 0
float v3 = v0 < 20 // boolean variable for a < 20
float v4 = v0 == v1 // boolean variable for a == b
int32+ v3 |= v4 // a < 20 || a == b
int32+ v2 &= v3 // a >= 0 && (a < 20 || a == b)
if (float v2 & 1) { // test bit 0 of boolean v2
    call _sqrtf
}
return
_condfunc2 end

```

We can reduce the number of instructions and make the code still faster by using a special feature of the compare instruction that uses the mask and fallback registers as extra boolean operands:

Example 14.6.

```

_condfunc3 function public
float v2 = v0 >= 0 // boolean variable for a >= 0
float v3 = v0 == v1 // boolean variable for a == b
// use mask and fallback register as extra boolean operands
float v4 = compare(v0, 20), options=0x22, mask=v2, fallback=v3
if (float v4 & 1) { // test bit 0 of boolean v4
    call _sqrtf
}
return
_condfunc3 end

```

Virtual functions

Virtual functions are used in C++ for polymorphous classes. This example shows how to implement a virtual class in ForwardCom. We can save space by using 32-bit relative pointers rather than 64-bit absolute pointers as other systems do.

Example 14.7.

```
/* C++ code:

class VirtClass {
public:
    // constructor:
    VirtClass() {x = 0;}
    // virtual functions:
    virtual void func1() {x++;}
    virtual int  func2() {return x;}
protected:
    int x;
};

int test() {
    VirtClass obj;          // create object
    obj.func1();           // call virtual function 1
    return obj.func2();    // call virtual function 2
}
*/

// ForwardCom code:
rodata section read ip align = 4
// table of virtual function pointers for VirtClass
// with REFPOINT as reference point
VirtClasstable:
int32 (VirtClass_func1 - REFPOINT) / 4
int32 (VirtClass_func2 - REFPOINT) / 4
rodata end

code section execute ip
// choose any reference point for the relative pointers,
// for example the beginning of the code section:
REFPOINT: nop

VirtClass_constructor function public
// The pointer 'this' is in r0
// At [r0] is a relative pointer to VirtClasstable,
// next the class data members, in this case: x
int32 r1 = VirtClasstable - REFPOINT
int32 [r0] = r1
int32 [r0+4] = 0
return
VirtClass_constructor end

VirtClass_func1 function
// The pointer 'this' is in r0
// x is in [r0+4]
int32 r1 = [r0+4]
int32 r1++
int32 [r0+4] = r1
return
VirtClass_func1 end
```

```

VirtClass_func2 function
int32 r0 = [r0+4]
return
VirtClass_func2 end

_test function public
push (r16)          // save register
// get the address of the reference point
int64 r16 = address([REFPOINT])
// the object 'obj' uses 8 bytes, allocate space on the
// stack for it
int64 sp -= 8
// call the constructor. The 'this' pointer must be in r0
int64 r0 = sp
call VirtClass_constructor
// r0 still points to the object because a constructor
// always returns a reference to the object.
// Get the address of the virtual table
int32 r1 = sign_extend_add(r16, [r0])
// call VirtClass_func1 as the first table entry
int32 call (r16, [r1])
// get a pointer to the object again
int64 r0 = sp
// Get the address of the virtual table
int32 r1 = sign_extend_add(r16, [r0])
// call VirtClass_func2 as the second table entry
int32 call (r16, [r1+4])
// release space allocated for 'obj'
int64 sp += 8
pop (r16)           // restore register
// the return value from callVirtClass_func2 is in r0
return
_test end

code end

```

Memory copying

The standard memcpy function can be implemented efficiently as a vector loop.

Example 14.8.

```

// C function:
// void *memcpy(void *destination, const void *source, size_t n);

_memcpy function public
// r0 = destination
// r1 = source
// r2 = n
int64 r3 = r0 + r2          // end of destination
int64 r1 += r2              // end of source
// the type int32 has the longest maximum vector length

```

```

// make vector loop with vector of int32
for (int32 v0 in [r1-r2]) {
    int32 v0 = [r1-r2, length = r2] // read from source
    int32 [r3-r2, length = r2] = v0 // write to destination
}
// destination is returned unchanged in r0
return
_memcpy end

```

String length

The standard `strlen` function finds the length of a zero-terminated string.

The ForwardCom standard specifies extra space at the end of user data to make it possible to read beyond the end of a data object of unknown size (see page 105). This is convenient here because we can read the string as a full vector before we know the length.

Example 14.9.

```

// C function:
// size_t strlen (const char * str);

_strlen function public
// r0 = str
// start at nearest preceding 16-bytes boundary for efficiency
int64 r1 = r0 & -16 // 16-bytes boundary
int64 r2 = r0 - r1 // length of unused part
int8 v0 = broadcast_max(0) // vector of zeroes to compare with
int64 r3 = get_len(v0) // maximum length
int64 v1 = mask_length(r2, v0, 0), options=1 // mask off unused part
int8 v2 = [r1, length = r3] // read vector
int8 v3 = v1 ? (v0 == v2) : v0 // compare with 0, masked
int8 v4 = bool_reduce(r3, v3) // horizontal OR is in bit 1
// loop as long as no zero is found
// (vector register can only be tested as float)
while (float !(v4 & 2)) {
    int64 r1 += r3 // next block of maximum length
    int8 v2 = [r1, length = r3] // read vector
    int8 v3 = (v0 == v2) // compare with 0, no mask now
    int8 v4 = bool_reduce(r3, v3)
}
// now v3 contains a bit index to the end of the string
int64 r2 = vec2gp(v3) // move to general purpose register
while (int64 r2 == 0) {
    // nothing in first 64 bits of v3. get next 64 bits
    int64 r2 = 8
    int64 r1 += r2
    int8 v3 = shift_reduce(r2, v3)
    int64 r2 = vec2gp(v3)
}
// get index to first bit
int64 r2 = bitscan_f(r2)
// add difference between current block start and string start

```

```

int64 r0 = r2 + r1 - r0
// the string length is returned unchanged in r0
return
_strlen end

```

High precision arithmetic

Function libraries for high precision arithmetic typically use a long sequence of add-with-carry instructions for adding integers with a very large number of bits. A more efficient method for big number calculation is to use vector addition and a carry-look-ahead method. The following algorithm calculates $A + B$, where A and B are big integers represented as two vectors of $n \cdot 64$ bits each, where $n < 64$.

Example 14.10.

```

uint64 v0 = A // first vector, n*64 bits
uint64 v1 = B // second vector, n*64 bits
uint64 v2 = carry_in // single bit in vector register
uint64 v0 += v1 // sum without intermediate carries
uint64 v3 = v0 < v1 // carry generate = (SUM < B)
uint64 v4 = v0 == -1 // carry propagate = (SUM == -1)
uint64 r0 = get_len(v0) // length of vector in bytes
uint64 v3 = bool2bits(r0,v3) // compressed to bitfield
uint64 v4 = bool2bits(r0,v4) // compressed to bitfield
// calculate propagated additional carry:
// CA = CP ^ (CP + (CG<<1) + CIN)
uint64 v3 <<= 1 // shift left carry generate
uint64 v2 = v2 + v3 + v4
uint64 v2 ^= v4
uint64 v1 = bits2bool(r0,v2) // expand additional carry to vector
uint64 v0 += v1 // add correction to sum
uint64 r0 >>= 3 // n = number of elements in vectors
uint64 v3 = gp2vec(r0) // copy to vector register
uint64 v2 >>= v3 // carry out
// v0 = sum, v2 = carry out

```

If the numbers A and B are longer than the maximum vector length then the algorithm is repeated. If the vector length is more than $64 * 8$ bytes then the calculation of the additional carry involves more than 64 bits, which again requires a big number algorithm.

Matrix multiplication

Matrix operations can be difficult because they involve a lot of permutations. The following example shows the multiplication of two 4×4 matrixes of floating point numbers, assuming that the vector registers are long enough to contain an entire matrix.

Example 14.11.

```

float v1 = first_matrix // first matrix, 4x4 floats
float v2 = second_matrix // second matrix, 4x4 floats
int64 r1 = 64 // size of entire matrix in bytes
int64 r2 = 1 // shift count, elements

```

```

int64 r3 = 4 // shift count, elements
float v0 = replace(v1,0) // make a matrix of zeroes
for (int64 r0 = 0; r0 < 4; r0++) { // repeat 4 times
    float v3 = repeat_within_blocks(r1,v1,16) // repeat column
    float v4 = repeat_block(r1,v2,16) // repeat row
    float v0 = v0 + v3 * v4 // multiply rows and columns
    float v1 = shift_down(r2,v1) // next column
    float v2 = shift_down(r3,v2) // next row
}
// Result is in v0.

```

You may roll out the loop and calculate partial sums separately to reduce the loop-carried dependency chain of v0

14.1 Optimization tricks

The ForwardCom system is designed with high performance as a top priority. The following guidelines may help programmers and compiler makers obtain optimal performance.

Minimize instruction size

The assembler will automatically pick the smallest possible version of an instruction. Instructions can have different versions with 4-bit, 8-bit, 16-bit, 32-bit, and 64-bit integer constants. A large integer constant with few significant bits can be represented as a smaller constant with a left shift. For example, the constant 0x5000000000 can be represented as $5 \ll 36$. The assembler will do this automatically when possible. Small floating point constants with no decimals can be represented as 8-bit signed integers. Simple rational numbers where the denominator is a power of 2 can be represented as half precision without loss of precision. The assembler does this automatically, too. It is recommended to check the output listing from the assembler to see how much space each instruction takes. Sometimes, you can reduce the instruction size by simple changes in the code. Many instructions will be smaller if the first source register is the same as the destination register. You can also see from the output listing if there are unpaired tiny instructions that can be paired if you reorder the instructions.

Optimize jumps

The assembler will merge a jump with a preceding arithmetic instruction if possible, unless optimization is turned off. For example, an integer addition followed by a conditional jump that compares the result with zero may be merged into a single instruction. This is only possible when a number of conditions are fulfilled:

- the two instructions have the same data type
- the destination of the arithmetic instruction is the same register as the source of the branch instruction
- the arithmetic instruction uses the same register for source and destination
- there are no other instructions between the two, and no jump label.

The output listing will show if the two instructions have been merged.

Optimization of jumps to jumps, etc., is generally the responsibility of the compiler. The assembler will do only a few simple jump optimizations.

Optimize cache use

Memory and cache throughput is often a bottleneck. You can improve caching in several ways:

- Optimize code caching by minimizing instruction sizes, pairing tiny instructions, and inlining functions.
- Optimize data caching by embedding immediate constants in instructions.
- Use register variables instead of saving variables in memory.
- Use registers as function parameters.
- Avoid spilling registers to memory by using the information about register use in object files, as described on page 121.

Use general purpose registers for control code

Any variables that control program execution, such as branch conditions and loop counters, should preferably be stored in general purpose registers rather than memory or vector registers. This may enable the microprocessor to resolve branches early and prefetch the code after the branch earlier.

Use the vector loop feature

Array loops are particularly efficient if the vector loop feature described on page 8 is used. Loops containing function calls can be vectorized if the functions allow vector parameters.

Avoid long dependency chains

ForwardCom may be implemented on a superscalar processor that can execute multiple instructions simultaneously. This works most efficiently if the code does not contain long dependency chains.

Chapter 15

Conclusion

The proposed ForwardCom instruction set architecture is a consistent, modular, flexible, orthogonal, scalable, and expandable instruction set offering a good compromise between the RISC principle that gives fast decoding and efficient pipelining, and the CISC principle that gives a more compact code and more work done per instruction. Support for efficient out-of-order execution and vector processing is a basic part of the design rather than a suboptimal patch added later as we have seen in other systems.

There are relatively few instructions, but each instruction can be coded in many different variants with integer operands of different sizes and floating point operands of different precisions. The operands can be scalars or vectors of any length. Operands can be registers, immediate constants in various compact forms, or memory operands with different addressing modes. Instructions can have predicates or masks with specified fallback values. Vectors have variable lengths with unlimited room for future expansions.

All in all, the same basic instruction can have many different variants with the same operation code where other instruction sets have many different instructions to cover the same diversity. Everything fits into a consistent template format that simplifies the hardware implementation. The design also has plenty of space for single-format instructions with fewer variants.

The instructions are designed so that the microprocessor pipeline can be simple and efficient. All instructions fit into the same pipeline structure to facilitate an efficient hardware design.

The decoder front-end can load multiple instructions per clock cycle because it is easy to detect the length of each instruction, and the decoder needs only distinguish between a few different instruction sizes. Actually, the only instruction size that must be supported is single-word. It is possible to make a working program with only single-word (32 bits) instructions, but it is highly recommended to also support double-word and triple-word instructions. Tiny instructions (two in one code word) are useful for making the code more compact.

It is possible to add support for longer instructions in future extensions, but the priority has been to avoid any bottleneck in the decoding of instruction length (which is a serious bottleneck in the x86 architecture).

The code format is designed to be compact in order to save code cache space. This compactness is obtained in several ways. The same instruction can be coded in different sizes with two- and three-operand forms, different sizes of immediate constants, shifted immediate constants, and relative addresses with different sizes of offsets and scale factors, while avoiding absolute addresses that would require 64 bits for the address alone. It is always possible to choose the smallest version of an instruction that fits the particular need. The load on the data cache can be reduced by storing immediate constants in the code rather than in memory operands.

Most instructions can have a mask register which is used for predication in scalar instructions and masking in vector instructions. The same mask register is also used for specifying various options such as rounding mode, exception handling, etc., that would otherwise require extra bits in the in-

struction code.

The introduction of vector registers with variable length is an important improvement over the most common current architectures. The ForwardCom vector system has the following advantages:

- The system is scalable. Different microprocessors can have different maximum vector lengths with no upper limit. It can be used for small embedded systems as well as large supercomputers with very long vectors.
- The same code can run on different microprocessors with different maximum vector lengths and automatically utilize the full vector capabilities of each microprocessor.
- The code does not have to be recompiled when a new microprocessor version with longer vectors becomes available. Software developers do not have to maintain multiple versions of their software for different vector lengths.
- The software can save and restore a vector register in a way that is guaranteed to work with future processors with longer vectors. The inability to do so is a big problem in current architectures.
- Only the part of a vector register that is actually used needs to be saved and restored. Each vector register includes information about how many bytes of it are used. Therefore, no unnecessary resources are wasted on saving a full-length vector if it is unused or only partially used.
- A special addressing mode supports a very efficient loop structure that will automatically use the maximum vector length on all but the last iteration of an array loop. The last iteration will automatically use a shorter vector to handle the remaining array elements in case the array size is not divisible by the maximum vector length. There is no need to handle the remaining elements separately outside the main loop and no need to make separate versions of the loop for different special cases.
- Functions can have variable-length vector registers as parameters. This makes it easy for the compiler to vectorize loops that contain function calls.
- Instructions with vector register operands need no extra information about the vector length because this information is included in the vector registers. This makes these instructions more compact. Instructions with vector memory operands do need this extra information, though.
- The system takes into account the hardware requirements of microprocessors with very long vectors where transport delays across a vector may depend on the vector length.
- Strong security features are built into the design.
- Software standards guarantee compatibility between different compilers, programming languages, user interface frameworks, and operating systems.

The memory model is flexible with relative addresses. Everything is position-independent. Memory management is simpler than in many current systems with less need for virtual address translation. There is no translation lookaside buffer (TLB) and no memory paging, but a simple on-chip memory map. Problems with stack overflow, memory fragmentation, etc. can be avoided completely in most cases. Task switches will be fast because of the small memory map and because of the efficient mechanism for saving vector registers.

The principle that a fundamental redesign enables us to learn from history and integrate late additions into the basic design also applies to the whole ecosystem of ABI standard, function libraries, compilers, linkers, and operating system. By defining not only an instruction set, but also ABI standard, binary file format, interface library standard, etc. we get the further advantage that different compilers and different programming languages will be compatible with each other. It will be possible to write different parts of a program in different programming languages and to use the same function libraries

with all compilers. Even different operating systems will be compatible to some degree. A feature for re-linking an executable file makes it possible to run the same binary program file in different operating systems or on different platforms where the appropriate user interface framework is selected when the program is installed.

We have also learned from past mistakes that it is difficult to predict future needs. While the ForwardCom instruction set is intended to be flexible with room for future extensions, we may ask whether the future will bring needs for new features that are difficult to integrate into our design and standards. The best way to prevent such unforeseen problems is to allow input and suggestions from the entire community of hardware and software developers. It is important that the design and standards are developed through an open process that allows everybody to comment and make suggestions. We have already seen the problems of leaving this to a commercial industry. The industry often makes short-term decisions for marketing reasons. Patents, license restrictions, and trade secrets harm competition and prevent niche operators from entering the market. The microprocessor industry keeps new features and instruction set extensions secret for competitive reasons until it is too late to change them in case the IT community comes up with better proposals.

The ForwardCom project is being developed in an open process with contributions and ideas from various people based on the philosophy that the problems mentioned above can best be avoided through openness and collaboration.

Chapter 16

Revision history

Version 1.07, 2017-11-03.

- The first version of binary tools is published, including a high-level assembler and disassembler. A manual is included in chapter 13. The assembly language is modified.
- A feature for re-linking of executable files replaces the previous idea of load-time library dispatching.
- Triple size formats with template E added. Some formats renumbered.
- Some support for half precision floating point vectors.
- Function calling convention allows return using two registers.
- 4-bit constants in tiny format changed from signed to unsigned.
- Many small changes in instruction list.
- Object file format modified to indicate IP, DATAP, or THREADP addressing for sections and symbols

Version 1.06, 2017-02-14.

- Added chapter: Proposal for reducing branch misprediction delay
- Added instruction: `increment_jump_sabove`.
- Modified various conditional jump instructions. More detailed descriptions.

Version 1.05, 2017-01-22.

- Systematic description of all instructions.
- Instruction list updated.
- Added chapter: Support for multiple instruction sets.
- Added chapter: Software optimization guidelines.
- Bit manipulation instructions improved.
- Shift instructions can multiply float by power of 2.
- Integer division with different rounding modes.
- Source of option bits for `mul_add`, `add_add` and `compare` instructions modified.

Version 1.04, 2016-12-08.

- Instruction formats made more consistent. Template E2 modified.
- Masking principle changed. Fallback value option. r0 and v0 allowed as masks.
- Compare instruction has additional features.
- Conditional jumps modified
- Several other instructions modified.

Version 1.03, 2016-08-01.

- Minor changes and additions to manual.
- Three new instructions added.

Version 1.02, 2016-06-25.

- Name changed to ForwardCom.
- Moved to github.
- Various security features added.
- Support for dual stack.
- Some instruction formats modified, including more formats for jump and call instructions.
- System call, system return and trap instructions added.
- New addressing mode for arrays with bounds checking.
- Several instructions modified or added.
- Memory management and ABI standards described in more detail.
- Instruction list in comma separated file instruction_list.csv.
- Object file format defined in file elf_forwardcom.h

Version 1.01, 2016-05-10.

- The instruction set is given the name CRISC1.
- The length of a vector register is stored in the register itself. The basic code structure is modified as a consequence of this. Function calling conventions are also simplified as a consequence of this.
- All user-level instructions are defined.
- The entire text has been rewritten and updated.

Version 1.00, 2016-03-22.

This document is the result of a long discussion on Agner Fog's blog , starting on 2015-12-27, as well as input from the RISC-V mailing list and the Opencores forum.

Additional inspiration was found in various sources listed on page 5.

Version 1.00 of this manual was published at www.agner.org/optimize.

Chapter 17

Copyright notice

This document is copyrighted 2016-2017 by Agner Fog with a Creative Commons attribution-share alike license. creativecommons.org/licenses/by-sa/4.0/legalcode.