

HIVE

**32 Bit – 8 Thread – 8 Register/Stack Hybrid – Barrel Pipelined
Verilog Soft Processor Core**

INTRODUCTION

Hive is a general purpose soft processor core intended for instantiation in an FPGA when CPU functionality is desired but when an ARM or similar would be overkill. The Hive core is complex enough to be useful, with a wide data path, a relatively full set of instructions, high code density, and good ALU utilization – but with very basic control structures and minimal internal state, so it is simple enough for a human to easily grasp and program at the lowest level without any special tools. It fits in the smallest of current FPGAs with sufficient resources left over for peripherals (as well as other unrelated logic) and operates at or near the top speed of the device DSP hardware.

Hive isn't an acronym, the name is meant to suggest the swarm of activity in an insect hive: many threads sharing the same program and data space, individually beavering away on separate tasks, and cooperating together to accomplish larger goals. Because of the shared memory space, thread intercommunication is facilitated, and threads can all share a single instance of code, subroutines, and data sets which enables code compaction via global factoring.

The novel hybrid stack / register construct employed reduces the need for a plethora of registers and allows for small operand indexes in the opcode. This construct, coupled with explicit stack pointer control in the form of a pop bit for each stack index, minimizes the confusing and inefficient stack gymnastics (swap, pick, roll, copying to thwart auto-consumption, etc.) normally associated with conventional stack machines, and also minimizes the saving and restoring of register contents normally associated with conventional register machines.

Hive employs a naturally emergent form of multi-threaded scheduling which eliminates all pipeline hazards and provides the programmer with as many equal bandwidth threads – each with its own independent interrupt – as pipeline stages. Processors that employ this form of pipelining are classified as “barrel” processors.

Hive is a largely stateless design (no pipeline bubbles, no registered ALU flags that may or may not be automatically updated, no reserved data registers, no pending operations, no branch prediction, etc.) so subroutines require no overhead and interrupts consume a single branch cycle, and their calculations can be performed directly and immediately with complete disregard for what may be transpiring in other contexts.

This paper presents the design of Hive along with some general background, so if you don't find the architecture of Hive itself to your liking, you may possibly find something else of use in it. Enjoy!

“Tis the gift to be simple, 'tis the gift to be free...”

From “Simple Gifts” by Elder Joseph

HIVE FEATURE LIST

- ✖ A simple, compact, relatively stateless, high speed, barrel pipelined, multi-threaded design based on novel RASH (**R**egister **A**nd **S**tack **H**ybrid) technology.
- ✖ 2 operand machine with operand select and stack control fields in the opcode.
- ✖ 32 bit data path with extended width arithmetic results.
- ✖ 16 bit addresses, instructions, and memory data accesses.
- ✖ 8 equal bandwidth threads.
- ✖ 8 independent, isolated, general purpose LIFO data stacks per thread with parameterized depth and fault protections.
- ✖ 8 fully independent interrupts with no hierarchical limitations (one per thread).
- ✖ 8 stage pipeline with no stalls or hazards.
- ✖ All instructions execute in a single pipeline cycle, including 32 x 32 = 64 bit signed / unsigned multiply (lower or extended).
- ✖ Common data & instruction memory space (Von Neumann architecture) enables dynamic code / data partitioning, combined code and data constructs, code copy & move, etc.
- ✖ All threads share the entire common data / code space, which facilitates global code factoring and thread intercommunication.
- ✖ Internal register set with highly configurable base register module that may be easily modified / expanded to enhance I/O, debug, etc.
- ✖ 32 bit GPIO with atomic full-width read.
- ✖ Double buffered UART with BAUD generator and several parameterized options.
- ✖ Written in 100% highly portable verilog (no vendor specific or proprietary language constructs).
- ✖ Partitioned into easy to understand and verify verilog modules.
- ✖ May be programmed via a verilog initial text file, no complex tool chain is necessary.
- ✖ Achieves aggregate throughput of ~200 MIPS in a bargain basement Altera EP3C5E144C8 FPGA (Cyclone 3, speed grade 8 – the target device for initial development) while consuming ~2600 logic elements, or ~50% of an EP3C5E144C8.
- ✖ Free to use, modify, distribute, etc. (but only for the greater good, please see the copyright).

CONTENTS

Introduction
Hive Feature List
Contents
Motivation
Register Machines vs. Stack Machines
Background: LIFOs
Register / Stack Hybrid
Operands
Stacks
Data Width
Access Width
Address Width
Arithmetic Results Width
Signed vs. Unsigned Arithmetic
Background: FPGA Resources
ALU Design
Pipelined Core
Instructions / Opcodes
Internal Register Set
Verification
Speed
Programming
Bzzz!
Etc.
Document Change Control
Copyright

MOTIVATION

As a (mostly) digital designer who works primarily in FPGAs, I'm always on the lookout for simple processor cores because projects often underutilize the hardware due to low data rates (e.g. a UART, or a sampled audio stream). If latency isn't a big issue, then why not multiplex the hardware with a processor construct? But the core needs to be really simple, not consume too much in the way of logic (LUTs, block RAMs, multipliers), have compact op codes (internal block RAM isn't limitless nor inexpensive), keep the ALU sufficiently busy, and be easy to program at the machine code level without the need for a compiler or even an assembler.

FPGA vendors have off-the-shelf designs that are quite polished and bug-free, but they, and therefore the larger design and the designer, are generally legally chained to that vendor's silicon and tool set. There are lots of free cores available on the web, but one may end up getting exactly what one paid for.

The Hive core is my offering for this problem area. The essentially free and naturally emergent multi-threading / rigid scheduling mechanism in Hive is not unique; I believe it was implemented as far back as 1964 on the CDC 6000 series peripheral barrel processors. Hive shift distances are treated as signed, which works out rather nicely, but the ancient PDP 10 does this as well. The notion of multiple stacks isn't original, nor is the explicit control over the processor stack pointer, but I believe the register/stack hybrid as implemented and described here (indexed stacks with conservative top-entry-only access and pop bit override) is something relatively new. And the way extended arithmetic results are dealt with uniformly in Hive may possibly be somewhat novel as well. But who knows? Processors have been around long enough that most of the good ideas have been mined out and put to the test in one form or another, which makes it difficult / unlikely to bring something fundamentally new or innovative to the table.

REGISTER MACHINES VS. STACK MACHINES

Most modern processors are register based, and so have some form of register set tightly bound to the ALU – a tiny fast triple port memory in a sense. This conveniently continues the memory hierarchy of faster and smaller the closer to the core, and has the advantage of being a very mature target for compilers.

Many registers are generally available because the register space grows exponentially with register address width. But register opcode indexes can still consume significant opcode space, particularly in a 3 operand machine, and register count is a limited resource that doesn't scale with the rest of the design. Registers are often reserved for special purposes, and some may be invisible to non-supervisory code. It would seem the more registers available, particularly of the "special" variety, the more the programmer has to juggle in his/her head. And a general purpose register may only be used if the programmer is absolutely certain that any data there is moot, or if the register contents are first saved to memory and later restored, which is something else to keep track of.

Almost since my first exposure to data stacks via an HP calculator (won in a high school engineering contest) I've been fascinated with stack languages and stack machines. With no explicit operands, a data stack, a return stack, and almost no internal state, a stack machine can have incredibly compact op codes - often 5 bits will do. Due to the stacked registers, interrupts and subroutines and other forms of code factoring can be quite efficient; all that is required is that they clean up after themselves. I've studied many of these, and have coded a few of my own and had them running on an FPGA demo board. They are surprisingly easy to implement but surprisingly cumbersome to program - one has to stick loop indices, conditional test values, and branch addresses under the operands on the stack or in memory somewhere, so there are a lot of operations and much real time wasted on stack manipulation which can get very confusing very quickly. Laborious hand optimization of stack code leads to "write only" procedural programs that are difficult to decipher later, and with catastrophic stack faults all too likely. The tiny opcode widths produce a natural instruction caching mechanism, but having multiple opcodes per word is awkward when they aren't powers of 2 wide, a nuisance when one must manually change the code by hand (one usually end up inserting no-ops to pad out the space), and interrupts / subroutines must either return to a word boundary (more no-ops / wasted program space) or the hardware must somehow store and retrieve a sub index into the return word (more state).

Stack machines are, perhaps somewhat inadvertently, portrayed as a panacea for computing ills, but with little in the way of formal analysis to back up these assertions. They are something very different and on the fringe and as such don't get addressed by the mainstream, so there aren't many technical comparisons (speed, code density, etc.) to more conventional architectures – or detractors for that matter, so the stack machine noob encounters a situation rather like serving on a jury and hearing only the one side of the case. My conclusion is their biggest strength – implicit operands – is also their biggest weakness. One has to follow the intricate stack manipulations closely and with some kind of idea of what the programmer originally had in mind in order to make any sense of the code. One can't rely on, say, the loop index residing and staying put in register 4 and the like. There are of course stack machines out there that have register sets tacked on, but this tends to bloat out the opcodes and doesn't strike me as a very elegant solution.

Another thing that isn't discussed much regarding stack machines is that auto consumption of *all* input values is generally necessary. While it is obvious that ALU operations pop the input operand(s) and push the result, it isn't always emphasized that conditional branches generally consume the branch test value(s) and the branch address regardless of whether the branch is taken or not. Auto consumption is a big issue because it leads to much copying or restoring of values to be used both now and later, and it also means most instructions cannot be made individually conditional (ala the ARM, or via a skip instruction) because the stack pointer(s) will likely be different depending on whether the instruction was executed or not, something the programmer can't generally track.

So my own personal conclusion is this: a single data stack is a good fit for hand calculator data input and manipulation. But even with the inclusion of a second general purpose data stack or dedicated return stack, it's not so great a paradigm on which to base processors and programming languages.

BACKGROUND: LIFOS

Since the discussion is about stack machines, it helps to fully understand stacks themselves, which are based on the LIFO (Last In First Out) construct.

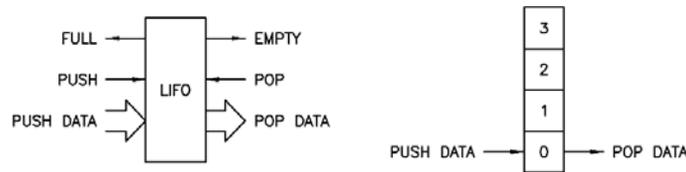


Figure 1. LIFO symbols.

The figure above shows two LIFO symbols, the one on the left is I/O centric, the one on the right more of a schematic memory view. Unlike FIFOs, which need separate read and write side pointers, LIFOs only require a single pointer, which may be implemented in such a way as to conveniently reflect the fullness of the LIFO. The push side is only concerned with whether the LIFO is full or not, the pop side only concerned if it is empty or not. Push when full is an error because it may drop the input data on the floor and may corrupt the LIFO pointer. Pop when empty is an error because it gives false read data and also may corrupt the LIFO pointer.

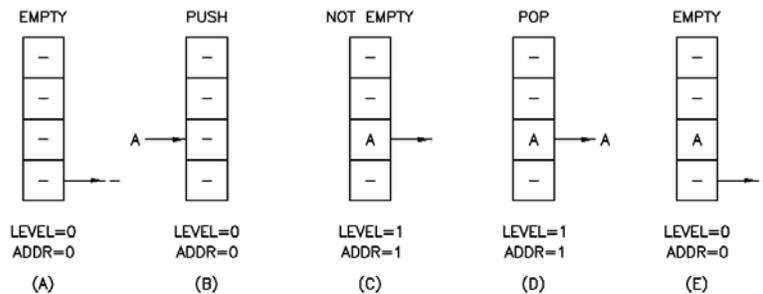


Figure 2. LIFO stack operations – push then pop from empty state.

The figure above shows LIFO operation from empty, to not empty, to empty again. Note that the first write to memory is address 1 rather than address 0, which may seem a bit counter-intuitive. This convention allows the level and pointer values to be the same.

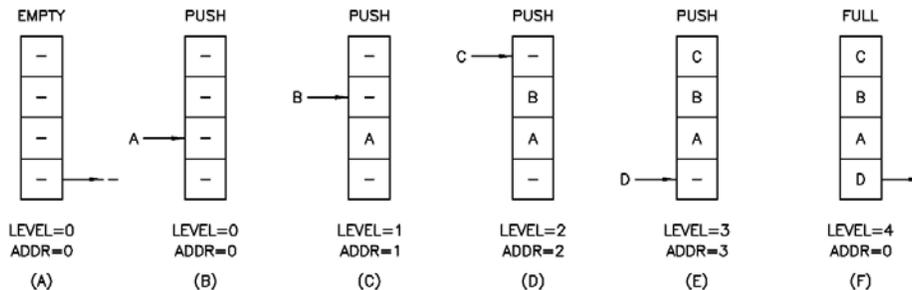


Figure 3. LIFO stack operations – push from empty state to full state.

The next figure shows LIFO operation from empty to full. Note that the last write to memory is at address 0, which may also seem a bit counter-intuitive. It helps here to think of the address as modulo (one MSB less than) the level value. For this 4-deep LIFO there are actually 5 distinct states corresponding to levels 0 through 4. Indeed, when fully utilizing the LIFO memory space there will always $2^n + 1$ levels, and it is easiest and most straightforward to handle them with an extra MSB in the level counter, and present the LSBs of this counter to the LIFO memory address input (i.e. the stack pointer).

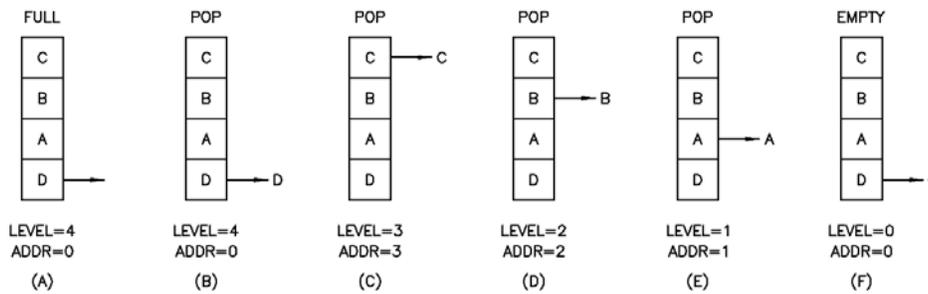


Figure 4. LIFO stack operations – pop from full state to empty state.

The next figure shows the previously filled LIFO operation from full to empty. At the end (in this case) the value D at memory location 0 is presented as output, but it is flagged as invalid by the empty indicator so the pop side knows not to use it.

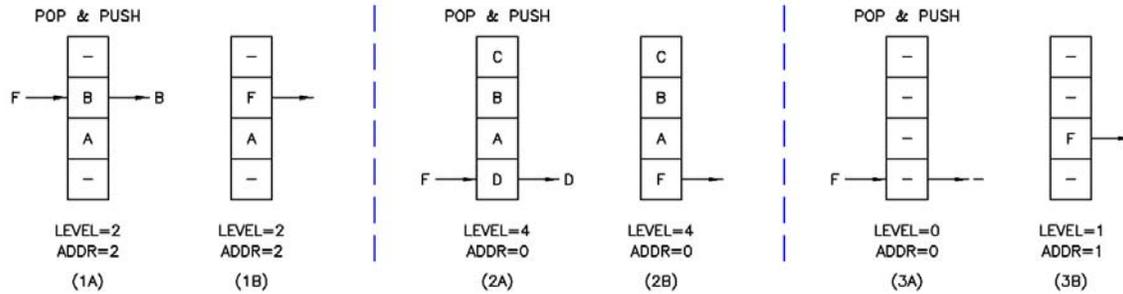


Figure 5. LIFO stack operations – three pop & push scenarios.

What happens if we pop and push at the same time? For a canonical stack machine we need to read the pop side value, pop it off the stack, and then push the result onto the stack. This is a pop & push (as opposed to a push & pop, which is nonsensical for this application). At the above left we see a pop & push in action, the value B at address location 2 is overwritten with the value F, and there is no net pointer change. In the center we see a pop & push when full, which is not an error because pop, which decrements the pointer, can be thought of as preceding push, which increments the pointer. Finally, on the right we see a pop & push when empty. This is obviously a pop error because the read data is invalid, but it is a pop error only! If the pointers are internally protected from corruption then the correct net result is a push.

Stack Protection

Is it always best to protect the stack against the corruption of the pointer or memory contents? It may seem that the answer to this is always “yes” but consider the following scenario. Say a stack is almost full and a data value is pushed to it, making it full. Then an address is pushed to the stack and the thread attempts to branch to this address. If the pointer and stack memory are overflow protected then the address was dropped on the floor and the thread instead branches to the location given by the previously pushed *data* – off into the weeds it goes with one stack that is essentially stuck and can’t accept new data (unless it is a pop & push, or unless a pop is otherwise performed first). The thread could be returned to sanity with an external clear (perhaps issued by another thread on supervisory duty) but the stuck stack means the thread itself has limited ability to fix its own problems. Would it be better to *not* protect against push errors, and just let them corrupt the first stack entries so the thread could continue? Granted this kicks the problem down the road, but perhaps the thread wasn’t going to use the earlier entries on the stack anyway and was about to issue a routine thread clear? Or perhaps it was about to check itself for stack errors and if it found one would have cleared itself? At least it isn’t immediately derailed and off corrupting the contents of main memory.

In contrast, I believe pop (underflow) protection is generally good because it prevents the stack from rolling under and thereby offering up completely unrelated, non-local data and addresses to the thread.

Contemplating how to deal with these “what if” conditions that shouldn’t happen (but likely will, at least during SW development) can drive you a little crazy. In any case, pop and push protections are individually configurable build options in Hive so you can set them however you like. And regardless of the protection settings, stack errors are always reported to the local register set.

REGISTER / STACK HYBRID

Many register based machines have a return stack, and many stack machines have a one or more registers stuck somewhere, but beyond this is there some kind of middle ground between stack based and register based machines? If a register based machine were designed with a LIFO stack under each register, then perhaps the programmer could accomplish the same goals with fewer indexed register locations, meaning the register index could be made narrower giving a more compact and efficient opcode. Multiple stacks would be more convenient than a single stack for complex algorithms, and would help keep inefficient and confusing stack thrash to a minimum. Unlike register count, LIFO depth can easily scale as required by other aspects of the design. Could the stacks indeed be indexed as register operands? If so, how might multiple stacks be implemented and how would the stack push/pop mechanism behave?

I recently encountered the J1 stack based processor (<http://www.excamera.com/sphinx/fpga-j1.html>) which is quite intriguing in that it has a two bit wide signed stack increment field in the opcode. This idea inspired me to investigate explicit rather than implicit stack control. I decided that an array of simple stacks, where only the top stack values are presented to the ALU (as opposed to the top *and* second values as in a conventional stack machine) would suffice. The stacks could then be indexed normally as register locations, with the usual one or two sources and one destination. I then came up with a simple, inherently conservative stack mechanism: whenever anything is read from a stack, the value and stack pointer remain unchanged. Whenever anything is written to a stack the value already there is pushed in to make room for the new value. Each stack index is provided with an associated pop bit that influences this default conservative behavior:

pop bit	read / write	Stack	Behavior
0	read	no change	Register type read.
1	read	pop	Stack type read.
0	write	push	Stack type write.
1	write	pop & push	Register type write.

Figure 6. Hybrid register / stack behavior.

This arrangement accommodates the full range of stack / register behaviors. For example, say the operand source of an ALU single operand operation is stack index B and the result destination is stack index A:

Case	B pop	A pop	B stack	A stack	Behavior
0	0	0	no change	push	Register type read, stack type write.
1	0	1	no change	pop & push	Register type read & write.
2	1	0	pop	push	Stack type read & write.
3	1	1	pop	pop & push	Stack type read, register type write.

Figure 7. One and two operand hybrid register / stack behavior.

Cases 1 and 2 respectively give the normal pure register and pure stack behaviors, while cases 0 and 3 give useful variations. What about the two input operand case? Say the primary input operand is stack index A, the secondary input operand is stack index B, with the result going to stack index A (e.g. a two operand opcode architecture). It turns out that the same table above works for both scenarios. How do we handle the case where both of the sources and the destination point to the same stack? The solution is to simply OR the two pop bits together. Remember that there is no access to the value below the top LIFO entry as in most stack machines, so when index A = index B for a two operand instruction such as multiply, the result will be A^2 pushed to A. And in this case, if both of the A and B pop bits are set this won't cause a double pop because the pop bits are simply ORed, causing a single pop of A (a pop & push, actually).

Now that we have simpler stacks and more control over them, the conditional execution of single operations is suddenly a viable option. Conventional stack machines generally don't have conditional single operations because operands are always consumed – the programmer wouldn't be able to tell what state the stack is in after a conditional two operand operation, leading to stack faults. With no auto-consumption of the input, and by setting the pop bit of the register being conditionally written to, we can ensure the stack pointers don't change during a single conditional operation.

OPERANDS

How many operands should be in the opcode? I picked 2 to keep the opcode small, so Hive is a 2 operand machine. Here are the rules:

- For single input ALU operations the source is B and the result destination is A. For example: $\text{not } B \Rightarrow A$.
- For two input ALU operations the primary source is A, the secondary source is B, and the result destination is A. For example: $A - B \Rightarrow A$.
- For single input conditional branch statements A is tested against zero, and the absolute address is B, or the address increment is B or is supplied as immediate data. For example: $(A > 0) ? B \Rightarrow \text{PC}$.
- For two input conditional branch statements, A is tested against B, and the address increment is supplied as immediate data. For example: $(A \neq B) ? \text{PC} + I \Rightarrow \text{PC}$.
- For memory reads the base address is B, the read value is written to A, and there is an immediate 4 bit positive address offset. For example: $\text{mem}(B + I) \Rightarrow A$.
- For memory writes the base address is B, the write value is read from A, and there is an immediate 4 bit positive address offset. For example: $A \Rightarrow \text{mem}(B + I)$.
- For subroutines the subroutine absolute address is B and the return address (the PC) is pushed to A.
- When an interrupt is taken the return address (the PC) is automatically pushed to stack S0 (this is the only "special" stack, and this is the only way in which it is "special").

So A is the primary data source and destination for two operand operations, is the primary data tested, receives subroutine return addresses, and is the only thing that can be written to. B is the primary data source for one operand operations, the secondary data source for two operand operations, is the secondary data that A is tested against, and always provides the address or address offset.

STACKS

How many stacks are needed? I picked 8. This gives 4 bits for each operand (one pop bit, three stack index bits) for a total of 8 bits of opcode consumed. How deep should the stacks be? I've read 32 entries are generally deep enough for single stack machines to not require auto spill-to-memory mechanisms and the like. Since we have 8 stacks, and since coding for this core is likely to be done by hand, we could doubtless get by with less depth. In any case the use of FPGA block RAM for the stacks sets a fairly generous practical lower limit (32 entries per stack per thread in our target device).

DATA WIDTH

Hive data is 32 bits wide. Byte data is too narrow for many applications, and 16 bit data doesn't have sufficient resolution to directly perform the internal computations required for audio DSP. 64 bit data is overkill for most applications that would be running on a small FPGA processor. Non-power-of-2 widths can be excluded for efficiency reasons, which leaves us with 32 bits. Data width directly dictates the top speed vs. pipelining depth because wider data requires more deeply cascaded combinatorial logic to perform adds, multiplies, etc.

ACCESS WIDTH

Hive access width, and by that I mean main (data and program) memory read / write – which includes instruction fetch and in-line literal access – is 16 bits. One generally desires everything to be the same width, but an intentionally compacted opcode can easily wind up narrower than the ALU data path. And much can be done with 16 bit data which can directly lead to code compaction.

ADDRESS WIDTH

Hive addresses are limited to 16 bits, giving an address space of 65k 16 bit words.

One can agonize over the small address space and the narrow access, but the target functionality here is a small processor in an FPGA, so large address space isn't often necessary, and beyond a certain depth FPGA-based memories are slower. At any rate, I believe only trivial edits would be required to provide Hive with a full 32 bit address, and I've been very careful to exclude opcode constructs that limit the address width to 16 bits.

ARITHMETIC RESULTS WIDTH

Some ALU arithmetic operations invariably produce wider results than the input operands. Traditional processors stick the extended results of add and subtract (carry, overflow, sign, etc.) in dedicated bit flag registers, and then have rules and special instructions that govern the updating, clearing, saving, and restoring of them. The results of full width multiplies are often sent to special concatenated register pairs. These practices may be efficient, but they introduce complexity and internal state.

A simple and uniform method of handling wide arithmetic results is to treat them as double width regardless of operation (add, subtract, multiply) and select either the lower (i.e. normal) half of the result or the upper (extended) half of the result via instructions. The obvious downside here is that obtaining the full width result takes two cycles even when the operation is actually performed in one. For the full result it may seem wasteful to perform the same internal calculation both times, but one probably shouldn't think of this as major effort for the ALU or as a huge opportunity lost. All processors have to perform a full width subtraction in order to generate the arithmetic comparison flags between two numbers. By examining the extended result of add / subtract first one can know beforehand if the result will overflow and perhaps not perform it (e.g. restoring division), and often only the lower or extended arithmetic result is required.

Interestingly, the extended result of signed and unsigned subtraction always forms a convenient all ones or all zeros flag (easily negated with a NOT instruction). The extended result of unsigned addition is a bit more complex. Here are some 4 bit corner case examples to get a flavor of how this works:

+ unsigned	15 + 15 = 30 = 0001,1110	max
	0 + 0 = 0 = 0000,0000	min
+ signed	7 + 7 = 14 = 0000,1110	max
	-8 + -8 = -16 = 1111,0000	min
- unsigned	15 - 0 = 15 = 0000,1111	max
	0 - 15 = -15 = 1111,0001	min
- signed	7 - -8 = 15 = 0000,1111	max
	-8 - 7 = -15 = 1111,0001	min
* unsigned	15 x 15 = 225 = 1110,0001	max
	0 x 0 = 0 = 0000,0000	min
* signed	-8 x -8 = 64 = 0100,0000	max
	7 x -8 = -54 = 1100,1000	min

Figure 8. 4 bit input / 8 bit corner results.

SIGNED VS. UNSIGNED ARITHMETIC

Although addresses are generally thought of as unsigned, unsigned subtraction will produce negative numbers whether one likes it or not. The programmer obviously needs the resources to handle both, so the impact of signed vs. unsigned arithmetic is largely one of default behavior and instruction naming conventions. Signed multiply is more basic due to sign/zero extension needs (hence Altera's FPGA multiply hardware primitives being signed). I feel that signed immediate, half-width memory read, and half-width in-line literal data is more useful than unsigned because it influences the MSBs above. Given the way that Hive deals with extended results, lower arithmetic operations are sign neutral (i.e. give the same results regardless of signed / unsigned operation) so only the right shift operations and the arithmetic operations that produce extended results as output need to be differentiated with respect to sign.

BACKGROUND: FPGA RESOURCES

The available physical resources and their detailed behavior, limitations, and timing characteristics in the target FPGA will strongly influence the top speed, size, and other important bulk metrics of any soft processor. One may as well exploit these resources up front rather than be stymied by them later.

Block RAM (BRAM)

The primary FPGA component the soft processor designer needs to understand is block RAM.

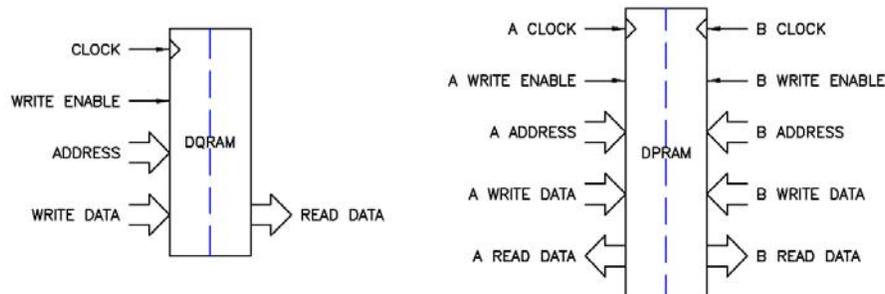


Figure 9. Block RAM: simple (DQ) on left, true dual port (DP) on right.

The figure above shows two common forms of block RAM: a simple dual port (DQ) on the left and a true dual port (DP) on the right. Because it uses a single address, the DQ variant is a good fit for the LIFO stacks. The DP variant is useful for main memory as it gives two independent accesses which enables a data read / write along with instruction fetch per cycle (thus sidestepping the “Von Neumann bottleneck”). Main memory access is a huge driver in any processor design, and often the limitation encountered is insufficient address ports, not so much data ports.

Block RAM resources have configurable variable widths, from some maximum down to a single bit. For widths of 8 and above an additional bit per byte (8+1, 16+2, 32+4) is provided for out-of-band signaling, individual byte enables, CRC, error correction, and other common uses. I believe it is a mistake to employ these extra bits in order to increase the data width of the ALU or instructions, as this precludes the efficient use of conventional 2^n width memory to store internal data / control information.

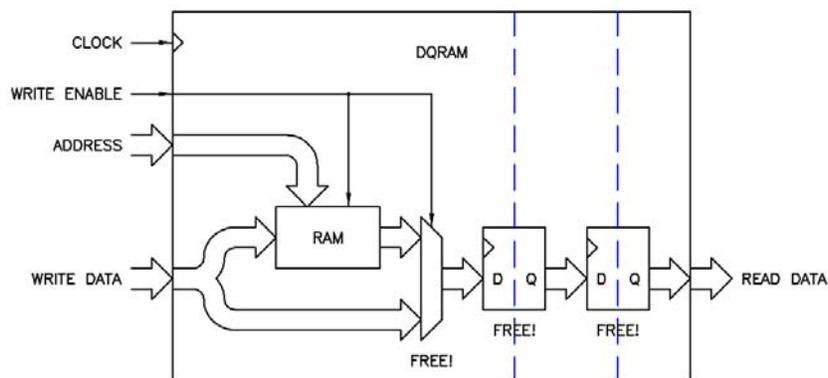


Figure 10. Block RAM internal resources.

What resources are available within block RAMs? The figure above shows a schematic view of the “inside” of a typical DQ RAM, though it applies to DP RAM as well. Even though FPGA block RAMs are *always* fully synchronous, it is sometimes helpful to think of the base RAM entity inside of the block as asynchronous. This RAM entity is supplemented with “read through” logic in the form of a multiplexer, which enables two types of configurable (at build time) read-during-write behavior. Without the multiplexer, a read-during-write delivers the old memory data to the read data port. With the multiplexer, a read-during-write conveys the data being written to the read data port. The register following this optional multiplexer is *always* present. Following this is yet another register; it is optional and generally part of the block RAM circuitry because it can dramatically speed up read clocking at the expense of one additional clock of latency.

In terms of read-during-write behavior, Hive needs write-through mode for the LIFO stacks to function correctly. This mode is unimportant for the main memory however because we will never be simultaneously reading from and writing to the data port, and the fetch port is read-only. In terms of speed, the write side is often capable of near fabric flip flop performance, while the read side is rather slow if the additional output register isn't used. So a certain amount of combinatorial logic can be placed directly in front of the write side, and if our architecture can tolerate the latency of the additional read side output register we should certainly use it because it speeds things up and is essentially free.

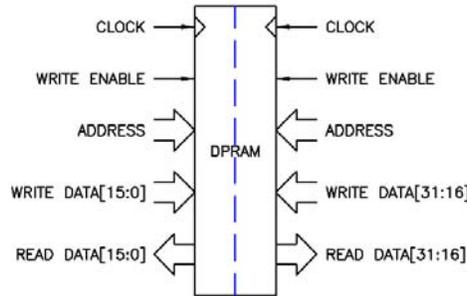


Figure 11. True dual port block RAM utilized as DQ RAM.

There is a way to convert DP RAM to DQ RAM, and this is shown in the figure above. Feeding the same clock, address, and write enable to both sides, along with splits / concatenations of the read / write data, accomplishes this simple transformation. In fact the tool will do this automatically when necessary. For our target Cyclone 3 device, DP data ports are limited to a maximum of 16 (+2) bits wide, and DQ data ports to a maximum of 32 (+4) bits wide – and the 1:2 ratio of these width limits makes sense given the above transformation. Since our LIFO stacks can employ DQ RAM (due to the single pointer) we can make them 32 bits wide using a single device.

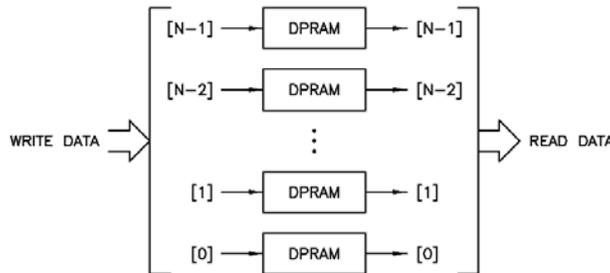


Figure 12. Block RAMs combined via bit-slice.

We may need our main memory to be considerably larger than a single 9k bit block RAM found in our target device. The tool will automatically combine multiple block RAM devices together, and often with no speed decrease – how does it accomplish this? The trick to making the largest and fastest block RAM amalgam is to configure the block RAMs to be one bit wide and maximum depth, 8k in this case, and then simply split / concatenate the write / read data by bit slicing the blocks together. Going above this size requires write steering and output multiplexing, which will also be inserted automatically by the tool when needed, but this extra logic tends to slow things down, particularly on the read side (though pipelining this logic could certainly get it back up to speed).

DSP Hardware

Since even quite low-end FPGAs these days have fairly fast hardware multipliers in some form of a DSP block, we should undertake any new designs with the knowledge and trust that they will be there. There is little point in leaving multiply operations out of our instruction set, and no point in trying to outsmart the FPGA manufacturers by constructing what would inevitably be slower and larger multipliers out of shifters, adders, etc. – both of which would needlessly strand this valuable resource. So it behooves us to understand the dedicated multiply hardware.

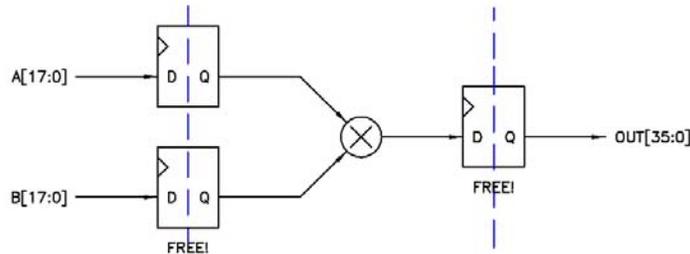


Figure 13. Signed multiplier hardware typically found in an FPGA.

Basic hardware multiplier width is 18 bits, which follows the convention of block RAM widths ($2^n + 1$ extra bit per byte). Being a full multiplier, the result is obviously double this, or 36 bits wide. As with add hardware, leaving some of the MSBs or LSBs unused will allow the remaining utilized multiply hardware to run faster due to fewer carry propagations, etc.

Altera multiplier blocks are signed by default, which makes sense because this convention simplifies sign extension of the inputs. To make a signed multiplier perform unsigned math all that is necessary is to construct it one MSB wider at the inputs and force those MSBs to zero (zero extension). Conveniently, this same construct can be used to do signed multiplication simply by driving these MSBs with the signs of the inputs (sign extension). Though of course this requires an extra bit and therefore negatively impacts top speed slightly. The extra output MSBs generated with this scheme are unused (left unconnected) which may generate tool warnings.

The multiplier hardware can be used in a purely combinatorial sense, but registering will speed it up considerably so manufacturers provide “free” internal registers at the inputs and outputs that are not part of the general FPGA fabric. As in the case of block RAM output registers, if our architecture can tolerate the latency of the additional multiplier I/O registering we would be crazy not to use it. But this leads one almost inexorably to ALU pipelining.

Digital Clock Managers (DCMs)

Virtually all FPGAs have some kind of DCM in the form of one or more PLLs (**P**hase **L**ocked **L**oops), and/or DLLs (**D**elay **L**ocked **L**oops) which may be used for a variety of purposes. A DCM can move the clock edge around to change external setup / hold / data out timing, trade internal cycle time margins for tighter external I/O timing, condition the input clock duty cycle, multiply and divide the input clock, generate multiple clocks with phase offsets, etc. Probably the main use for a DCM in a processor core is to manipulate the input clock frequency (multiply / divide) so that the clock feeding the core is at or a bit below the top theoretical speed of the core in order to get the best performance from it.

Note that there is some lower frequency limit below which a DCM will not be able to lock to or otherwise process the input clock, and this figure is given in the AC specifications datasheet for the FPGA. Also note that running the core at high frequencies will increase dynamic power consumption, and may make other logic which is not in the core but supplied by the core clock more difficult to construct due to the tighter timing constraints. It is entirely possible to have multiple clock domains inside the FPGA, but then one must take special care to condition data (particularly vectors) that cross domain boundaries.

ALU DESIGN

Building an ALU for all but the most trivial of processors is more involved than “compute all results and pick the one you want” (though in the future we may see sufficient unification of DSP blocks across devices and manufacturers and new HDL constructs that allow for more naïve ALU instantiations). Arithmetic and logical calculations aside, the wide output multiplexer itself can be a speed bottleneck. The design of the ALU drives much of the rest of the processor design, particularly one that is pipelined, so it’s not surprising if it takes a while to fully develop and implement it.

Multiplication

Let’s start with the elephant in the room – the multiply unit. If we want to do audio DSP we need $16 \times 16 = 32$ bits signed as a fairly unsuitable absolute bare minimum. We could probably get by with $16 \times 32 = 48$ bits signed, with 16 bit samples, 32 bit filter coefficients, and a 48 bit result. For the sake of symmetry and simplicity, let’s set the goal as full $32 \times 32 = 64$ bits signed and unsigned. The use of a signed base entity requires $33 \times 33 = 66$ to accommodate unsigned, which conveniently is slightly less than twice the width of a single 18×18 FPGA hardware multiplier.

Just as multiplication is performed using pencil and paper, addition and concatenation enable the utilization of several hardware multipliers in parallel, thus increasing the input and output widths. Xilinx and Altera both have nice application notes describing how to do this. Consider the following base 10 example:

$$\begin{array}{r}
 98 \\
 * 67 \\
 \hline
 56 \\
 + 630 \\
 + 480 \\
 + 5400 \\
 \hline
 6566
 \end{array}
 \Rightarrow
 \begin{array}{r}
 56 \\
 + 54 \\
 \hline
 5456
 \end{array}
 \Rightarrow
 \begin{array}{r}
 63 \\
 + 48 \\
 \hline
 111
 \end{array}
 \Rightarrow
 \begin{array}{r}
 111 \\
 + 5456 \\
 \hline
 6566
 \end{array}$$

Figure 14. Multiplication example.

On the left 98 and 67 are multiplied together in the usual manner, 7×8 , 7×90 , 60×8 , and 60×90 . All of the results of multiplication are added together to get the final answer, which requires three additions – or does it? Looking closely, 5400 and 56 can be simply concatenated, which eliminates one addition. 630 and 480 will always have zero as their least significant digits, so this addition is simplified to adding 63 and 48 giving 111. The result 1110 will also always have a zero as the least significant digit, so adding it to 5456 simplifies to adding 545 and 111 and concatenating the 6 to the least significant digit location. So 4 half width multiplications must be performed, but the three additions have been reduced to two, narrowed, simplified, and therefore likely sped up.

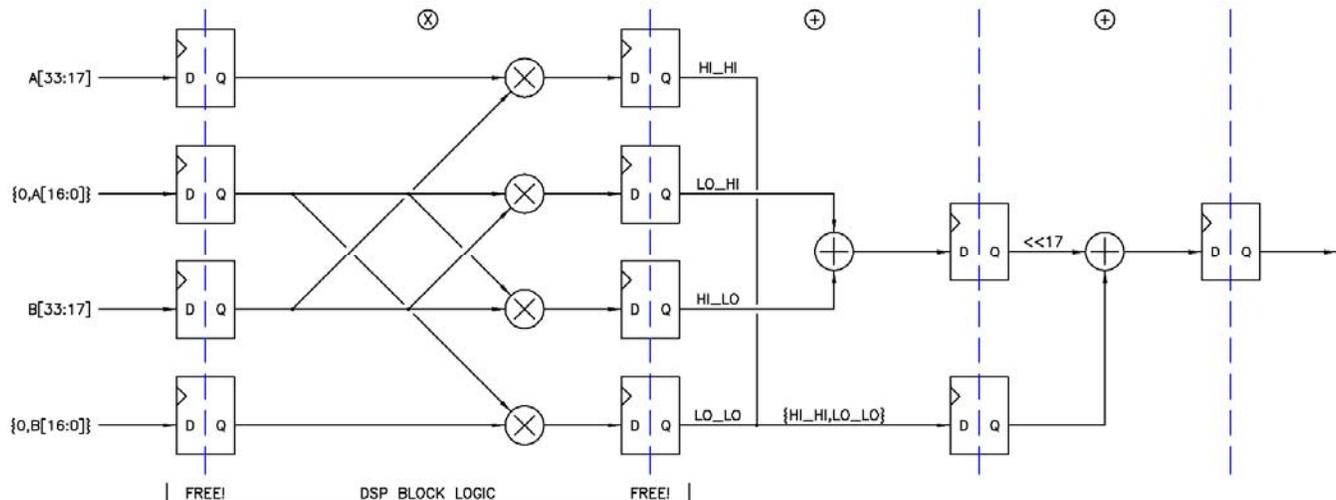


Figure 15. Three stage $33 \times 33 = 66$ bit signed pipelined multiplication.

The figure above shows these same methods implemented in binary 2s complement logic. The inputs are split in half, with the lower parts zero extended to make them unsigned (interpreting their MSBs as signs would give incorrect results). In the first stage the cross multiplications are performed, in the second stage the outer concatenation and inner add are performed, and in the third stage the final add / concatenation is carried out (the 17 LSBs of the add are automatically implemented by the compiler as a concatenation).

In terms of speed, the 18 bit multiplies in the first stage will likely be the slowest logic in the entire design, though the 47 bit add in the third stage may be close or possibly slightly worse. In the target EP3C5E144C8 device the multiply is restricted to 200 MHz, which means we should endeavor to make all of the other logic at least somewhat faster in order to have a chance of hitting 200 MIPS with the final design. The dedicated I/O registering in the multiplier hardware should certainly be used, with interstage registering to isolate the addition hardware, giving three stages and four clocks of latency.

Shifting

One thing that really nagged me about my earlier designs was that their rudimentary ALUs didn't exploit the overlapping properties of shift and multiply. It takes a fair amount of FPGA fabric logic to shift a number to the right and left some arbitrary distance and the result isn't super speedy. Having a multiplier just sitting there doing nothing useful during the shift is a missed opportunity.

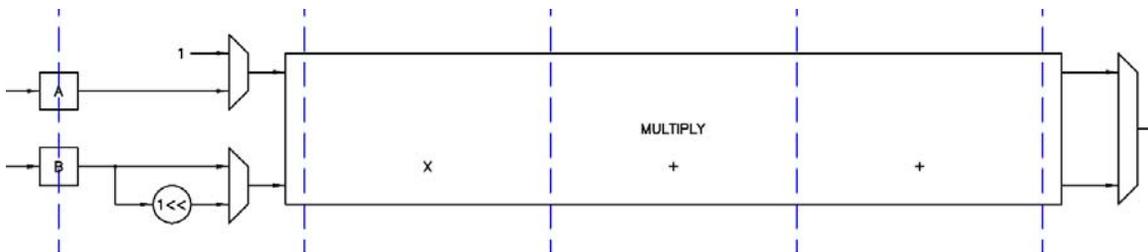


Figure 16. The Multiply and Shift unit.

When a number is multiplied by a power of 2, say 2^5 , it is shifted to the left 5 bit positions. So if a full multiplier is already present, the positioning of a simple one-hot shifter at the front ($1 \ll n$) can eliminate the left shift hardware. Can a right shift be accomplished with the same hardware? Yes, the trick is to consider the shift distance input as signed, with positive inputs causing shifts to the left and negative inputs shifts to the right. The shift distance MSB (the sign bit) is stripped off and used to select the upper (or extended) multiplication result when set (negative), and the lower result when zero (non-negative). The remaining shift distance LSBs are treated as unsigned and simply routed to the ($1 \ll n$) unit at the input as before. Here is an 8 bit example that may help clarify things:

Shift {s,n}	MSB (s)	LSBs (n)	LSBs (n)	B input ($1 \ll n$)	A input	X output
+7	0	111	7	10000000	10110111	01011011, 10000000
+6	0	110	6	01000000	10110111	00101101, 11000000
+5	0	101	5	00100000	10110111	00010110, 11100000
+4	0	100	4	00010000	10110111	00001011, 01110000
+3	0	011	3	00001000	10110111	00000101, 10111000
+2	0	010	2	00000100	10110111	00000010, 11011100
+1	0	001	1	00000010	10110111	00000001, 01101110
0	0	000	0	00000001	10110111	00000000, 10110111
-1	1	111	7	10000000	10110111	01011011, 10000000
-2	1	110	6	01000000	10110111	00101101, 11000000
-3	1	101	5	00100000	10110111	00010110, 11100000
-4	1	100	4	00010000	10110111	00001011, 01110000
-5	1	011	3	00001000	10110111	00000101, 10111000
-6	1	010	2	00000100	10110111	00000010, 11011100
-7	1	001	1	00000010	10110111	00000001, 01101110
-8	1	000	0	00000001	10110111	00000000, 10110111

Figure 17. 8 bit example of left and unsigned right shifting using a full multiplier.

Though we are thinking of the shift distance input as signed, the shifted one must be presented to the multiplier as unsigned for the 100...000 case to work correctly. Then presenting the input data to be shifted as unsigned or signed will conveniently produce unsigned ("logical" or zero extended) and signed ("arithmetic" or sign extended) right shifts. (Note that independent control over the input signedness is required for this to work, global signedness is not sufficient, which restricts the construction of signed shift from a series of more basic instructions.) So we have left shift covered, which is sign neutral, as well as unsigned and signed right shift.

Other Uses

Can more be done with this construct? A multiplexer on port A with a fixed input value of one can be used for a couple of things. The first is copying the B input shifted one result to the output of the multiplier, which is useful for generating powers of 2, bit setting & masking, etc. The second is even simpler – multiplication by one replicates the B input to the output of the multiplier, which provides us with a free and convenient "copy B" route through the ALU (though this copy feature is currently unused in Hive).

Note that signed and unsigned left shift are identical (zero padding from the right). With a bit of logic governing the input multiplexers, one of these redundant modes may be replaced with the power of 2 described above. I chose to replace immediate unsigned shift left, non-negative input shift value, with power of 2, which makes it something of an odd man out in terms of operations but hopefully not too confusing. Signed shift left works as expected. These are summarized below:

Shift Value	Instruction	Operation	Example
-	Shift left, signed	Shift right, signed	B=-3, A=10110111, Out=11110110
+,0	Shift left, signed	Shift left, signed	B=+3, A=10110111, Out=10111000
-	Shift left, unsigned	Shift right, unsigned	B=-3, A=10110111, Out=00010110
+,0	Shift left, unsigned	Power of 2	B=+3, A=xxxxxxxx, Out=00001000

Figure 18. Shifting and power of 2 functions as implemented.

Addition and Subtraction

Next we need to consider addition and subtraction. Signed and unsigned can be handled with the same method employed in the multiplier, i.e. by making the inputs one MSB wider and sign or zero extending them depending on whether that input value is to be considered signed or not. As with multiplication, overflow / carry out is extended into the double width data space and selected with instructions. Note that the lower word result is sign neutral, so only the extended result will vary based on input signed / unsigned status. The add / subtract unit is also used to compare (A<B) and (A<0) for conditional branching.

Logical Functions

For logical functions, the usual suspects are implemented:

Operation	Description	Examples
AND	A & B	A=1100, B=0101 : A=0100
OR	A B	A=1100, B=0101 : A=1101
XOR	A ^ B	A=1100, B=0101 : A=1001
NOT	~B	A=xxxx, B=0101 : A=1010
BRA	&B	A=xxxx, B=0101 : A=0000 A=xxxx, B=1111 : A=1111
BRO	B	A=xxxx, B=0101 : A=1111 A=xxxx, B=0000 : A=0000
BRX	^B	A=xxxx, B=0101 : A=0000 A=xxxx, B=0111 : A=1111

Figure 19. Logical functions as implemented (examples here limited to 4 bits).

Note that "BR" stands for "bit reduction" though it is also a mnemonically convenient reminder that B is the input to these single operand functions. The logical unit is also used to compare (A!=B) and (A!=0) for conditional branching.

Pulling It All Together

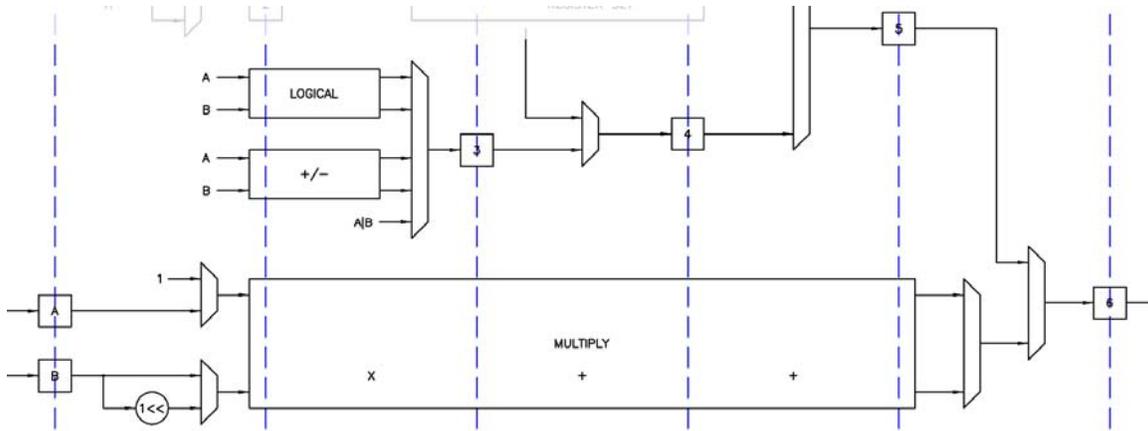


Figure 20. The Arithmetic and Logic Unit (ALU).

The figure above shows the complete ALU. The blue dashed lines represent register boundaries of a pipeline. Data enters from the left and proceeds through the pipe, with the result emerging on the right 6 clocks later. Inputs are multiplexed in, and the desired results multiplexed out. The PC (Program Counter) is multiplexed in between stages 3 and 4 for reading and subroutine / interrupt return address use. Read and literal data from main memory and the local register set is multiplexed in between stages 4 and 5. This pipeline structure provides natural intermediate value storage, so the ALU can be presented with new input data on every clock without worry that the new data will be somehow mixed in or confused with previous or later data. Pipeline interstage registering speeds things up and is an otherwise largely stranded FPGA resource, so it might as well be used (my earlier processor designs only employed a few percent of the fabric registers, and not surprisingly were relatively slow).

A somewhat thorny issue with ALU design is working out what the control inputs should be and how they should be implemented. So as not to slow things down with elaborate encoding and decoding, I decided to encode them one-hot, but with a precedence that is not actually relied upon in practice. The control signals are also pipelined, so the data and the desired operation on it may be conveniently presented together on the left. The multiply and shift unit is complex enough to have its own controls internally pipelined.

PIPELINED CORE

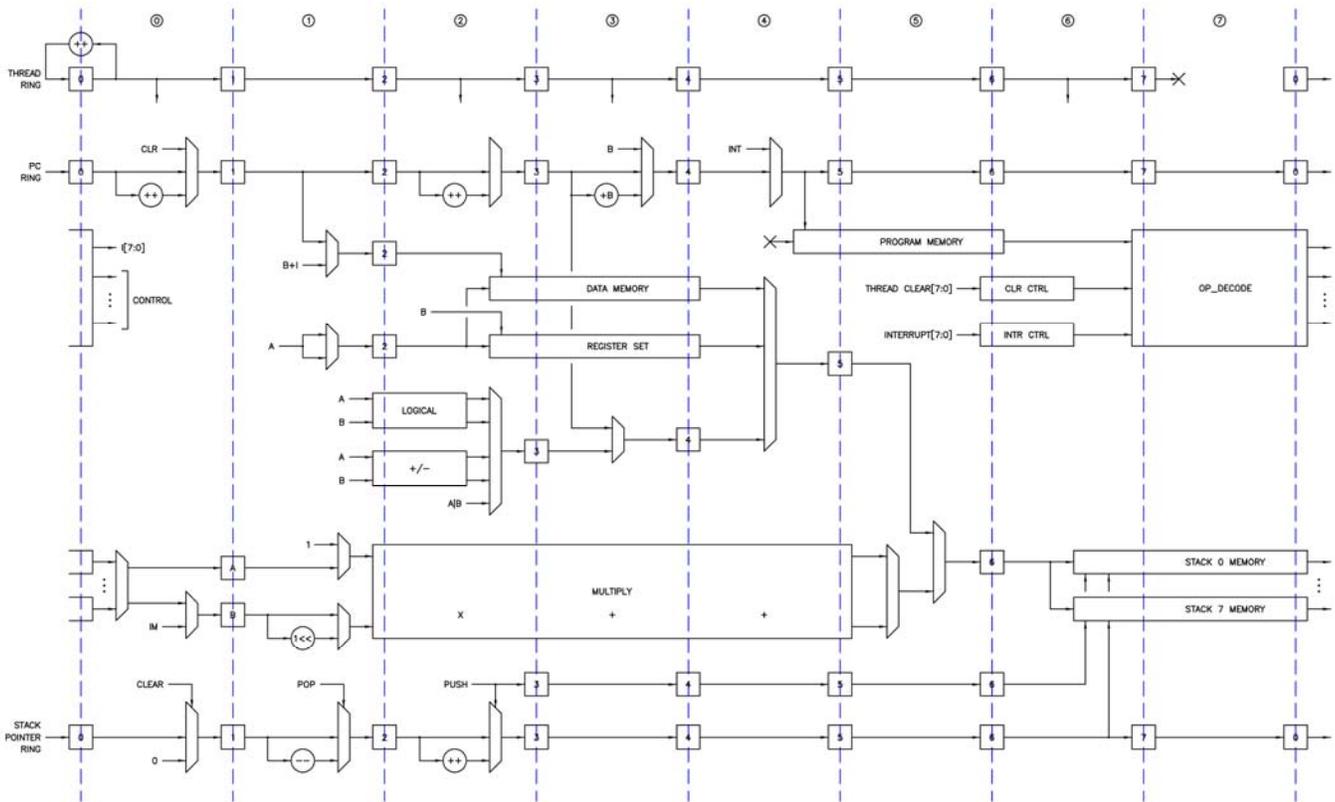


Figure 21. Hive core – view from 100 feet up.

Shown above is the full Hive core. The dotted lines and numbered boxes represent interstage registering. I'll refer to the logic *following* a line of registers with the same numbering as the registers to the left, e.g. stage 3 logic is located between the "3" and "4" register lines. Pipeline stage numbering is relative to data path operations, rather than control path operations. I chose this convention because the lion's share of the core logic – ALU & LIFOs – is contained in the data path.

It is vitally important to note that the left and right edges of the figure are connected, which converts the horizontal paths into loops, and so the core may be thought of as one large ring structure. As with the ALU, the pipeline interstage registering provides natural storage for intermediate results. With the pipelines configured as rings, values such as the PC and the LIFO pointers are not only buffered but actually *stored* in the interstage registering. Clearly this also forms a natural and simple scheduling mechanism, with packets of data and associated control information spinning around a global ring like horses on a carousel, all independent of one another, isolated by and stored within the pipe interstage registering, passed from stage to stage in a circular bucket brigade fashion. Let's call these packets "threads" – each stage of the core pipeline can receive and temporarily store, process, and pass on data and control information for a single thread, and there are 8 stages, so we have 8 threads. (Given extra buffering, one could have more threads than pipeline stages with this scheme, but not vice-versa.)

The core may then be thought of as eight processors running at 1/8 the clock speed, sharing a memory (data and address) space which facilitates intercommunication between them as well as code compaction / factoring (the sharing of common data and subroutines). The ring structure of the core forms a "barrel" type scheduler for the threads. Each thread is unique, has as much real time as the next, and gets equal access to the core resources in a strictly offset / overlapped but non-interfering manner. It is up to the programmer to keep the threads busy doing something, though of course unused threads could simply loop, perhaps waiting for an interrupt or a semaphore in memory to change (i.e. "camping on a bit").

Let's look at the individual rings in a bit more detail.

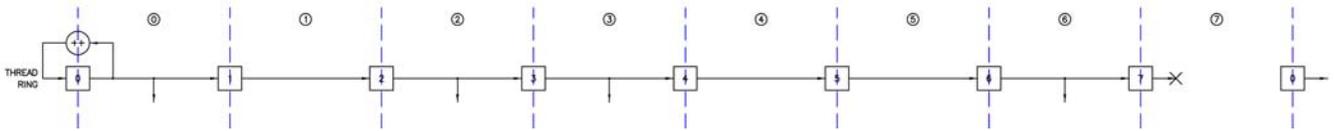


Figure 22. The Thread ID “Ring”.

Threads needs an identification number to correctly time the injection of thread clear and interrupt events into the ring, for stack error reporting, and to generate thread clear and interrupt addresses. (All threads *could* vector to the same clear and interrupt address, but that would require overhead for the thread when emerging from start up or when servicing an interrupt: read the thread ID from the local register set, use it to lookup or offset an address, jump there, etc.) A simple up-counter at the beginning of the ring generates the thread ID. A true ring structure sans counter could be used here, but that would rely on everything going well from hard reset to infinite time (never do this if you can avoid it) so we break the ring and use a counter and pipe construct instead because it is inherently self-correcting. The interstage registers emerge from asynchronous reset with the values they would normally have if previously fed by the counter, and thread ID 0 is the first to emerge from a global reset / clear, followed by 1, 2, etc. Note that this isn't a true scheduler, just a round-robin doling out of identifiers, and that any scheme which produces a continuously repeating fixed pattern where each and every ID is generated once and only once every 8 clocks would suffice.

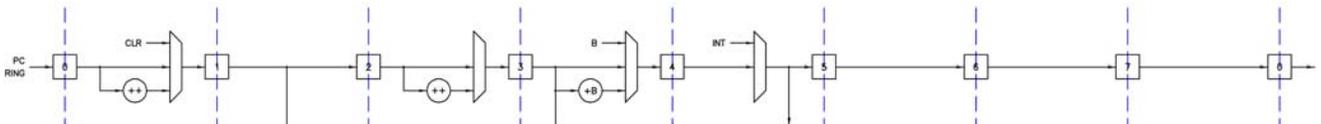


Figure 23. The Program Counter Ring.

Above is the program counter ring. At stage 0 the PC is replaced by the thread clear address if the thread is being cleared, left alone if the thread is taking an interrupt, or incremented to get the next instruction (or in-line literal). In stage 1 the PC is used as the address for the main memory data port if retrieving in-line literal data. In stage 2 the PC is incremented by 1 if retrieving an in-literal (to get the next instruction). In stage 3 the PC is sent to the data path for reading, or as a return address if taking a subroutine or interrupt. Also in stage 3 the PC is incremented by B (or an immediate value) if taking a relative jump, or replaced by B if performing an absolute jump or subroutine. In stage 4 the PC is replaced with the thread interrupt address if taking an interrupt, and then used as the address for the main memory instruction port to fetch the next instruction.

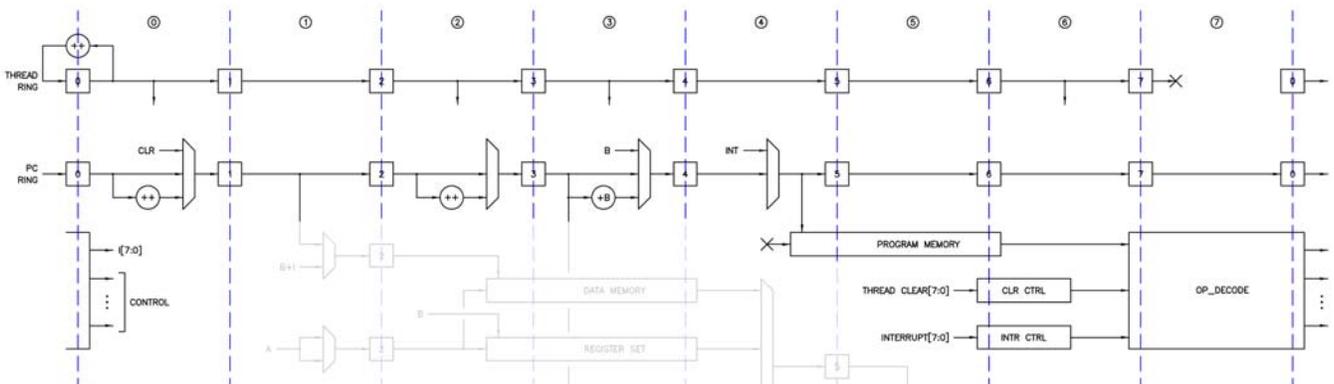


Figure 24. The Control Ring.

The thread ring and PC ring, together with the opcode decoding unit and the clear and interrupt event controllers, form the control ring. Opcode decoding takes place in several stages in order to speed it up, and as a consequence the instruction fetch must happen fairly early in the pipeline, which means conditional testing has to take place even earlier. The clear and interrupt event controllers use the thread ID to correctly inject thread clear and interrupt events into the control ring structure (and to simultaneously retire these events once injected); these events are handed off to the opcode decoder where they are prioritized and decoded. Note that each thread has its own separate clear and interrupt. The clearing or interruption of one or more threads won't disturb the other normally functioning threads. The abundance of independent interrupts means that hierarchical interrupt logic / code won't likely be necessary for most applications.

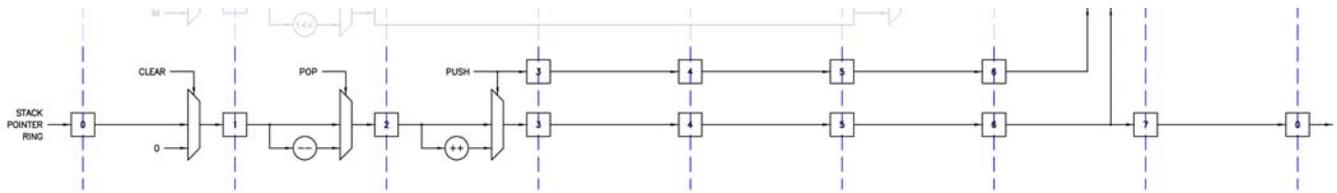


Figure 25. The Stack Pointer Ring.

Shown above is the stack pointer ring. In stage 0 the stack pointers are cleared if the thread is being cleared. In stage 1 valid pop events decrement the relevant stack pointer(s). In stage 2 valid push events increment the relevant stack pointer. Not shown in stages 1 and 2 is logic that measures fullness and prevents push when full / pop when empty from corrupting the stack pointers (if so configured at build time). These error events are reported to the local register set for debugging purposes. Separating the clear, pop, and push logic in this manner actually simplifies combined pop & push actions, as well as error tracking and reporting. Valid pushes also generate write enables for the LIFO memories, which are pipelined and applied in stage 6. In stage 6 the stack pointers are concatenated with the thread ID to form the LIFO memory write / read address and the ALU result is written to one of the stack memories. This pointer / thread ID concatenation scheme gives each thread its own private set of stacks in shared block RAM, and renders stack corruption from one thread to another impossible. Stack to stack corruption within a thread is also impossible due to the physically separate block RAMs employed for each stack.

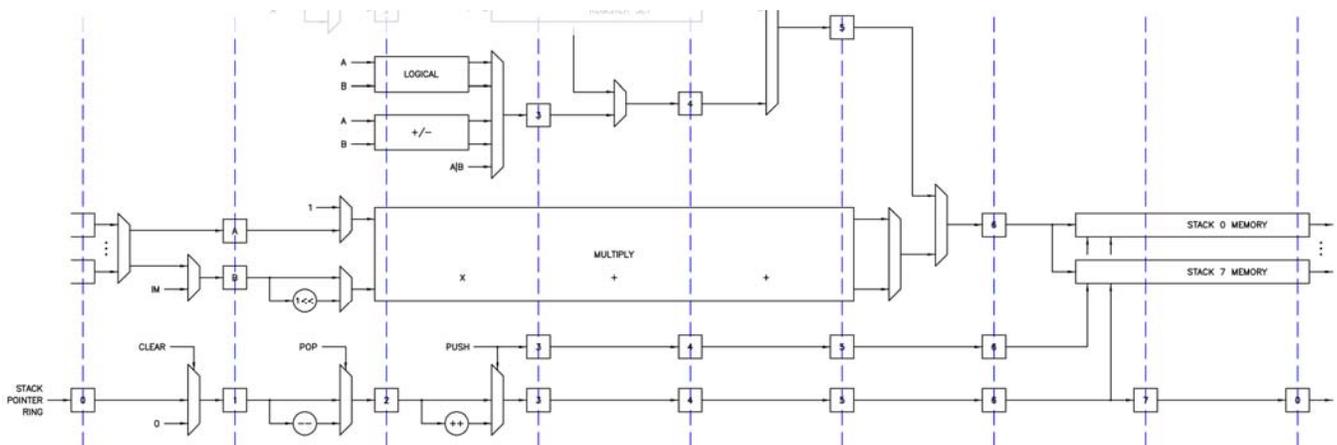


Figure 26. The Data Ring.

As shown above, the stack output multiplexer, ALU, LIFO memories, and stack pointer ring constitute the data ring.

The control ring, data ring, main memory, local register set, and register set / memory shim (not shown) make up the Hive core.

INSTRUCTIONS / OPCODES

With the basic hardware structure in place we can now decide on the instructions and opcode formats. In actuality the design process isn't this cut and dried, and the inclusion and format of certain instructions will obviously ripple back into the hardware structure.

Beyond Turing completeness, selecting a self-contained and balanced set of instructions for general purpose computing is something of a conundrum – you want to at least minimally accommodate your own code, but how do you guarantee efficient coverage for all future code you may write without resorting to a “kitchen sink” design? Studying the instruction sets of other processors is useful here, as is coding up often used simple functions like division and others which are not supported directly by the ALU. If the coding process feels particularly laborious, or the resulting code strikes you as overly awkward or unusually cryptic, then you likely have more work to do.

Determining how to best fit the instructions into the opcode space can be a challenge, particularly when the opcode is intentionally compact. From the previous discussion, we know there are at most 2 stack indexes of 3 bits each, with 1 pop bit for each index. This consumes 8 bits of opcode space, leaving 8 bits remaining. Most processors utilize the operand select field room freed up when fewer operands are required for a particular operation, and Hive does this as well.

Immediates

Instructions that contain an immediate data or address offset field can be quite effective, though they quickly gobble up opcode space so they need to be firmly in the frequent use category to earn their keep. The immediate field width and position within the opcode need not be fixed, and I decided to implement immediate 6 and 4 bit signed instructions, as well as immediate 4 bit unsigned instructions. A 6 bit wide signed immediate value when coupled to the shift instructions is useful because it allows for full 32 bit left or right shifting in a single cycle, and one or more shifts can perform many chores that would otherwise require dedicated instructions and hardware (full width MSB flag, arbitrary width sign / zero extension, isolation of contiguous bit fields, 2^n integer modulo, etc.). The 4 bit unsigned immediate is used exclusively as a memory access address offset.

Immediate instruction types and 16 bit opcode space consumption:

- Four 4 bit immediate unsigned address instructions: memory access – 16384 codes.
- Six 4 bit immediate signed address instructions: conditional (A?B) jump – 24576 codes.
- Four 6 bit immediate signed address instructions: conditional (A?0) jump – 4096 codes.
- Four 6 bit immediate signed data instructions: data, add, shift left, power / unsigned right shift – 4096 codes.

0	1	W	H	IM[3:0]				PB	B		PA	A			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

1	C		N	IM[3:0]				PB	B		PA	A			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

1	1	1	0	C	N	IM[5:0]				PA	A				
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 27. Immediate address instruction formats (top to bottom): memory access; (A?B) conditional jump; (A?0) conditional jump.

1	1	1	1	OP	IM[5:0]				PA	A					
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 28. Immediate data instruction format: literal data, add, and shift.

The bandwidth consumed by immediate / literal data is quite important; some processor designs devote (literally!) half of the opcode space to a single immediate data operation. With Hive, the immediate signed data instruction is one way to insert data from the instruction stream, another is via literal instructions that employ an in-line mechanism (the value immediately follows the literal instruction in program space). The data instruction obviously uses 16 bits and one cycle to push 6 signed bits of data. The in-line literal instructions use 32 bits (the 16 bit literal instruction followed by 16 bits of data – the data is used “literally” rather than decoded) but still just one cycle to push 16 bits of data. There are signed and unsigned literal low instructions, and a high instruction that preserves the lower 16 bit value to fill 32 data bits in two cycles.

There are 4 instructions for memory access: two for read and two for write. All have an unsigned 4 bit immediate field in order to provide a group of 16 convenient memory slots off of a base address. The low read is signed, and there is no corresponding unsigned read instruction. The high read (like the high in-line literal) preserves the lower 16 bit value to load all 32 bits in two cycles. Write and write high instructions are provided to store 32 data bits in two cycles. Data and code space are shared, which enables the programmer to freely allocate and partition it, and enables the copying in of new code via this data read / write mechanism.

There are two forms of shifts, immediate and variable, so shifts distances / directions known at “compile time” can be more efficiently coded, and so variable shift distances are accommodated. The immediate unsigned shift performs a power of 2 (one hot bit) function when the shift distance is non negative, and unsigned right shift when the shift distance is negative. This dual functionality replaces the redundant left shift (signed and unsigned left shift produce identical results) with a useful secondary operation. For clarity, variable shifts are simply signed and unsigned - I felt it was important to keep the variable shifts unmixed so that there could not be unexpected behavior, and so sign testing of the shift variable would not be necessary. There is also a separate variable power function which only generates powers of 2 and is sign agnostic.

An immediate signed add is provided for small quick increments and decrements (+31/-32).

Arithmetic & Logical

Add, subtract, multiply, shift, and all of the logical operations have been described previously. Functions also performed by the logic unit are move / copy, flipping 32 bits end-for-end, negation of the sign bit, and the counting of leading zeros.

0	0	OP						PB	B				PA	A		
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	

Figure 29. General (non-immediate) instruction format.

Pop

By repurposing the entire stack selection / pop area of the opcode, the pop instruction is able to pop all, none, or any combination of stacks at once.

0	0	0	0	0	0	0	1	P7	P6	P5	P4	P3	P2	P1	P0
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 30. Pop instruction format.

Branches

There are three types of branches – jump, go to, and subroutine:

- **JMP** (jump) is relative to the PC and is either conditional or unconditional, and either immediate or not. The non-immediate unconditional and conditional forms jump a signed distance given by B unconditionally or if the test (A?0) is true. There are also conditional (A?B) signed and unsigned tests and (A?0) tests for the immediate JMP.
- **GTO** (go to) is absolute and unconditional. It loads the PC with the value given by B.
- **GSB** (go to subroutine) is absolute and unconditional. It loads the PC with the value given by B and stores the return address to A.

Relative branching is relative to the *next* instruction address, and not *this* instruction address, which seems like the most natural convention: a relative jump of 0 does nothing, a relative jump of +1 skips over the next instruction, and a relative jump of -1 is an infinite loop.

The only thing conditional about a conditional instruction is whether or not the branch is taken. Pops are *always* performed if pop bits are set in the conditional instruction.

Note that there is no explicit *return* operation, a GTO is used to return from subroutines and interrupts. The return address can be simultaneously popped at this point as well for cleanup.

Conditional Tests

Other than equality, the most useful conditional tests tend to split the numerical space under test in half via *sign* (less than zero, not less than zero), and *subtraction sign* (less than, not less than). I've found other random combinations of less than, equal to, and greater than testing to be less useful, and so are not directly supported in Hive. Odd / even testing could be included but the need for it seems less pressing, I ran out of immediate room, and it can be performed when necessary via a flip and sign test. (A?B) testing seems to happen more rarely than (A?0) testing, though it unfortunately consumes more Hive opcode space overall due to the necessary inclusion of the second operand selector. Less than and not less than comparisons need both signed and unsigned variants. Note that there is an unconditional jump, but there is no unconditional immediate jump – for this you can use the (A==B) immediate jump with both A and B fields pointing to the same stack (even an empty stack will work here), or use one of the zero test jumps pointing to a stack that has a known quantity for short unconditional immediate jumps.

Some conditional sign conventions / observations:

- All conditional comparisons of A to zero, L (A<Z), and NL (A!<Z) treat A as signed.
- All conditional comparisons of A and B that are signed | unsigned treat *both* A and B as signed | unsigned.
- The equivalency comparisons Z (A=0), NZ (A!=0), E (A=B), and NE (A!=B) are obviously sign neutral.

0x3			0	0	C	N	PB	B			PA	A			
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 31. (A?0) conditional jump instruction format.

Other Instructions

There are several stack and miscellaneous instructions:

- **PGC** pushes the current program counter (pointing to the next instruction) to A.
- **NOP** is a do nothing instruction, all functionality including pops is disabled.
- **REG** instructions allow access to the processor register set. Read signed, read high, write, and write high are supported.

Naming Conventions

Consistency is important with instruction naming conventions so that one can easily remember them or, failing that, quickly construct them knowing some basic rules. The letters “op_” precede all Hive instructions, and this is mainly to avoid conflict with verilog reserved words. After this is the three letter operation, usually followed (but not necessarily) by a second underscore and one or more option letters. Obviously not all operations are defined for all of the options.

Op_*	Function
nop	No Operation (no pops either)
pop	POPs : one hot bit per stack
pgc	ProGram Counter : A=PC
lit	In-line LIT eral data : A=mem(PC)
reg	REG ister operations : A=reg(B); reg(B)=A
cpy	CoPY : A=B
nsg	Negate SiGn : A={~B[31], B[30:0]}
not	Logical NOT : A=~B
flp	FLiP : A=B[0:31]
lzc	Leading Zero Count : A=lzc(B)
bra	Bit Reduction And : A=&B
bro	Bit Reduction Or : A= B
brx	Bit Reduction Xor : A^B
and	Logical AND : A=A&B
orr	Logical OR : A=A B
xor	Logical XOR : A=A^B
add	Arithmetic ADDITION : A=A+B
sub	Arithmetic SUB traction : A=A-B
mul	Arithmetic MULT iplication : A=A*B
shl	Shift Left : A=A<<(B or I)
pow	POWER : A=1<<B
jmp	JuMP : PC=PC+(B or I)
gto	Go TO : PC=B
gsb	Go SuB routine : PC=B, A=PC
mem	MEM ory operations : A=mem(B+I); mem(B+I)=A
dat	DAT a : A=I
psu	Power Shift right Unsigned

Figure 32. Instruction operations.

Op_ *_ ?	Function
i	Immediate
x	eX tended
r	Read
w	Write
n	Not
e	Equal
l	Less
z	Zero
s	Signed
u	Unsigned
h	High

Figure 33. Operation options.

The rules for these options are:

- No underscore between options.
- The option i if present comes first.
- The options x, r, or w if present come next.
- The conditional options n, e, l, and z if present come next, and in that order.
- The options s, u, or h if present go last.
- Some operations exist only in an immediate form (dat, mem) and the rule here is to always include the immediate option regardless.

Encoding

When assigning the actual numerical values to the instructions – the operational encoding or opcodes – it is important to make the decoding as straightforward and orthogonal as possible, which is much easier to say than to do. I use a spreadsheet to keep track of them, with a column for each control signal that is an output of the opcode decoder. This helps to reveal similar decoding patterns that may then be grouped together or otherwise advantageously arranged for ease of interpretation by the decoding logic.

Codes	Instruction	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
256	nop	0x0				0x0				-				-			
256	pop	0x0				0x1				P7	P6	P5	P4	P3	P2	P1	P0
256	pgc	0x0				0x4				PB	B			PA	A		
768	lit	0x0				1	0	U	H	PB	B			PA	A		
1024	reg	0x0				1	1	W	H	PB	B			PA	A		
256	cpy	0x1				0x0				PB	B			PA	A		
256	nsg	0x1				0x1				PB	B			PA	A		
256	not	0x1				0x2				PB	B			PA	A		
256	flp	0x1				0x4				PB	B			PA	A		
256	lzc	0x1				0x5				PB	B			PA	A		
256	bra	0x1				0x8				PB	B			PA	A		
256	bro	0x1				0x9				PB	B			PA	A		
256	brx	0x1				0xA				PB	B			PA	A		
256	and	0x1				0xC				PB	B			PA	A		
256	orr	0x1				0xD				PB	B			PA	A		
256	xor	0x1				0xE				PB	B			PA	A		
768	add	0x2				0	0	X	U	PB	B			PA	A		
768	sub	0x2				0	1	X	U	PB	B			PA	A		
768	mul	0x2				1	0	X	U	PB	B			PA	A		
512	shl	0x2				1	1	0	U	PB	B			PA	A		
256	pow	0x2				0xE				PB	B			PA	A		
1024	jmp (A?0)	0x3				0	0	L	N	PB	B			PA	A		
256	jmp	0x3				0xC				PB	B			PA	A		
256	gto	0x3				0xD				PB	B			PA	A		
256	gsb	0x3				0xE				PB	B			PA	A		
16384	mem_i	0	1	W	H	IM[3:0]			PB	B			PA	A			
24576	jmp_i (A?B)	1	C		N	IM[3:0]			PB	B			PA	A			
4096	jmp_i (A?0)	0xE				L	N	IM[5:0]			PA	A					
1024	dat_i	0xF				0	0	IM[5:0]			PA	A					
1024	add_i	0xF				0	1	IM[5:0]			PA	A					
1024	shl_is	0xF				1	0	IM[5:0]			PA	A					
1024	psu_i	0xF				1	1	IM[5:0]			PA	A					

Figure 34. Opcode encoding.

As seen in the table above, the single instructions are arranged by functionality in four groups of 16, with the first group something of a catch-all, the logical group second, arithmetic third, and branching fourth. The immediates follow, with memory access first, (A?B) jumps next, (A?0) jumps following, and immediate data, add, and shifts bringing up the rear. Most instructions are conveniently segmented into hex fields, which facilitates reading and interpretation of them.

Only 40 of the 64 single instruction codes are actually assigned, leaving 24 remaining. Opcode space utilization is therefore:

$$(2^{16} - 24 \cdot (2^8)) / (2^{16}) = 90.625\%$$

So there are some opcode slots open for future expansion, but the opcode space is otherwise largely consumed by instructions with immediate fields, and this probably is as it should be for a processor with compact opcodes.

It can be a long road leading up to the final selection of operations and their encoding, with much inserting and deleting of operations, resizing of immediate fields, and reshuffling of the encoding space, and I'm not sure there is any way to abbreviate this activity and still really do it justice. A compact opcode can be a rather harsh mistress due to the cramped working environment.

The use of a multi-file text search and replace tool such as TextCrawler (Digital Volcano) during development can be quite useful for making global changes to opcode names, etc. reaching all the way into the boot code text as well.

INTERNAL REGISTER SET

Any processor core will need a local, or internal register set to manage things like the reporting of basic operational errors, enabling and disabling of interrupts, general purpose I/O communications, timers, UARTs, watchdog sanity timers, shoot yourself in the head resets, etc. But register set implementation can be a dull, repetitive, and bug prone exercise. To automate this to some degree and to reduce the chance of errors, at the foundation of the Hive register set is a configurable multi-function single base register component with many parameter-based options.

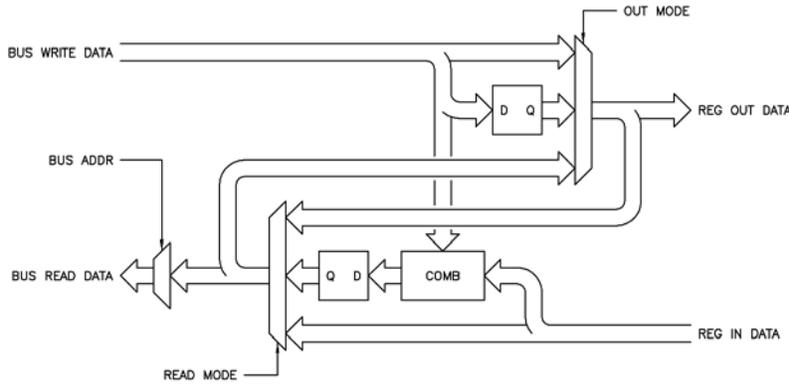


Figure 35. Configurable base register component.

The figure above shows a schematic view of the single base register component. On the left is the common processor expansion bus, on the right are the individual inputs and outputs of the single register. Not shown is logic that detects the address match, nor the read / write / in / out signal conditioning. The two large multiplexers set the read and output modes via parameters at build time. Any number and combination of bits can be “live” (provided with functional logic) and initialized to a known value at reset. Register input data can be optionally resynchronized and/or made edge sensitive. Most common register types can be formed via various combinations of the modes, most others can be implemented by adding a bit of circuitry to this base construct. Mixed mode bits in a single register aren’t directly supported.

Parameter	Output Mode
“ZERO”	zero output
“THRU”	no latch, direct connect
“LTCH”	write latch
“LOOP”	no output latch, output selected read data

Figure 36. Output mode options.

Parameter	Read Mode
“THRU”	no latch, direct connect
“CORD”	set on input one, clear on read
“COW1”	set on input one, clear on write one
“DFFE”	D type flip flop with enable
“LOOP”	no read latch, read selected out data

Figure 37. Read mode options.

Multiple base register components are assembled into a register set by using a big OR gate to combine their bus side data read port bits. Placement of each individual base register within the register set address space is governed by an input parameter to each base register component. Access to the register set is via special instructions, a shim component, and the internal ALU mux.

Unlike their ASIC brethren, one nice thing that soft processor cores have going for them is they don’t need huge gobs of configuration registers. Want a timer? Write one in verilog, connect it to an interrupt, and it’s game over. Want 8 UARTs? No problem – you don’t have to put them in until you really need them and you can take them out later when you don’t.

The Hive internal register set is located at the bottom of the address space (base 0x0000) and includes the following basic functionality:

Decode:

- 0x0 : Core version register - ver_reg
- 0x1 : Thread ID register - thrd_id_reg
- 0x2 : Clear register - clr_reg
- 0x3 : Interrupt enable register - intr_en_reg
- 0x4 : Opcode error register - op_er_reg
- 0x5 : Stack error register - stk_er_reg
- 0x6 - 0x7 : UNUSED
- 0x8 : I/O low register - io_lo_reg
- 0x9 : I/O high register - io_hi_reg
- 0xA - 0xB : UNUSED
- 0xC : UART RX register - uart_rx_reg
- 0xD : UART TX register - uart_tx_reg
- 0xE - 0xF : UNUSED

=====

- 0x0 : Core version register - ver_reg

bit	name	description
7-0	ver_min[7:0]	minor version info
15-8	ver_maj[7:0]	major version info

Notes:

- Read-only.
- Nibbles S/B BCD (0-9; no A-F) to be easily human readable, and to eliminate confusion between decimal and hex here.
- Major version changes when op_code binary decode changes (incompatibility).

=====

- 0x1 : Thread ID register - thrd_id_reg

bit	name	description
2-0	thrd_id[2:0]	thread ID
15-3	-	0000000000000

Notes:

- Read-only.
- Threads can read this to discover their thread ID.

=====

- 0x2 : Clear register - clr_reg

bit	name	description
7-0	clr[7:0]	0=>1 clear thread; 1=>0 no effect;
15-8	-	00000000

Notes:

- Read / write.
- Per thread clearing.
- All bits cleared on async reset.

=====

- 0x3 : Interrupt enable register - intr_en_reg

bit	name	description
7-0	intr_en[7:0]	1=thread interrupt enable; 0=disable
15-8	-	00000000

Notes:

- Read / write.
- Per thread enabling of interrupts.
- All bits cleared on async reset.

=====
- 0x4 : Opcode error register - op_er_reg

bit	name	description
7-0	op_er[7:0]	1=opcode error; 0=OK
15-8	-	00000000

Notes:
- Clear on write one.
- Per thread opcode error reporting.

=====
- 0x5 : Stack error register - stk_er_reg

bit	name	description
7-0	pop_er[7:0]	1=lifo pop when empty; 0=OK
15-8	push_er[7:0]	1=lifo push when full; 0=OK

Notes:
- Clear on write one.
- Per thread LIFO stack error reporting.

=====
- 0x6 - 0x7 : UNUSED

- 0x8 : I/O low register - io_lo_reg

bit	name	description
15-0	io_lo[15:0]	I/O data

Notes:
- Separate read / write.
- Reads of io_lo_reg freeze data in io_hi_reg, so read io_lo_reg first then read io_hi_reg for contiguous wide (32 bit) data reads.
- Writes function normally.

=====
- 0x9 : I/O high register - io_hi_reg

bit	name	description
15-0	io_hi[15:0]	I/O data

Notes:
- Separate read / write.
- Reads of io_lo_reg freeze data in io_hi_reg, so read io_lo_reg first then read io_hi_reg for contiguous wide (32 bit) data reads.
- Writes function normally.

=====
- 0xA - 0xB : UNUSED

- 0xC : UART RX register - uart_rx_reg

bit	name	description
7-0	uart_rx_data[7:0]	RX UART data
8	rx_rdy	1=RX UART ready (has new data); 0=not ready
15-9	-	0000000

Notes:
- Reads from this register pop data from the RX UART.
- To avoid RX data loss, read soon after RX UART is ready.
- UART ready bit will self clear after associated register operation.

=====

- 0xD : UART TX register - uart_tx_reg

bit	name	description
7-0	uart_tx_data[7:0]	TX UART data
8	tx_rdy	1=TX UART ready (for new data); 0=not ready
15-9	-	0000000

Notes:

- Writes to this register push data to the TX UART.
- To avoid TX data loss, restrict writes to when TX UART is ready.
- UART ready bit will self clear after associated register operation.

=====

- 0xE - 0xF : UNUSED

=====

VERIFICATION

Job #1 when building a processor is obviously wringing out all the bugs. Processors that have caches, pipeline hazards and stalls, and lots of internal state are notoriously difficult to verify (and therefore fundamentally trust – Pentium division bug anyone?). Much of engineering is the exercise of complexity management, and processor architectures themselves should be governed by this guiding principle as well. Simplicity allows one to more easily juggle the processor model in one's head, but can also greatly ease the verification problem. Hive has relatively simple control structures, minimal internal state, and the entire design is partitioned into right sized modules that are as self-contained as possible, making verification a straightforward and relatively painless task.

It goes without saying that all basic blocks should be fully tested before being assembled together. With Hive, most module port widths and associated internal logic are parameterized so, for example, full verification of the ALU may be accomplished by shrinking the data port widths to a trivial size and manually examining the results of all possible inputs. The multiplier may be verified separately at full width by comparing its results to a second naively instantiated multiplier, both supplied with corner cases and random input (the inclusion of this test hardware is a parameterized option for the multiplier base module). The intermediate control and data ring constructs allow for the testing of lower level aggregate functionality.

Once basic functionality is up (thread clearing, immediate byte, jumps) boot code enables the processor to essentially verify itself. Stack functioning and error reporting may be fully tested for all threads. Jump distances and all associated conditionals should be confirmed. Each opcode should be tested to make sure it is decoded and functioning correctly – distinctive signatures may be used here rather than exhaustive testing. This is also a good way to get some early experience hand coding the processor (the point at which I've become largely disillusioned with my past designs) which can often lead to changes in the op codes or other parts of the fundamental design.

Finally, several simple algorithms should coded up, first in a spreadsheet and then in the simulated core, with the results compared. When working on this phase of the design I find that I have to fight a strong inclination to tailor the instruction set to the algorithm du jour and keep my eye on the big picture. For instance, after coding the log₂ algorithm, a leading zero count instruction (lzc) seemed like it would be a valuable addition (it has floating point normalization uses as well). So I coded up a fully parameterized verilog module and speed / functionally tested it. But only after I recognized the general use of this function did I include it in the logical unit of the Hive ALU.

SPEED

Job #2 when building a processor is getting the most speed out of it as possible. To this end most Hive modules have configurable registering on inputs and outputs which can effectively isolate timing to the fabric rather than the FPGA I/O pins when doing individual module speed trial builds. The component *pipe.v* can go from a single wire to any desired width and registering depth, and is used throughout the design for general registering and pipelining. (A downside to this approach is that useful internal signal names get reduced to vector indexes, which can make them difficult to differentiate in simulation.)

It is important during early testing to identify the slowest low level hardware path. This then is the lower speed target for the remaining circuitry, which should be written / implemented at least 10% or so faster so as to have a bit of a cushion when it all comes together – the more margin the better because modules have a tendency to slow down considerably when spattered willy nilly onto the fabric with all of the other logic. Just as C compilers can often beat the best human hand coders (particularly when targeting x86 and other insanely complex processors) there are various FPGA synthesis options that will likely produce a faster top speed, and an automated seed hunt with multiple options (e.g. Altera's "Design Space Explorer") will usually produce a faster point in the design space if you've got some time in your life to spare. This is worth doing if only to know the top speed easily attainable.

Watch the fitter resource allocation like a hawk, particularly for any extra block RAM creeping into your design. It seems the synthesis / fitter likes to replace pipe stages with block RAM, which can sometimes slow things down.

Use a DCM to get your board clock up to the maximum speed of the core if you want performance, or keep the clock speed lower to conserve power.

PROGRAMMING

The verilog hardware description language has an “initial” construct that can be used along with other verilog syntax features to write fairly legible boot code, comments and all. Hive boot code text resides in a text file (boot_code.h) that gets inserted into the main memory module with an include statement. Let's take a look at some sample boot code:

```
`include "boot_code_defs.h"
```

This include pulls in our opcode encoding and internal address register locations so we can refer to them by name rather than by their rather cryptic numerical encoding.

```
integer i;  
initial begin
```

The above declares an integer we'll use to keep from having to name each and every address, and marks the beginning of the initialization code.

A Simple Example

```
// clr space //  
i='h0; ram[i] = { `lit_u,          `__, `s7 }; // s1='h0040  
i=i+1; ram[i] = {                16'h0040 ; //  
i=i+1; ram[i] = { `gsb,          `P7, `s3 }; // gsb, pop s7  
i=i+1; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
// all others : loop forever  
i='h04; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
i='h08; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
i='h0c; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
i='h10; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
i='h14; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
i='h18; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
i='h1c; ram[i] = { `jmp_ie,      -4'd1, `s0, `s0 }; // loop forever  
  
// intr space //  
  
// code & data space //  
  
// sub : read core version & write to GPIO, return to (s3)  
i='h40; ram[i] = { `dat_is,        `VER, `s1 }; // s1=reg addr  
i=i+1; ram[i] = { `reg_rs,        `P1, `s0 }; // s0=(s1), pop s1  
i=i+1; ram[i] = { `dat_is,        `IO_LO, `s1 }; // s1=reg addr  
i=i+1; ram[i] = { `reg_w,         `P1, `P0 }; // (s1)=s0, pop both  
i=i+1; ram[i] = { `gto,           `P3, `__ }; // return, pop s3
```

The first line is located at address 0, which is where thread 0 vectors to when cleared. The instruction puts an unsigned literal in S7, the value of which is the address of a subroutine. The second line is the unsigned in-line literal value, 0x0040. The third line calls the subroutine and pushes the return address to S3, and it simultaneously pops the subroutine address in S7 (stack cleanup). The fourth line is an immediate jump -1, which is an infinite loop.

The next seven lines are for threads 1 through 7, which are instructed to twiddle their thumbs by looping infinitely. Note that the clear addresses are spaced 4 apart (both this distance and the base address are configurable at build time for the clear and interrupt vector groups). The interrupt instruction address space is blank because the interrupts won't be enabled nor used for this program.

The subroutine code at address 0x40 loads the processor version register address to S1, reads the core version and pops the address, loads the processor lower I/O register address to S1, then writes the core version to the lower I/O port, pops both address and data simultaneously with the write (stack cleanup), then issues a gto S3 and pops S3 (stack cleanup), which is the way subroutines and interrupts are returned from in Hive.

Binary Search Division Subroutine Example

A somewhat meatier example is division. The binary search division algorithm is shown below:

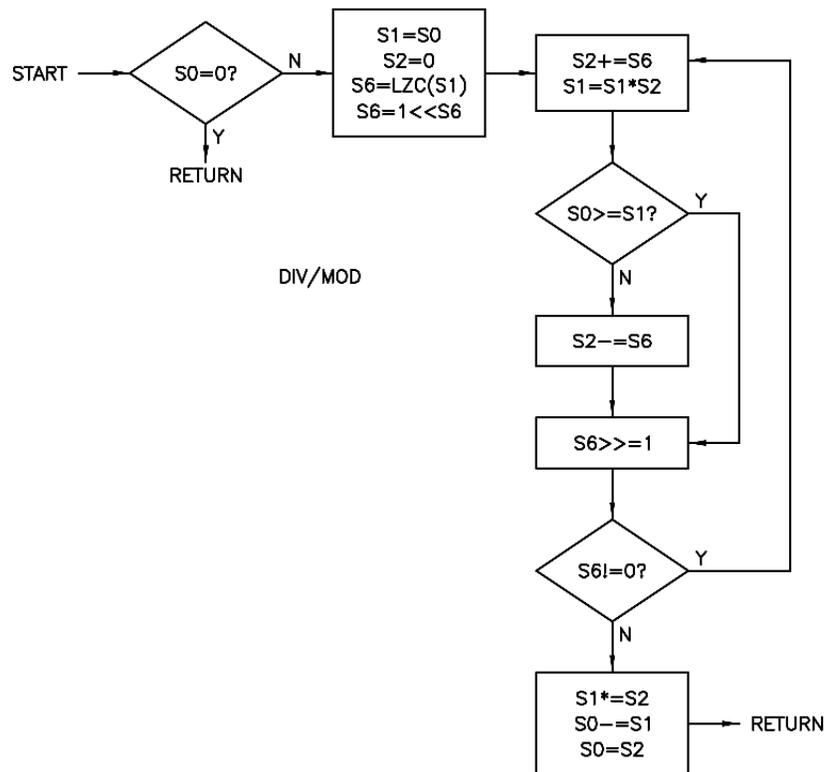


Figure 38. Binary search division algorithm flow chart.

This algorithm divides two unsigned inputs using a binary search. The idea is to light up one-hot bits going from MSB to LSB in a trial number Q , multiply it by the denominator D , and compare it to the numerator N .

Setting the one-hot (OH) start value to $LZC(D)$ prevents internal overflow at $D*(Q+OH)$ and speeds up the average case by reducing the number of loops.

The shifted one-hot value one used in the loop can also be used as the loop counter: when the one is completely shifted out (the vector = 0) exit the loop. This saves one step in the loop and one storage register.

```

// sub : unsigned divide & modulo remainder, return to (s7)
//
// algorithm: binary search
//
// s0 : N, D(top)/N(under) input, Q(top)/R(under) output
// s1 : D
// s2 : Q
// s3 :
// s4 :
// s5 :
// s6 : one-hot (& loop test)
// s7 : sub return address
//
// (D=0)? is an error, return
i='h90; ram[i] = { `jmp_inz,      6'd1, `s0 }; // (s0!=0) ? skip return
i=i+1; ram[i] = { `goto,        `P7, `__ }; // return to (s7), pop s7
// loop setup
i=i+1; ram[i] = { `cpy,         `P0, `s1 }; // s0=>s1 (s1=D, s0=N)
i=i+1; ram[i] = { `dat_is,     6'd0, `s2 }; // s2=0 (s2=init Q)
i=i+1; ram[i] = { `lzc,       `s1, `s6 }; // s6=lzc(s1)
i=i+1; ram[i] = { `pow,       `s6, `P6 }; // s6=1<<s6, pop s6 (s6=init OH)
// loop start
i=i+1; ram[i] = { `add,        `s6, `P2 }; // s2+=s6 (s2=new trial Q)
i=i+1; ram[i] = { `mul,        `s2, `s1 }; // s1=s1*s2 (s1=D*Q)
// jump start
i=i+1; ram[i] = { `jmp_inlu,   4'd1, `P1, `s0 }; // (s0>=s1) ? skip restore, pop s1 (N>=D*Q)
i=i+1; ram[i] = { `sub,        `s6, `P2 }; // s2-=s6 (s2=restored Q)
// jump end
i=i+1; ram[i] = { `psu_i,     -6'd1, `P6 }; // s6>=1 (new OH)
i=i+1; ram[i] = { `jmp_inz,   -6'd6, `s6 }; // (s6!=0) ? do again
// loop end
// calc remainder, move Q
i=i+1; ram[i] = { `mul,       `s2, `P1 }; // s1*=s2 (s1=D*Q)
i=i+1; ram[i] = { `sub,       `P1, `P0 }; // s0-=s1, pop both (s0=N-D*Q=R)
i=i+1; ram[i] = { `cpy,       `P2, `s0 }; // s0=s2, pop s2 (s0=Q)
// return
i=i+1; ram[i] = { `goto,      `P7, `P6 }; // return to (s7), pop s7 & s6
// end sub

```

The subroutine code is above. The return is not skipped if the denominator value is zero, which is mathematically undefined (+/- infinity). The denominator is pushed to S1, S2 is initialized to zero, the LZC of the denominator is converted to a one-hot value and is pushed to S6, and the loop begins.

The one-hot value is added to the quotient, which is then multiplied by the denominator and compared to the numerator. If greater, the one-hot value is subtracted from the quotient (restoring it). The one-hot value is shifted right once and the loop is performed again until the one-hot one is shifted out of the LSB position, leaving all zeros and the loop is exited.

Finally the remainder is calculated and pushed onto S0, the quotient is pushed on top of it, and some cleanup is performed at subroutine return.

In terms of real time, assuming the denominator isn't zero and the return is skipped, it takes 5 cycles to test the input and setup the loop, 5 cycles per loop best case and 6 cycles worst case, with 4 cycles after the loop. For 32 worst case iterations with all being 6 worst case loops this gives:

$$5 + 6 \cdot 32 + 4 = 201 \text{ cycles worst case}$$

For a 200 MHz clock and 8 clocks per cycle, this is 8.04 us worst case.

Binary Search SQRT Subroutine Example

A somewhat similar example to division is the calculation of the square root shown below. The binary search algorithm in a processor that has single cycle unsigned multiply is likely faster than other methods due to fewer necessary loop instructions:

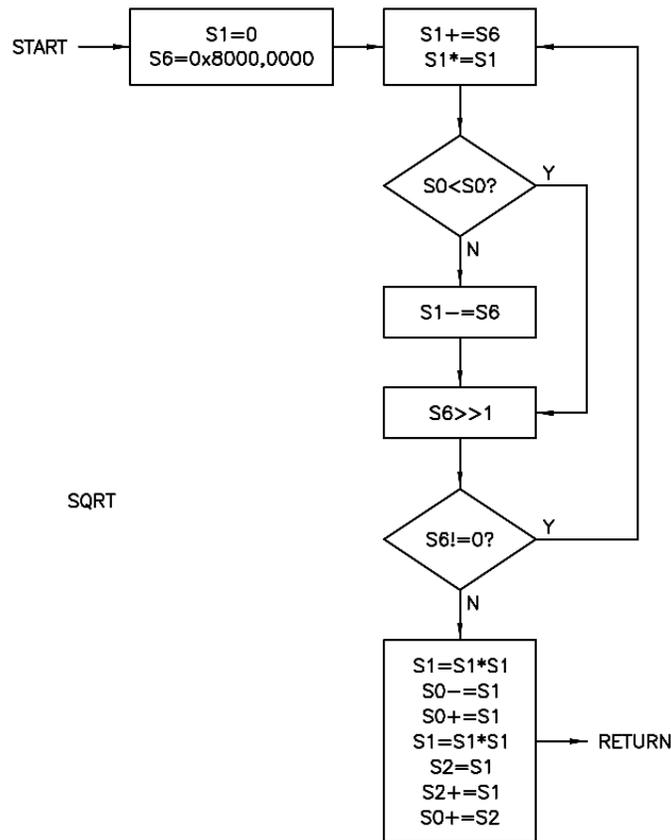


Figure 39. Binary search square root algorithm flow chart.

The idea is to light up the bits going from MSB to LSB in a trial number, square it, and compare it to the input. If the input x is an 8 bit unsigned integer (i.e. $N=8$) then the output needs at most 4 bits for the integer portion, and carrying out the looping 4 more times gives a result with 4 decimal places. So N bits in gives $n.n$ bits out where $n = N/2$.

Note that picking the upper half of the square result will truncate the lower half, or the decimal portion. The negative influence of this can be largely ameliorated by the use of ' $<$ ' rather than ' $<=$ ' for comparison to the input, and this gives an underestimated result which is too small by at most 1 LSB. Computing the squared result and the square of the incremented result, and picking the one closest to the input corrects this and obviously adds complexity, but provides a significant benefit by making the loop smaller.

Since the loop underestimates q , we know the q is either OK or needs at most 1 LSB of value added to it. The generic expression to evaluate this is:

$$\text{abs}\{x - q^2\} < \text{abs}\{x - (q+\text{LSB})^2\} ?$$

If the result is true we pick q as the output, if false we pick $q+\text{LSB}$ ($= q++$).

Since $\text{LSB} > 0$, then $q < q+\text{LSB}$, and $q^2 < (q+\text{LSB})^2$. So we don't need to use absolute value functions to evaluate the expressions. Also, q^2 will always be in the range $[x-2:x]$ and $(q+\text{LSB})^2$ will always be in the range $[x:x+2)$.

What we really want to know is the influence of adding $\text{LSB}/2$ to the input. We can square the result and compare it to the input to know whether or not to increment q .

$$(q + \text{LSB}/2)^2 \geq x ?$$

Expanding the above gives:

$$q^2 + 2q\text{LSB}/2 + (\text{LSB}/2)^2 \geq x ?$$

$$q^2 + q\text{LSB} + (\text{LSB}/2)^2 \geq x ?$$

Since $(\text{LSB}/2)^2$ is very tiny we can safely discard it which gives:

$$q^2 + q\text{LSB} \geq x ?$$

The term $q\text{LSB}$ is a simple right shift of n ; even more simply it is the uncorrected value of q as it sits in the processor register! This gives:

$$q^2 + q \gg n \geq x ?$$

Which can be rearranged as:

$$q \gg n \geq x - q^2 ?$$

If true we pick q as the result, otherwise we pick $q + \text{LSB}$ ($= q++$).

We can actually eliminate the above conditional:

$$Q = q + x - \text{integer}[q^2] - \text{int}[q \gg n + \text{decimal}[q^2]]$$

Note that the shifted one-hot one used in the loop can also be used as the loop counter: when the one is completely shifted out (the vector = 0) exit the loop. This saves one storage register, but more importantly eliminates one step in the loop.

```

// sub : s0=sqrt(s0), return to (s7)
//
// input is unsigned 32 integer
// output is unsigned 16.16 integer
//
// algorithm: binary search
// iterate 32 times
//
// s0 : input (x), output (Q)
// s1 : running q
// s2 : temp: q^2 + q
// s3 :
// s4 :
// s5 :
// s6 : one-hot (& loop test)
// s7 : sub return address
//
// loop setup
i='h90; ram[i] = { `dat_is,          6'd0, `s1 }; // s1=0 (init Q)
i=i+1; ram[i] = { `psu_i,          6'd31, `s6 }; // s6 MSB=1 (init OH)
// loop start
i=i+1; ram[i] = { `add,            `s6, `P1 }; // s1+=s6
i=i+1; ram[i] = { `mul_xu,        `s1, `s1 }; // s1=s1*s1 (square, integer portion)
// jump start
i=i+1; ram[i] = { `jmp_ilu,       4'd1, `s0, `P1 }; // (s1<s0) ? jump 1 pop s1 (skip restore)
i=i+1; ram[i] = { `sub,           `s6, `P1 }; // s1-=s6 (restore)
// jump end
i=i+1; ram[i] = { `psu_i,         -6'd1, `P6 }; // s6>>=1 (new OH)
i=i+1; ram[i] = { `jmp_inz,       -6'd6, `s6 }; // (s6!=0) ? do again
// loop end
i=i+1; ram[i] = { `mul_xu,        `s1, `s1 }; // s1=s1*s1 (square, integer portion)
i=i+1; ram[i] = { `sub,           `P1, `P0 }; // s0-=s1, pop s1 : x -= q^2
i=i+1; ram[i] = { `add,           `s1, `P0 }; // s0+=s1 : Q = q + x - q^2
i=i+1; ram[i] = { `mul,          `s1, `s1 }; // s1=s1*s1 (square, decimal portion) : (q>>n)^2
i=i+1; ram[i] = { `cpy,          `P1, `s2 }; // s2=s1, move
i=i+1; ram[i] = { `add_xu,       `P1, `P2 }; // s2+=s1 (carry out, integer portion), pop s1 :
int[q>>n + (q>>n)^2]
i=i+1; ram[i] = { `sub,          `P2, `P0 }; // s0+=s2, pop s2 : Q = q + (x - q^2) - int[q>>n
+ (q>>n)^2]
i=i+1; ram[i] = { `goto,         `P7, `P6 }; // return, pop s7 & s6
// end sub

```

The subroutine code is above. S2 is initialized to zero, the initial one-hot value and is pushed to S6, and the loop begins.

The one-hot value is added to q, which is then squared and compared to the input value. If greater than or equal, the one-hot value is subtracted from q (restoring it). The one-hot value is shifted right once and the loop is performed again until the one-hot one is shifted out of the LSB position, leaving all zeros and the loop is exited.

Finally q is corrected and pushed to S0, and some cleanup is performed at subroutine return.

In terms of real time, it takes 2 cycles to setup the loop, 5 cycles per loop best case and 6 cycles worst case, with 8 cycles after the loop. If all 32 iterations are 6 worst case loops this gives:

$$2 + 6 \cdot 32 + 8 = 202 \text{ cycles worst case}$$

For a 200 MHz clock and 8 clocks per cycle, this is 8.08 us worst case.

Log₂ Subroutine Example

Presented here is the calculation of the 32 bit base 2 logarithm of an unsigned 32 bit input number. The algorithm shown exploits the fact that $\log_2(x^2) = 2*\log_2(x)$, and is implemented by a looped squaring processes.

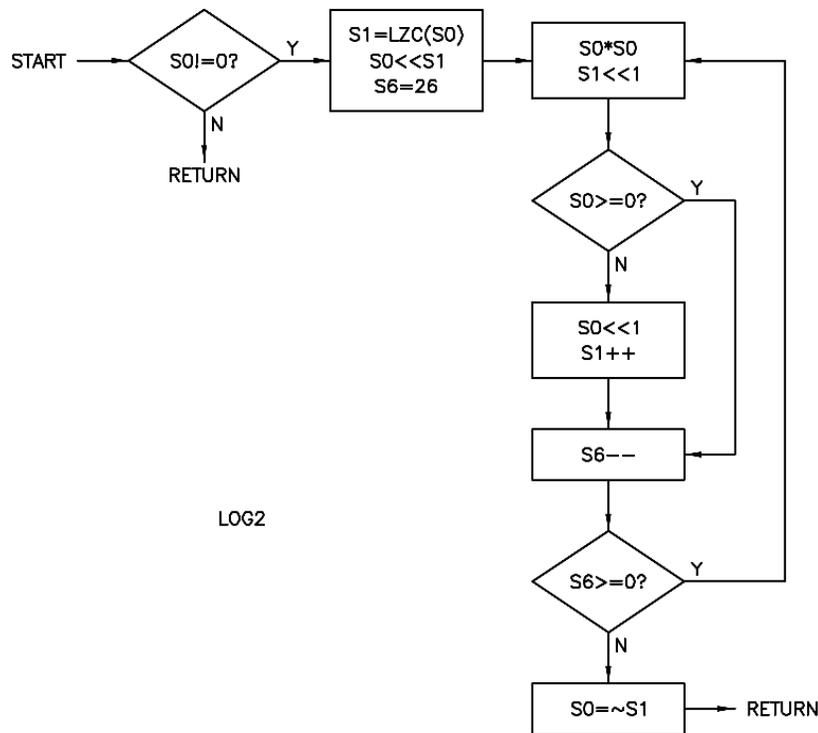


Figure 40. Log base 2 algorithm flow chart.

This algorithm is shown above as a flow chart. The initial test makes sure the input is non-zero because the log of zero is undefined. Next is the input normalization. Finally we have the square / mantissa loop, with the loop exit test and result negation at the end.

The first section normalizes the input by shifting it to the left until the MSB is equal to 1 (which tests as negative). The number of shifts necessary to accomplish this is subtracted from 31 to form the log characteristic, which is the 5 bit number to the left of the decimal place in the result.

After normalization, the normalized input is squared and the resulting MSB examined. If it is equal to 1 then a 1 is left shifted into the characteristic. If it is equal to 0 both the characteristic and the squared input are left shifted once (which should make the characteristic LSB = 0 and the squared input MSB = 1). This loop is executed $32 - 5 = 27$ times to find all bits to the right of the decimal place in the result, AKA the log mantissa.

It is actually possible to skip the input normalization subtraction and use the same jump test within the loop for both operations that are conditionally jumped over, then simply negate the result at the end. This somewhat tricky but more efficient algorithm is the one implemented.

```

// sub : s0=log2(s0), return to (s7)
//
// input is in[31:0] an unsigned 32 bit integer
// output is c[31:27].m[26:0] unsigned fixed decimal
//
// s0 : input, normalize, square, output
// s1 : lzc, result
// s2 :
// s3 :
// s4 :
// s5 :
// s6 : loop index
// s7 : sub return address
//
// (input=0)? is an error, return
i='h90; ram[i] = { `jmp_inz,          6'd1, `s0 }; // (s0!=0) ? skip return
i=i+1; ram[i] = { `gto,             `P7, `__ }; // return to (s7), pop s7
// normalize
i=i+1; ram[i] = { `lzc,              `s0, `s1 }; // s1=lzc(s0)
i=i+1; ram[i] = { `shl_s,           `s1, `P0 }; // s0<=s1 (normalize)
// loop setup
i=i+1; ram[i] = { `dat_is,          6'd26, `s6 }; // s6=26 (loop index)
// loop start
i=i+1; ram[i] = { `mul_xu,          `s0, `P0 }; // s0*=s0
i=i+1; ram[i] = { `shl_is,          6'd1, `P1 }; // s1<=1
// jump start
i=i+1; ram[i] = { `jmp_ilz,         6'd2, `s0 }; // (s0[31]==1) ? jump
i=i+1; ram[i] = { `shl_is,          6'd1, `P0 }; // s0<=1
i=i+1; ram[i] = { `add_is,          6'd1, `P1 }; // s1++
// jump end
i=i+1; ram[i] = { `add_is,          -6'd1, `P6 }; // s6--
i=i+1; ram[i] = { `jmp_inlz,        -6'd7, `s6 }; // (s6>=0) ? do again
// loop end
// cleanup, return
i=i+1; ram[i] = { `not,              `P1, `P0 }; // s0=~s1, pop both
i=i+1; ram[i] = { `gto,              `P7, `P6 }; // return, pop s7 & s6
// end sub

```

The subroutine code is above. The return is not skipped if the input value is zero. Input normalization shifts the input value to the left until the MSB is 1, the number of shifts necessary to do this gives the inverse of the characteristic.

The square loop uses a shift and a conditional immediate add to left shift either a 0 or 1 into the mantissa LSB. Unsigned extended multiplication is the operation used for squaring. After the loop has completed, the result in S1 is negated and copied to S0 with both popped to form a move. The return address and loop index are both popped at return to complete the cleanup.

In terms of real time, assuming the input isn't zero and the return is skipped, it takes 4 cycles to test the input, normalize it, and setup the loop, 5 cycles per loop best case and 7 cycles worst case, with 2 cycles after the loop. For 26 iterations with all being 7 worst case loops this gives:

$$4 + 7*26 + 2 = 188 \text{ cycles worst case}$$

For a 200 MHz clock and 8 clocks per cycle, this is 7.52 us worst case.

Exp₂ Subroutine Example

Presented is the calculation of the 32 bit base 2 exponentiation of an unsigned 32 bit fixed decimal input.

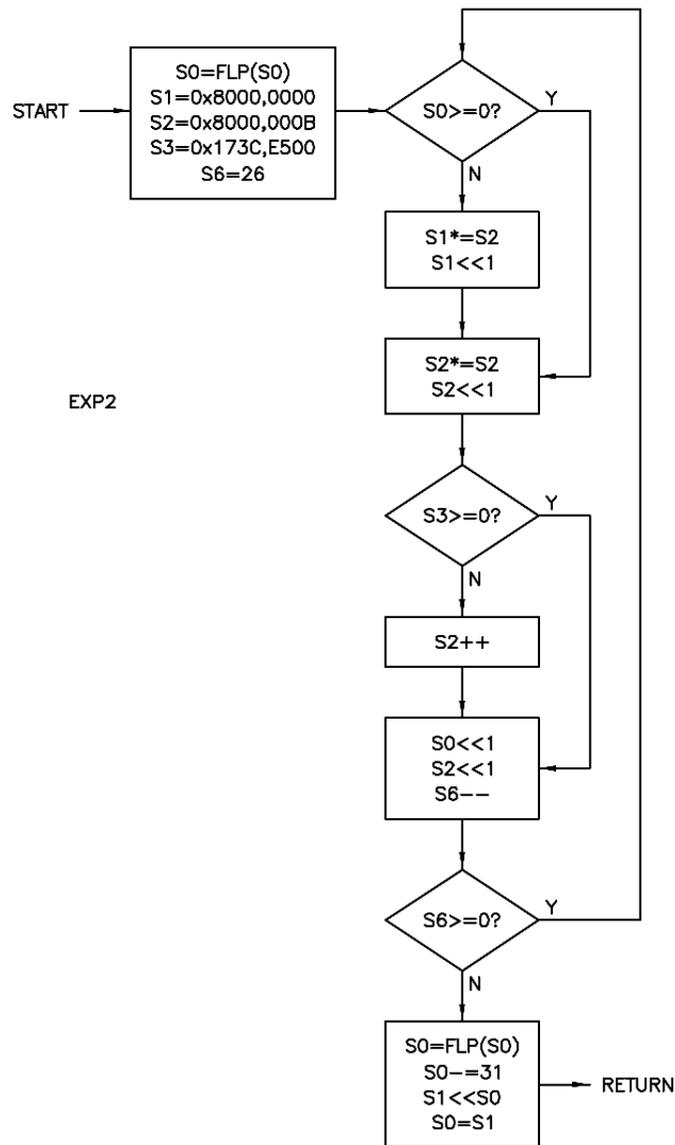


Figure 41. Exponential of 2 algorithm flow chart.

The idea is that exponentiation (2^n) where n is a real positive fixed point number is straightforward because binary numbers are themselves constructed of powers of 2.

For example:

$$5.5 = 101.1 = (1 \cdot 2^2) + (0 \cdot 2^1) + (1 \cdot 2^0) + (1 \cdot 2^{-1}).$$

So

$$2^{5.5} = 2^{[(1 \cdot 2^2) + (0 \cdot 2^1) + (1 \cdot 2^0) + (1 \cdot 2^{-1})]} = (2^{(2^2)}) \cdot (2^{(2^0)}) \cdot (2^{(2^{-1})}) = 16 \cdot 2 \cdot 1.414... = 45.254...$$

So we need to find (and selectively multiply together) successive square roots of 2 (difficult), or start at some $2^{(2^{-m})}$ root of 2 and square our way up (much easier but prone to error due to repeated squaring of the limited resolution base number).

For a 32 bit input, log2 gives 5 bits of characteristic and $32 - 5 = 27$ bits of mantissa. If we use this fixed decimal form to "undo" the log2 with 2^n , then the input n to the 2^n function is:

$$\begin{aligned} [31:27] &= 2^4, \dots, 2^0 \\ [26:0] &= 2^{-1}, \dots, 2^{-27} \end{aligned}$$

Therefore, the first root we need is $2^{(2^m)} = 2^{(2^7)} = 1.00000000516\dots$ which is approximated by the 32 bit number 0x8000000B. Squaring this and picking the upper 32 bits requires a left shift of one to re-align the result. Shifting truncates the LSB, and squaring doubles the initial error, so the results tend to be underestimated and deviate more from the ideal with each iteration. We can add a small (0 or 1 LSB) "fudge" factor after each multiply & shift to make up for this, and dramatically increase the overall precision. The fudge factor bits were arrived via trial and error with an eye towards minimizing both local and final error. Without the fudge factor full scale output is off by approx. -0.56%. With the fudge factor full scale output is good to approximately 8 significant (base 10) digits.

After we hit the square root of 2 (i.e. 2^{-1}) the remaining operations are simple right shifts which can be performed as a single bulk right shift (thus likely tossing away many calculated digits).

```

// sub : s0=exp2(s0), return to (s7)
//
// input is c[31:27].m[26:0] unsigned fixed decimal
// output is out[31:0] an unsigned 32 bit integer
//
// s0 : input, output
// s1 : running multiply
// s2 : running root
// s3 : fudge factor
// s4 :
// s5 :
// s6 : loop index
// s7 : sub return addr
//
// setup
i='h90; ram[i] = { `flp,          `s0, `P0 }; // flp(s0) (to examine lsbs via msb)
i=i+1; ram[i] = { `psu_i,      6'd31, `s1 }; // s1=0x8000,0000 (starting value = 1)
i=i+1; ram[i] = { `cpy,        `s1, `s2 }; // s2=0x8000,000b (starting root = 2^2^-27)
i=i+1; ram[i] = { `add_is,     6'hb, `P2 }; //
i=i+1; ram[i] = { `lit_u,      `_, `s3 }; // s3=0x173c,e500 (fudge factor bits)
i=i+1; ram[i] = {          16'he500 ; //
i=i+1; ram[i] = { `lit_h,      `_, `P3 }; //
i=i+1; ram[i] = {          16'h173c ; //
i=i+1; ram[i] = { `dat_is,     6'd26, `s6 }; // s6=26 (loop index)
// loop start
// jump 0 start
i=i+1; ram[i] = { `jmp_inlz,   6'd2, `s0 }; // (s0[31]==0) ? jump +2 (skip running mult)
i=i+1; ram[i] = { `mul_xu,     `s2, `P1 }; // s1*=s2
i=i+1; ram[i] = { `shl_is,     6'd1, `P1 }; // s1<<=1 (so msb=1)
// jump 0 end
i=i+1; ram[i] = { `mul_xu,     `s2, `P2 }; // s2*=s2 (square to get next root)
i=i+1; ram[i] = { `shl_is,     6'd1, `P2 }; // s2<<=1 (so msb=1 & lsb=0)
// jump 1 start
i=i+1; ram[i] = { `jmp_inlz,   6'd1, `s3 }; // (s3[31]==0) ? jump +1 (no fudge bit)
i=i+1; ram[i] = { `add_is,     6'd1, `P2 }; // s2++ (set lsb of running root)
// jump 1 end
i=i+1; ram[i] = { `shl_is,     6'd1, `P0 }; // s0<<=1 (get next input bit)
i=i+1; ram[i] = { `shl_is,     6'd1, `P3 }; // s2<<=1 (get next fudge bit)
i=i+1; ram[i] = { `add_is,     -6'd1, `P6 }; // s6-- (loop index--)
i=i+1; ram[i] = { `jmp_inlz,   -6'd11, `s6 }; // (s6>=0) ? jump -11 (loop again)
// loop end
// final shift
i=i+1; ram[i] = { `flp,        `s0, `P0 }; // flp(s0) (flip remaining bits)
i=i+1; ram[i] = { `add_is,     -6'd31, `P0 }; // s0--=31
i=i+1; ram[i] = { `shl_u,      `P0, `P1 }; // s1<<=s0, pop s0
// cleanup, return
i=i+1; ram[i] = { `cpy,        `P1, `s0 }; // s0=s1, pop s1 (move)
i=i+1; ram[i] = { `pop,        8'b01000110 }; // pop s2, s3, s6
i=i+1; ram[i] = { `gt0,        `P7, `_; // return, pop s7
// end sub

```

The subroutine code is above. The input is flipped to test the LSBs sequentially via the MSB (the “sign” bit). The running multiply is initialized to ‘1’ and pushed to S1. The initial running root of 2 is pushed to S2. The fudge factor is pushed to S3. The loop index is pushed to S6.

There are two conditional jumps in the loop. The first tests the input LSB and skips the running multiplication of the current root of two and shift if the bit is zero. After this the next root is found via squaring and shifting, and the second test applies the fudge bit. Then the input and fudge vectors are shifted left once to expose the next bits, the loop index is decremented, and the loop repeats until the loop index goes negative.

When the loop is exited, the remaining input bits are flipped back, -31 is added to them, and the running multiply is shifted left by this amount. The result is pushed to S0, some cleanup occurs, and the subroutine returns.

In terms of real time, it takes 7 cycles to setup the loop, 8 cycles per loop best case and 11 cycles worst case, with 6 cycles after the loop. For 26 iterations with all being 11 worst case loops this gives:

$$7 + 11 * 26 + 6 = 299 \text{ cycles worst case}$$

For a 200 MHz clock and 8 clocks per cycle, this is 11.96 us worst case.

“Bouncing Ball” LED PWM Display Example

An interesting digital concept that seems rather unlikely to work in actual practice is the cross coupled integrator sine / cosine oscillator. The output of an accumulator is fed to the input of a second accumulator, the output of this second accumulator is inverted fed back to the input of the first. By identically scaling the accumulator inputs by a number generally much less than one (call it “alpha”) the frequency of the oscillations may be controlled or set. The first accumulator generates cosine, the second sine. It may be best seen as a ringing state variable band pass filter with infinite Q. Sine and cosine amplitude is set by placing an initial value in one accumulator and zeroing the other one out, and then letting it “ring” for infinity. Sine and cosine frequency \approx total cycle frequency * alpha / Pi.

Mathematically this construct works because the integral of cosine = sine, and the integral of sine = -cosine. Numerically this construct works only if there is a single register delay in the loop, and given this condition truncation errors due to integrator input scaling rather mysteriously don’t build up or otherwise become a long-term problem.

If we feed a sine wave to an absolute value circuit (invert the entire value if the sign bit is negative) and find the power of 2 of this value we can make a one-hot “bouncing ball” type LED display. To add smoothness we can interpret the absolute sine wave value as a fixed decimal, use the integer portion to nominally select the LED, and use the decimal portion to perform PWM (pulse width modulation). First order PWM is most easily accomplished by accumulating the PWM value and looking for accumulator overflow (smaller inputs cause infrequent overflows, and larger inputs cause more frequent overflows). When there is overflow we select the next higher LED. Rather than a flow diagram, a block diagram of this is shown below:

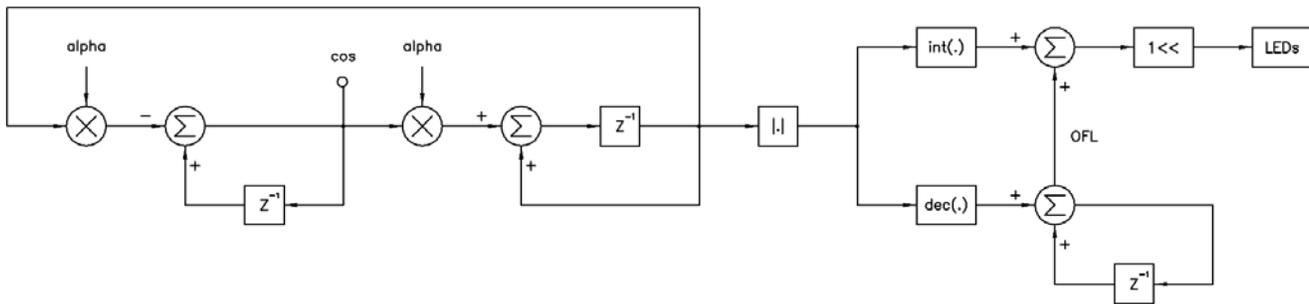


Figure 42. Bouncing ball block diagram.

In terms of design, we need to establish the frequency and amplitude of the sine wave, as well as fix the decimal place. The core clock is 160 MHz and dividing this by 8 clocks per cycle gives 20 MHz. Our loop takes 14 cycles to execute once, so to get a roughly 1 Hz bounce rate we need an alpha of $14 * 2 * \pi / 20 \text{ MHz} \approx 0.000044$ which is roughly 0x00003000. The Cyclone demo board I’m using has 4 LEDs, which correspond to the values 0, 1, 2, and 3, so if we set the most significant hex digit to be our integer portion, then the initialization value should be 0x30000000.

```

// "bouncing ball" 4 LED display w/ PWM
//
// s0 : sin
// s1 : cos
// s2 : alpha (attenuation factor = speed)
// s3 : rectified sin, val, one-hot(val)
// s4 :
// s5 : pwm counter
// s6 :
// s7 : i/o register address

i='h100; ram[i] = { `dat_is,          6'd0, `s0 }; // s0=0 (sin init)
i=i+1; ram[i] = { `lit_u,           `_, `s1 }; // s1=0x3000,0000 (cos init)
i=i+1; ram[i] = { `lit_u,           16'h3000 }; //
i=i+1; ram[i] = { `shl_is,          6'd16, `P1 }; //
i=i+1; ram[i] = { `lit_u,           `_, `s2 }; // s2=0x3000 (alpha init)
i=i+1; ram[i] = { `lit_u,           16'h3000 }; //
i=i+1; ram[i] = { `dat_is,          6'd0, `s5 }; // s5=0 (pwm init)
i=i+1; ram[i] = { `dat_is,          `IO_LO, `s7 }; // s7=reg addr
// loop start
// sin & cos
i=i+1; ram[i] = { `mul_xs,          `s2, `s0 }; // s0=s0*s2 (sin*alpha)
i=i+1; ram[i] = { `sub,             `P0, `P1 }; // s1-=s0 (cos-=sin*alpha)
i=i+1; ram[i] = { `mul_xs,          `s2, `s1 }; // s1=s1*s2 (cos*alpha)
i=i+1; ram[i] = { `add,             `P1, `P0 }; // s0-=s1 (sin+=cos*alpha)
// |sin|
i=i+1; ram[i] = { `cpy,             `s0, `s3 }; // s3=s0
i=i+1; ram[i] = { `jmp_inlz,        6'd1, `s3 }; // (s3!<0) ? jmp +1
i=i+1; ram[i] = { `not,             `s3, `P3 }; // s3~s3
// decimal( |sin| ) + pwm to update, + pwm to get ofl
i=i+1; ram[i] = { `shl_is,          6'd4, `s3 }; // s3=s3<<4
i=i+1; ram[i] = { `add,             `s3, `P5 }; // s5+=s3 (update pwm count)
i=i+1; ram[i] = { `add_xu,          `P3, `s5 }; // s5=s5+s3, pop s3 (get pwm ofl)
// one-hot( int( |sin| ) + pwm ofl )
i=i+1; ram[i] = { `shl_is,          -6'd28, `P3 }; // s3>>=28
i=i+1; ram[i] = { `add,             `P5, `P3 }; // s3+=s5, pop s5 (add pwm ofl)
i=i+1; ram[i] = { `pow,             `s3, `P3 }; // s3=1<<s3, pop s3 (one-hot)
// output
i=i+1; ram[i] = { `reg_w,           `s7, `P3 }; // (s7)=s3, pop s3
i=i+1; ram[i] = { `jmp_inz,        -6'd15, `s7 }; // loop forever
// loop end

```

Code is above. The initial values are loaded and the loop is entered. Cosine is updated first, then sine. The absolute value of sine is found and pushed to stack 3. This value is left shifted 4 places to obtain the decimal portion, which is added to the PWM counter to update it, after this it is added again but only to check for overflow. The absolute sine value is right shifted 28 places to obtain the integer portion, the overflow is added to it, the result of this converted to a power of 2, which is output to the LEDs.

There is a bit of cheating going on here. We really should check for PWM overflow *before* we update its value, but for slowly changing inputs the order isn't too important. Also, the total loop time varies by one cycle depending on the sign of the sine value (due to the negation jump) but this isn't noticeable even if you know about it.

Generic Boot Code

If one has a UART level shifter attached to one of the FPGA inputs (or the USB equivalent), the boot code may consist of a simple boot loader capable of uploading and storing executable code, and the boot code itself wouldn't need to be touched much after that. Simple scripting could be used to convert verilog boot text (or similar assembly) into uploadable binary data. A ~\$1 SPI FLASH device tacked onto a few spare FPGA pins could hold gobs of uploaded programming goodness.

BZZZ!

I would be completely remiss if I didn't point out the less positive aspects of Hive that I'm aware of:

- The instruction set of Hive was hatched more intuitively than scientifically by a person who doesn't exactly have loads of practical experience programming assembly (that probably just scared off most readers). I've put more time into selecting operations, sizing immediate fields, and shuffling things around in the opcode encoding space than I've spent actually programming Hive (at this point).
- Common data & instruction memory space (Von Neumann architecture) enables many good things, but it generally prevents code from executing directly from ROM, and it also makes it that much likelier for wild data writes to clobber code. A single thread caught out in the weeds means you should probably clear them all. (I should point out that the "Von Neumann bottleneck" is not an issue for Hive because it uses dual port BRAM for main memory.)
- With any stack machine, stack fullness is something the programmer must track carefully in order to avoid stack faults, and Hive has more stacks than usual to keep track of (though to be fair they are used in a simpler manner).
- Strict equal bandwidth multi-threading forces the programmer to implement some kind of load sharing arrangement for algorithms that require more real-time / less latency than a single thread can provide.
- Real-time response to an interrupt can be somewhat long and variable (though depending on the application this could perhaps be compensated for with additional interrupt time stamp & register set logic).
- FPGA logic will likely never be as fast / power efficient / inexpensive / etc. as an ASIC, so *any* soft processor core is in some sense a solution in search of a problem.

ETC.

- ☛ Hive was developed (including simulation / verification) with Altera Quartus II 9.1sp2 Web Edition (unfortunately the last edition with integrated simulator) running on WinXP Pro (sadly nearing the end of support).
- ☛ TextCrawler was used extensively to perform multi-file text search and replace (freeware from Digital Volcano).
- ☛ Pictures were drawn in AutoCAD 2006 (it is ironically nearly impossible to export good looking image files from AutoCAD) plotted to Adobe Generic PS printer (free from Adobe, good luck finding a suitable *.inf file) and rasterized with Paint Shop Pro X (pretty much broken in Win7/64).
- ☛ The Hive document was written in MS Word 2003 (also sadly nearing the end of support), and converted to PDF with Adobe Acrobat 8 Professional (which is free-ish due to license server retirement).
- ☛ There are inexpensive FPGA demo boards readily available on eBay. A very nice Cyclone 4 board can be had for \$27 USD or thereabouts. Comparable Xilinx Spartan offerings may require a few code changes here and there (I haven't run the code through ISE yet) and will likely run slower – Altera apparently uses faster, and consequently leakier and more power hungry, transistors.
- ☛ I'm working on a VBA MS Excel based simulator that will allow the user to easily enter and exercise code and spit out hex. I have the bulk of it up and running for the earlier 4 stack version of Hive but need to get on the stick and adapt it to v04.05.

Comments?

Found a bug in Hive (ha ha)? If you have questions, comments, criticisms, improvements, etc. regarding Hive I'd love to hear them! Contact me at: tammie_eric@verizon.net (note the '_' underscore).

DOCUMENT CHANGE CONTROL

(YYYY-MM-DD)	Hive version	Notes
2014-01-05	04.05	Major edits to reflect 8 stacks and somewhat different opcodes / resized immediates, UART, etc. Updated and expanded the coding examples.
2013-07-07	01.10	Edits to reflect code reshuffled opcodes. Fixed immediate add range on page 23. Added "barrel" processor classification and PDP 10 signed shift reference. Other sporadic minor edits.
2013-06-19	01.09	First public release.

COPYRIGHT

Copyright © Eric David Wallin, 2013 & 2014.

Permission to use, copy, modify, and/or distribute this design for any purpose without fee is hereby granted, provided it is not used for spying or "surveillance" uses, military or "defense" projects, weaponry, or other nefarious purposes. Furthermore the above copyright and this permission notice must appear in all copies.

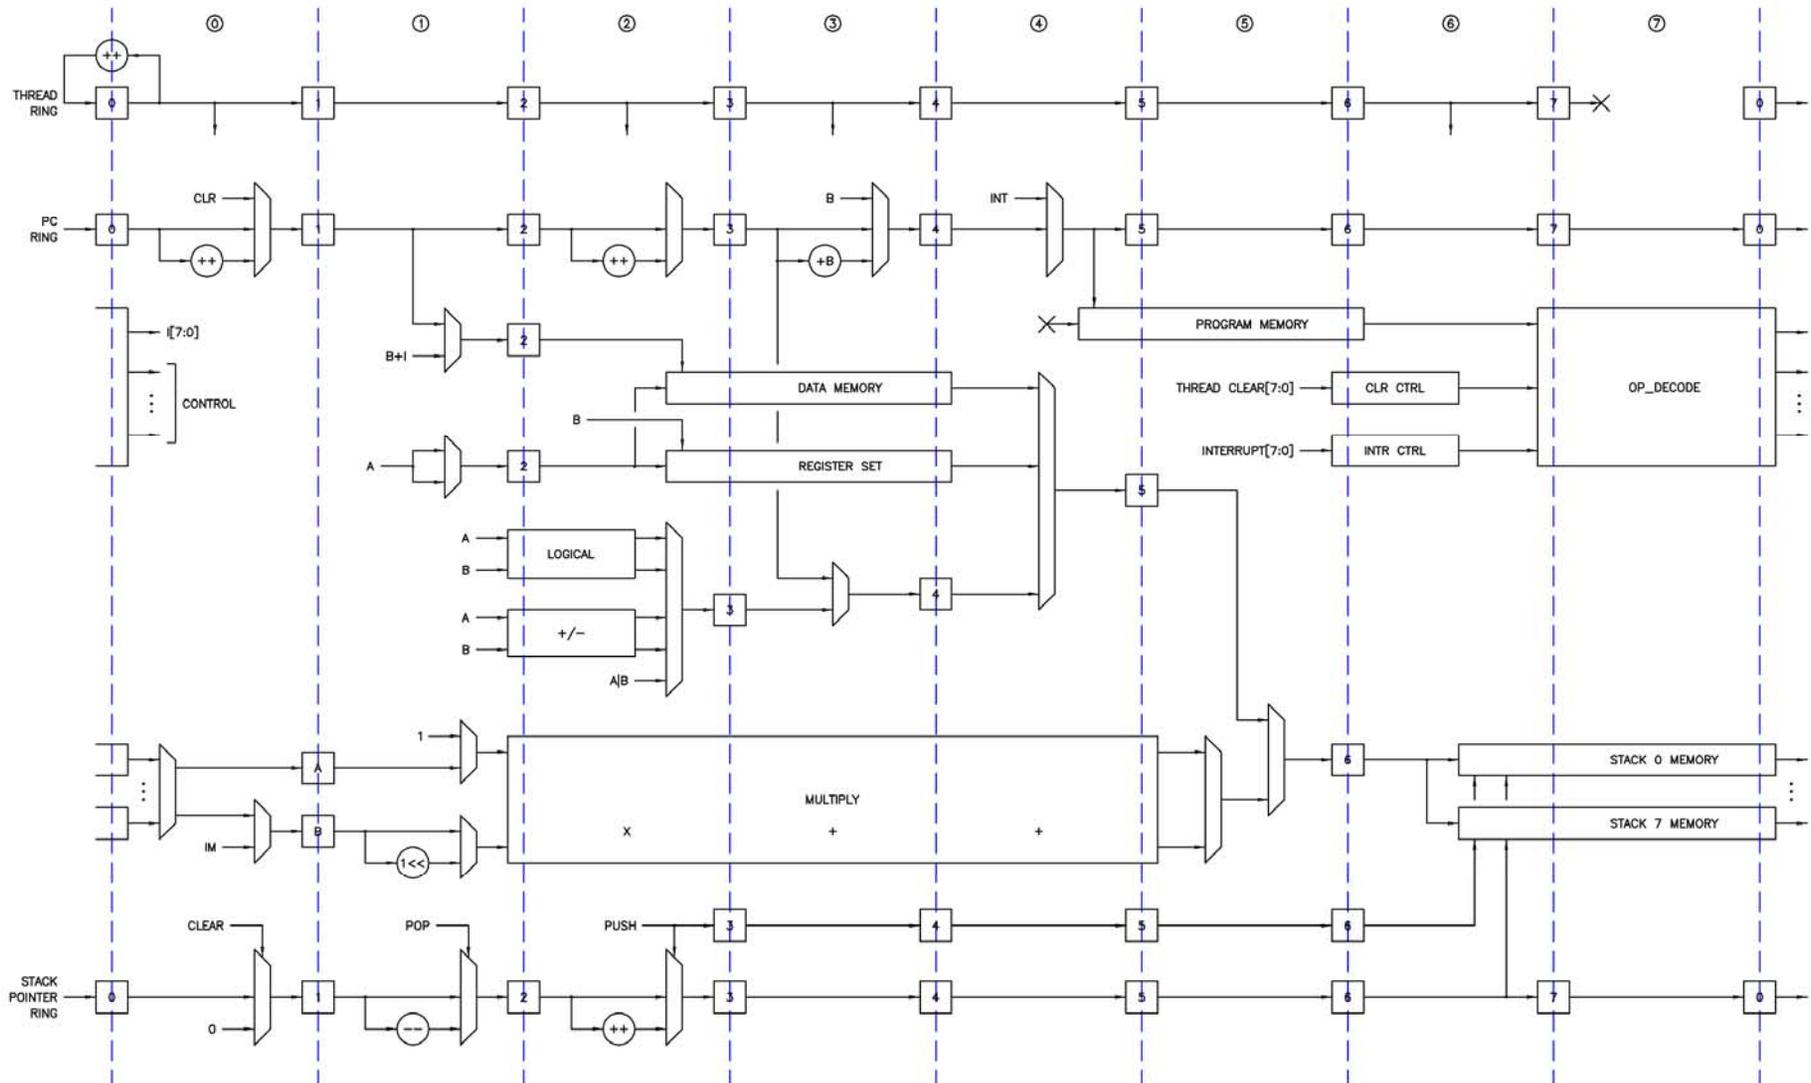


Figure 43. Hive core.

```

// misc - 16 x 16 x 16 = 4096 codes
localparam [OP_CODE_W-1:0] op_nop = { 4'h0, 4'h0, 4'hx, 4'hx }; // do nothing (no pops either)
localparam [OP_CODE_W-1:0] op_pop = { 4'h0, 4'h1, 4'hx, 4'hx }; // pop [7:0] none/one/some/all stacks
localparam [OP_CODE_W-1:0] op_pgc = { 4'h0, 4'h4, 4'hx, 4'hx }; // A=PC read PC (unsigned)
localparam [OP_CODE_W-1:0] op_lit_s = { 4'h0, 4'h8, 4'hx, 4'hx }; // A=mem(PC) literal data signed
localparam [OP_CODE_W-1:0] op_lit_h = { 4'h0, 4'h9, 4'hx, 4'hx }; // A={mem(PC),A[lo]} literal data high
localparam [OP_CODE_W-1:0] op_lit_u = { 4'h0, 4'ha, 4'hx, 4'hx }; // A=mem(PC) literal data unsigned
localparam [OP_CODE_W-1:0] op_reg_rs = { 4'h0, 4'hc, 4'hx, 4'hx }; // A=reg(B) register read signed
localparam [OP_CODE_W-1:0] op_reg_rh = { 4'h0, 4'hd, 4'hx, 4'hx }; // A={reg(B),A[lo]} register read high
localparam [OP_CODE_W-1:0] op_reg_w = { 4'h0, 4'he, 4'hx, 4'hx }; // reg(B)=A[lo] register write
localparam [OP_CODE_W-1:0] op_reg_wh = { 4'h0, 4'hf, 4'hx, 4'hx }; // reg(B)=A[hi] register write high
// logical & other - 16 x 16 x 16 = 4096 codes
localparam [OP_CODE_W-1:0] op_cpy = { 4'h1, `lg_cpy, 4'hx, 4'hx }; // A=B copy
localparam [OP_CODE_W-1:0] op_nsg = { 4'h1, `lg_nsg, 4'hx, 4'hx }; // A[MSB]=~B[MSB] invert sign
localparam [OP_CODE_W-1:0] op_not = { 4'h1, `lg_not, 4'hx, 4'hx }; // A=~B logical NOT
localparam [OP_CODE_W-1:0] op_flip = { 4'h1, `lg_flip, 4'hx, 4'hx }; // A=flip(B) flip bits end for end
localparam [OP_CODE_W-1:0] op_lzc = { 4'h1, `lg_lzc, 4'hx, 4'hx }; // A=lzc(B) leading zero count
localparam [OP_CODE_W-1:0] op_bra = { 4'h1, `lg_bra, 4'hx, 4'hx }; // A=&B logical AND bit reduction
localparam [OP_CODE_W-1:0] op_bro = { 4'h1, `lg_bro, 4'hx, 4'hx }; // A|=B logical OR bit reduction
localparam [OP_CODE_W-1:0] op_brx = { 4'h1, `lg_brx, 4'hx, 4'hx }; // A^=B logical XOR bit reduction
localparam [OP_CODE_W-1:0] op_and = { 4'h1, `lg_and, 4'hx, 4'hx }; // A=A&B logical AND
localparam [OP_CODE_W-1:0] op_orr = { 4'h1, `lg_orr, 4'hx, 4'hx }; // A=A|B logical OR
localparam [OP_CODE_W-1:0] op_xor = { 4'h1, `lg_xor, 4'hx, 4'hx }; // A=A^B logical XOR
// arithmetic - 16 x 16 x 16 = 4096 codes
localparam [OP_CODE_W-1:0] op_add = { 4'h2, 4'h0, 4'hx, 4'hx }; // A=A+B add
localparam [OP_CODE_W-1:0] op_add_xs = { 4'h2, 4'h2, 4'hx, 4'hx }; // A=A+B add extended signed
localparam [OP_CODE_W-1:0] op_add_xu = { 4'h2, 4'h3, 4'hx, 4'hx }; // A=A+B add extended unsigned
localparam [OP_CODE_W-1:0] op_sub = { 4'h2, 4'h4, 4'hx, 4'hx }; // A=A-B subtract
localparam [OP_CODE_W-1:0] op_sub_xs = { 4'h2, 4'h6, 4'hx, 4'hx }; // A=A-B subtract extended signed
localparam [OP_CODE_W-1:0] op_sub_xu = { 4'h2, 4'h7, 4'hx, 4'hx }; // A=A-B subtract extended unsigned
localparam [OP_CODE_W-1:0] op_mul = { 4'h2, 4'h8, 4'hx, 4'hx }; // A=A*B multiply
localparam [OP_CODE_W-1:0] op_mul_xs = { 4'h2, 4'ha, 4'hx, 4'hx }; // A=A*B multiply extended signed
localparam [OP_CODE_W-1:0] op_mul_xu = { 4'h2, 4'hb, 4'hx, 4'hx }; // A=A*B multiply extended unsigned
localparam [OP_CODE_W-1:0] op_shl_s = { 4'h2, 4'hc, 4'hx, 4'hx }; // A=A<<B shift left A signed
localparam [OP_CODE_W-1:0] op_shl_u = { 4'h2, 4'hd, 4'hx, 4'hx }; // A=A<<B shift left A unsigned
localparam [OP_CODE_W-1:0] op_pow = { 4'h2, 4'he, 4'hx, 4'hx }; // A=1<<B power of 2
// branching - 16 x 16 x 16 = 4096 codes
localparam [OP_CODE_W-1:0] op_jump_z = { 4'h3, 2'b00, `z, 4'hx, 4'hx }; // PC=(A?0)?PC+B jump zero conditional
localparam [OP_CODE_W-1:0] op_jump_nz = { 4'h3, 2'b00, `nz, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_lz = { 4'h3, 2'b00, `lz, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_nlz = { 4'h3, 2'b00, `nlz, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump = { 4'h3, 4'hc, 4'hx, 4'hx }; // PC=PC+B jump unconditional
localparam [OP_CODE_W-1:0] op_goto = { 4'h3, 4'hd, 4'hx, 4'hx }; // PC=B go to unconditional
localparam [OP_CODE_W-1:0] op_gsb = { 4'h3, 4'he, 4'hx, 4'hx }; // PC=B,A=PC subroutine call unconditional
// immediate memory access - 4 x 16 x 16 x 16 = 16384 codes
localparam [OP_CODE_W-1:0] op_mem_irs = { 4'h4, 4'hx, 4'hx, 4'hx }; // A=mem(B+I) memory read signed
localparam [OP_CODE_W-1:0] op_mem_irh = { 4'h5, 4'hx, 4'hx, 4'hx }; // A={mem(B+I),A[lo]} memory read high
localparam [OP_CODE_W-1:0] op_mem_iw = { 4'h6, 4'hx, 4'hx, 4'hx }; // mem(B+I)=A[lo] memory write
localparam [OP_CODE_W-1:0] op_mem_iwh = { 4'h7, 4'hx, 4'hx, 4'hx }; // mem(B+I)=A[hi] memory write high
// immediate conditional (A?B) jumps - 6 x 16 x 16 = 24576 codes
localparam [OP_CODE_W-1:0] op_jump_ie = { `e, 4'hx, 4'hx, 4'hx }; // PC=(A?B)?PC+I jump immediate conditional
localparam [OP_CODE_W-1:0] op_jump_ine = { `ne, 4'hx, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_ils = { `ls, 4'hx, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_inls = { `nls, 4'hx, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_ilu = { `lu, 4'hx, 4'hx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_inlu = { `nlu, 4'hx, 4'hx, 4'hx };
// immediate conditional (A?0) jumps - 4 x 64 x 16 = 4096 codes
localparam [OP_CODE_W-1:0] op_jump_iz = { 4'he, `z, 6'bxxxxxxx, 4'hx }; // PC=(A?0)?PC+I jump immediate conditional
localparam [OP_CODE_W-1:0] op_jump_inz = { 4'he, `nz, 6'bxxxxxxx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_ilz = { 4'he, `lz, 6'bxxxxxxx, 4'hx };
localparam [OP_CODE_W-1:0] op_jump_inlz = { 4'he, `nlz, 6'bxxxxxxx, 4'hx };
// immediate data - 1 x 64 x 16 = 1024 codes
localparam [OP_CODE_W-1:0] op_dat_is = { 4'hf, 2'b00, 6'bxxxxxxx, 4'hx }; // A=I data immediate signed
// immediate add - 1 x 64 x 16 = 1024 codes
localparam [OP_CODE_W-1:0] op_add_is = { 4'hf, 2'b01, 6'bxxxxxxx, 4'hx }; // A=A+I add immediate signed
// immediate shifts - 2 x 64 x 16 = 2048 codes
localparam [OP_CODE_W-1:0] op_shl_is = { 4'hf, 2'b10, 6'bxxxxxxx, 4'hx }; // A=A<<<I shift left A signed
localparam [OP_CODE_W-1:0] op_psu_i = { 4'hf, 2'b11, 6'bxxxxxxx, 4'hx }; // A=1<<<I power of 2; A=A<<<I shift A unsigned

```

Figure 44. Hive opcodes.