

Turtles all the Way Down: Running Linux on Open Hardware



Q: What parts of Linux Systems are Open and Under your Control?

- Modern laptops have more arm/mips processors than x86

- USB controller? Exploitable:

- <http://arstechnica.com/security/2014/07/this-thumbdrive-hacks-computer>
S
= badusb-exploit-makes-devices-turn-evil/

- Hard drive controller? Exploitable:

- http://hackaday.com/2013/08/02/sprite_tm-ohm2013-talk-hacking-hard-drive-controller-chips/


- System management mode, ACPI...

Taking Back Control of HW and SW



open source
hardware

What is the minimum system that can run Linux?



The definition of a 'CPU'

(for the purposes of this talk)

- Has a Flat 32bit memory space
 - All pointers 'just work' (no separate spaces)
- A port of GCC
 - Regular, efficient 32-16-8 bit int-short-char
- Executes instructions 'fast enough'
 - Servicing streaming Ethernet traffic: about 30MIPs

System Requirements

- CPU (of course)
- Memory
 - 8 megs RAM runs practical Linux workloads
 - less is possible, but awkward
- Storage (load kernel+initramfs into memory)
- Some form of I/O
 - Usually, and at first, just a UART
- A Timer (interrupt).

Things you don't need

- Anything else, actually
 - Video, audio, keyboard, persistent storage
- An MMU, or any 'fancy' CPU features.
 - FPU, SMP, even cache is optional

A History Lesson



The Linux/Microcontroller project is a port of Linux to systems without a Memory Management Unit (MMU).

uClinux first ported to the Motorola MC68328: DragonBall Integrated Microprocessor. The first target system to successfully boot is the [PalmPilot using a TRG SuperPilot Board with a custom boot-loader created specifically for our Linux/PalmPilot port](#)

•

(January 1998)

Why recreate existing architecture?

- Existing compiler, kernel, debugger, strace...
- Leverage massive R&D outlay in SuperH
 - 5 stage RISC (full Harvard architecture)
 - instruction set density (16 bit fixed length)
 - simple highly optimized design
 - original designers/implementers still around
- Old chips are prior art vs. "breathing is patented"

Patent Expiration

- SuperH ISA had a huge effort put into it
 - See above about efficient GCC.
- The SuperH architecture was developed by Hitachi a quarter century ago.
- Last patent on SH2 (Sega Saturn) expired in October 2014.
 - That's why we can release this now
 - last SH4 (Dreamcast) patent expires in 2016.
- SuperH ISA was the blueprint for ARM Thumb

Code Density : Efficiency

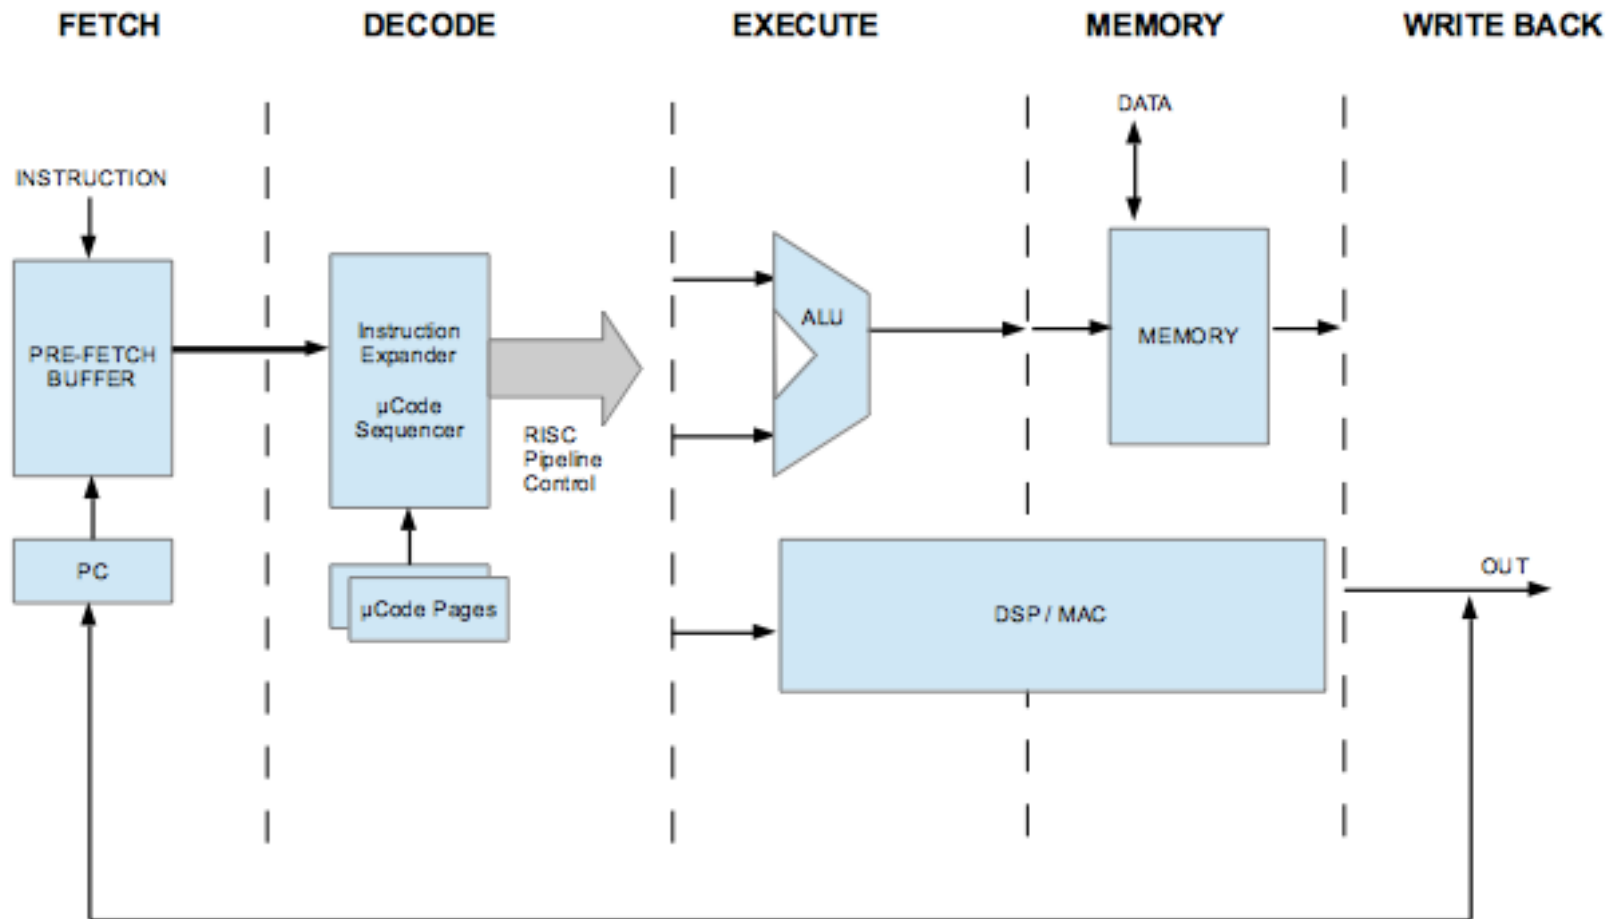
Fig. 2. Total size of benchmarks (includes some platform-specific code, so does not strictly reflect code density)



- There **Are** other metrics, but none actually matter more (unless something is broken)
 - Note: ARM paid millions to Hitachi to use SuperH patents in Thumb instruction set... which just expired.

The basic SuperH design:

5 stage 'classical' RISC pipeline, with some additions



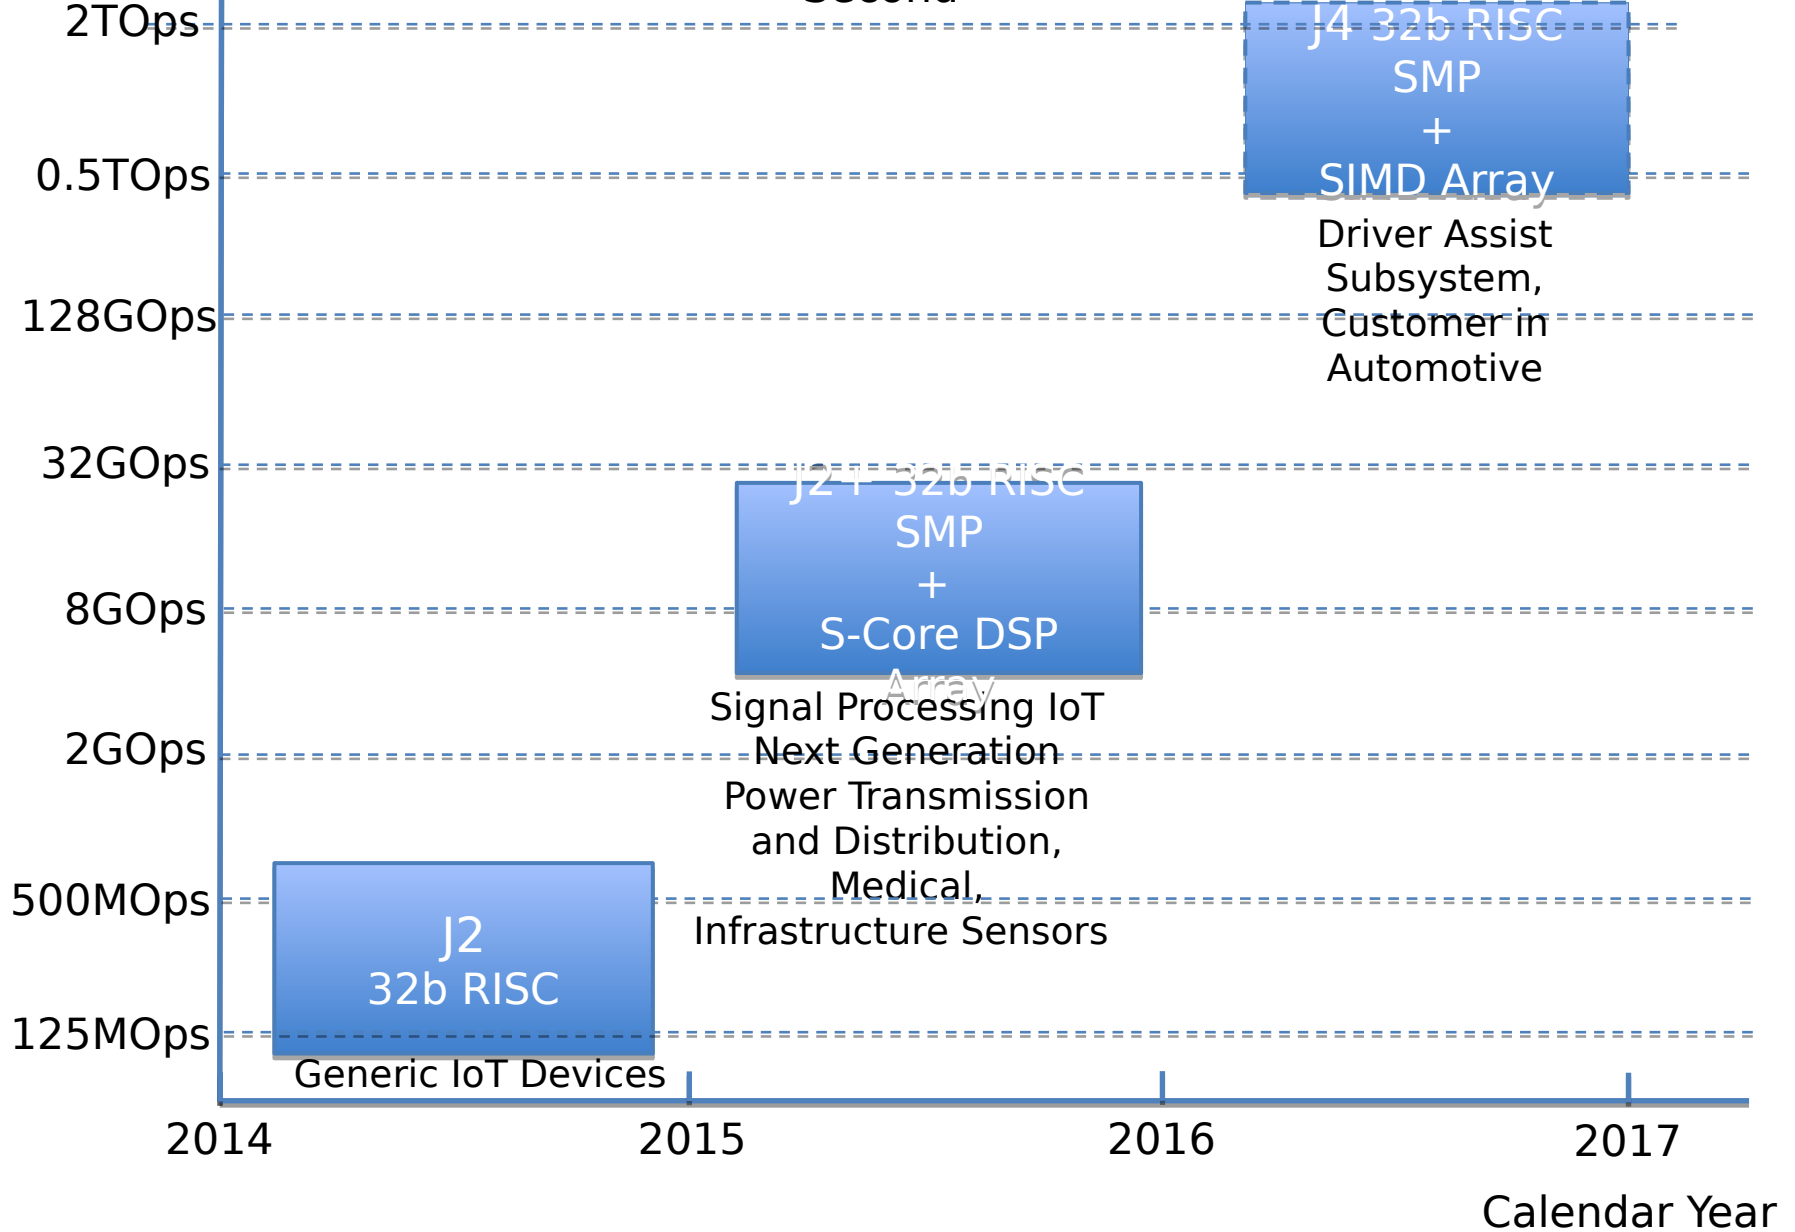
simple to implement, except the Instruction Decoder / Exp.

What did we do?

- New SH2 instruction set compatible core design kit
 - called "j2" because trademarks haven't expired.
 - clean-room implementation, initially by Canadian engineers
 - built for design reuse
 - Then hired SuperH architects to work on it afterwards
 - Source is VHDL using programming model developed by European Space Agency
 - Verilog is low level like assembly, VHDL is a HLL
 - Design Kit contains high level abstractions for future cores
- SOC builder system
 - Links together peripherals automatically
 - E.g. Serial, mmc, Ethernet
 - Can produce FPGA bitstream, ASIC RTL, C emulator source
 - SMP capable but not finished yet

J Series Computation Core Cluster Roadmap

Unit: Arithmetic Operations per
Second



J Series Computation Core Cluster Roadmap

Unit: Arithmetic Operations per
Second

2TOps
0.5TOps
128GOps
32GOps
8GOps
2GOps
500MOps
125MOps

2014

2015

2016

2017

Calendar Year

J2
32b RISC

Generic IoT Devices

J2+ 32b RISC
SMP
+
S-Core DSP
Array

Signal Processing IoT
Next Generation
Power Transmission
and Distribution,
Medical,
Infrastructure Sensors

J4 32b RISC
SMP
+
SIMD Array


Driver Assist
Subsystem,
Customer in
Automotive

First Device:

Smart Energy Instruments
IoT / Energy Management SoC
FPGA and ASIC implementation

Demo Platforms

- Our j2 processor core can run linux on low-cost FPGA Spartan6 platforms such as Numato Labs' mimas2, or Avnet's Lx9 microboard



- We are currently working to develop a custom development board with the same form factor as Raspberry Pi

So, how do you use it for anything?

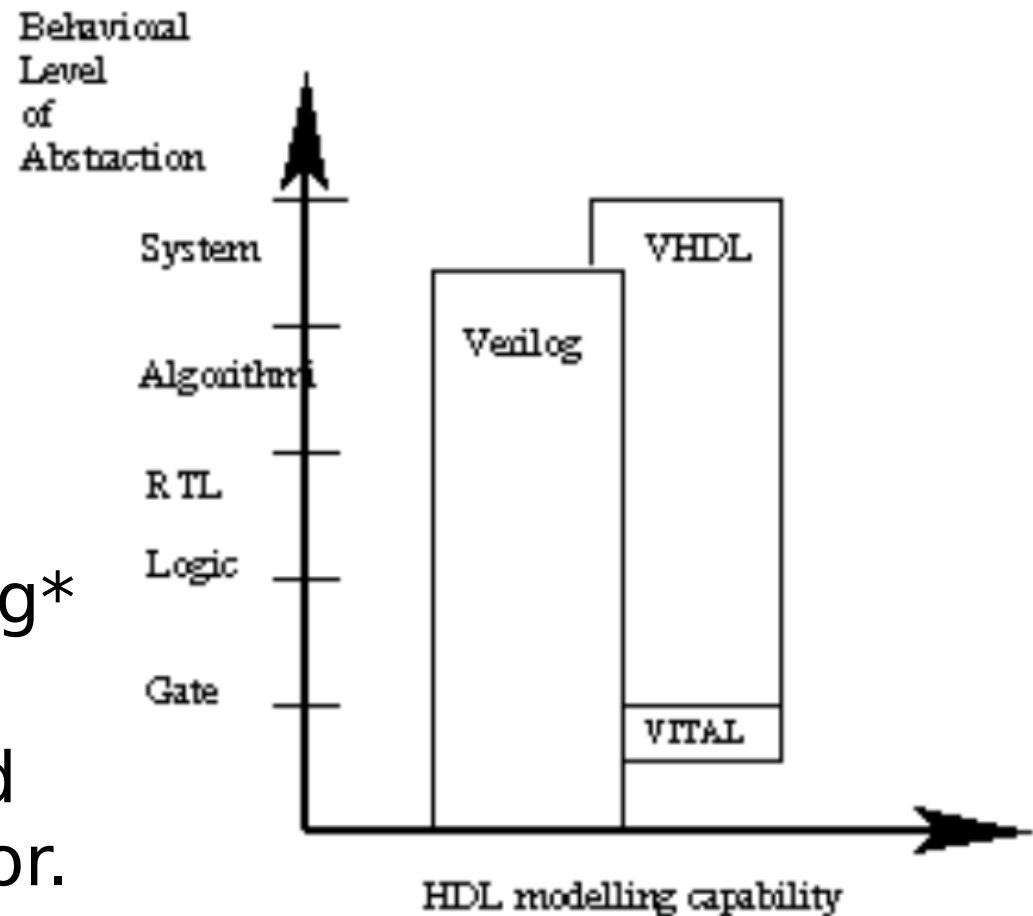
- Releasing VHDL and build system under BSD license
- Make any chip you want
 - Royalty free
 - 180nm ASIC of SOC we're demoing costs <10 yen
 - Processor only, about 2 and a half cents
 - Disposable computing at "free toy inside" level
 - Think IoT : 'Trillion Sensor Network' economics, but running Linux
- nommu.org (uclinux-ng), 0pf.org (zero-p-f)
 - Source, documentation, tutorials, mailing lists
 - We assume you've never done hardware before.
 - Still a bit sparse but filling sites out as we go
 - Patches welcome. No question too stupid.

How: VHDL, not Verilog

- Although Verilog is more commonly used in certain geographies..
- VHDL is the preferred language in developments initiated or led by the European Space Agency.
- The VHSIC Hardware Description Language (VHDL) is a formal notation intended for use in all phases of the creation of electronic systems. Because it is both machine readable and human readable, it supports the development, verification, synthesis, and testing of hardware designs, the communication of hardware design data, and the maintenance, modification, and procurement of hardware.

These 2 things are not the same...

- Actually...
- It's about the type system.
- Verilog don't have one (to speak of)
- In VHDL, everything* is a derived type.
- Even + is 'just' and overloaded operator.



Example code : Grey Vector Type

...

```
package gray_pack is
```

```
type gray_vector is array (natural range <>) of std_logic;
```

```
function "+" (L: gray_vector; R: integer) return gray_vector;
```

...

```
function "+" (L: gray_vector; R: integer) return gray_vector is  
variable res : gray_vector(L'range);
```

```
begin
```

```
    res := gray_vector(gr_inc(std_logic_vector(L)));
```

```
    return res;
```

```
end "+";
```

Example code : CPU Top Level

...

```
type cpu_instruction_i_t is record
    d    : std_logic_vector(15 downto 0);
    ack  : std_logic;
end record;
```

...

```
component cpu is port (
    clk          : in  std_logic;
    rst          : in  std_logic;
    db_o         : out cpu_data_o_t;
    db_i         : in  cpu_data_i_t;
    inst_o       : out cpu_instruction_o_t;
    inst_i       : in  cpu_instruction_i_t;
    debug_o      : out cpu_debug_o_t;
    debug_i      : in  cpu_debug_i_t;
    event_o      : out cpu_event_o_t;
    event_i      : in  cpu_event_i_t);
end component cpu;
```


Going further : Structured VHDL Design Method

- In order to overcome the limitations of the classical 'dataflow' design style (large number of concurrent VHDL statements and processes, leading to bad readability and increased simulation time), a '**two-process**' coding method is proposed: one process contains all combinational logic, whereas the other process infers all (and only) the registers.
- The use of record types to increase readability and the safe use of variables to reduce simulation time. The method has been applied on several designs made by or for ESA.

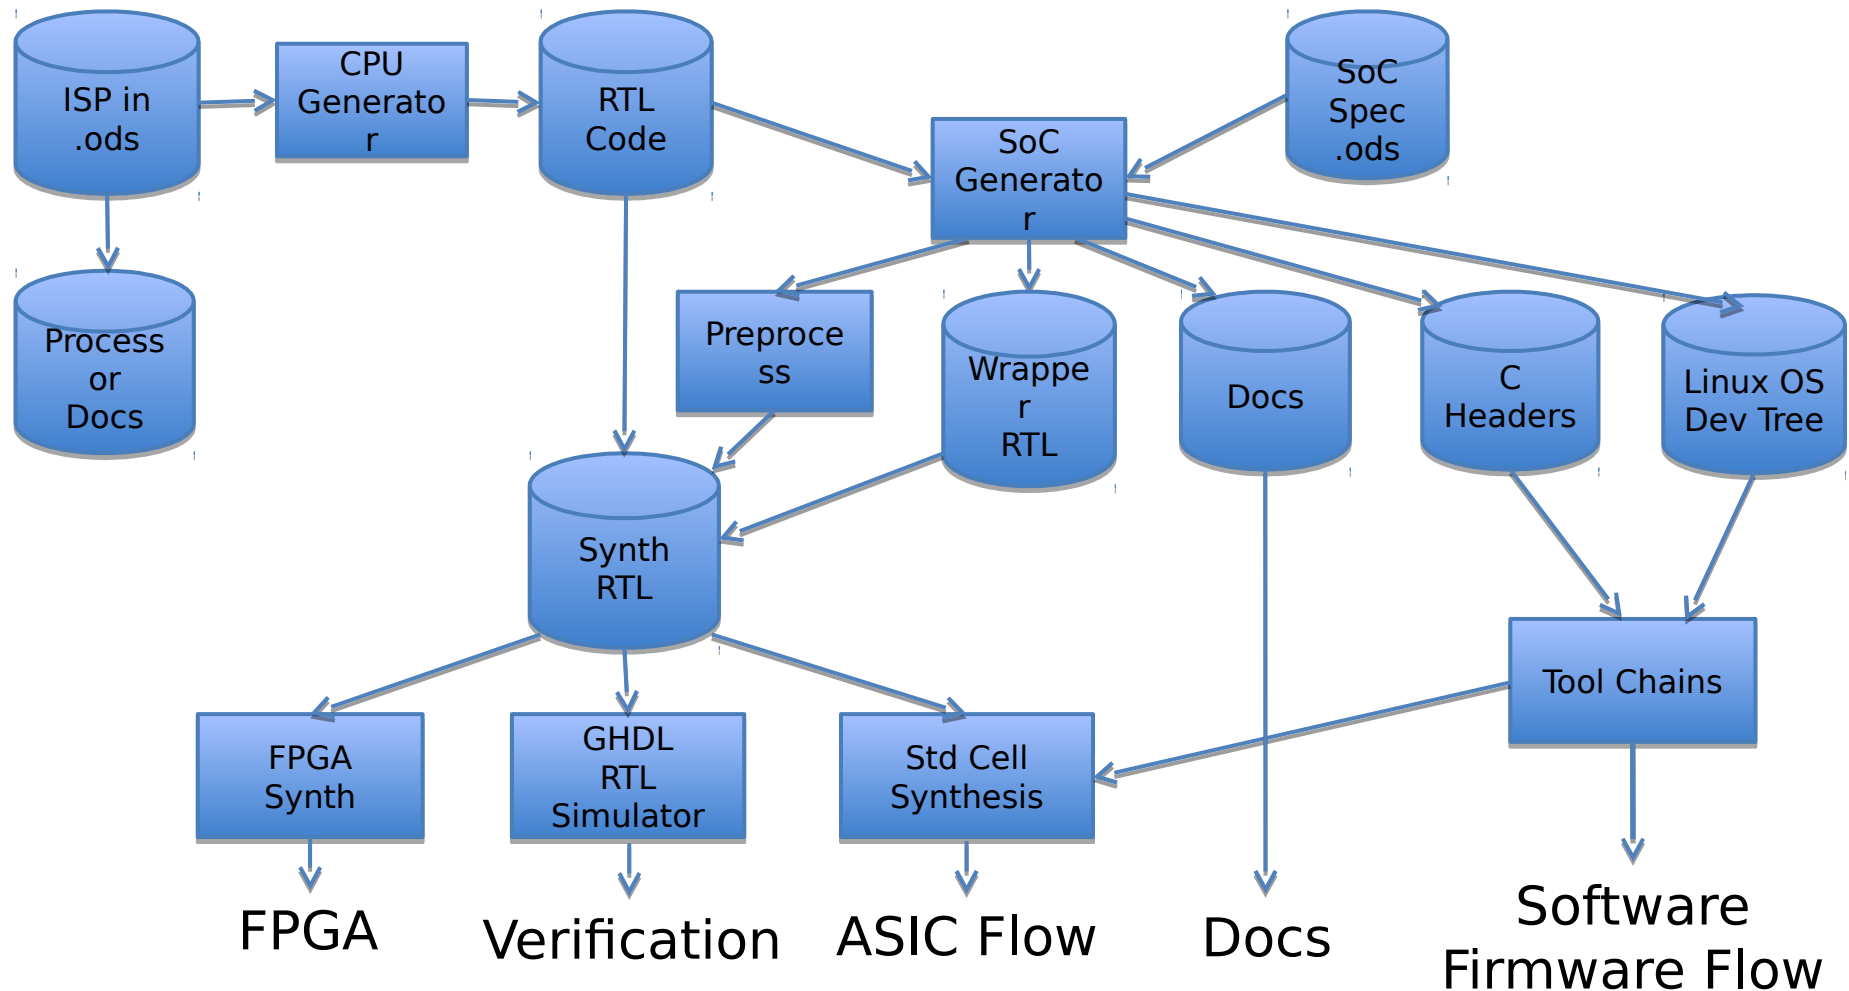
-Jiri Gaisler, <http://www.gaisler.com/doc/vhdl2proc.pdf>

Avoiding Common Errors

- We developed a pre-processor perl tool (v2p) to avoid latches, and other similar coding errors .
 - The `.vhm` file is a dialect of vhd; sensitivity lists are generated by the perl script
- This has resulted in highly reliable RTL, and greatly increased the efficiency of our internal development process

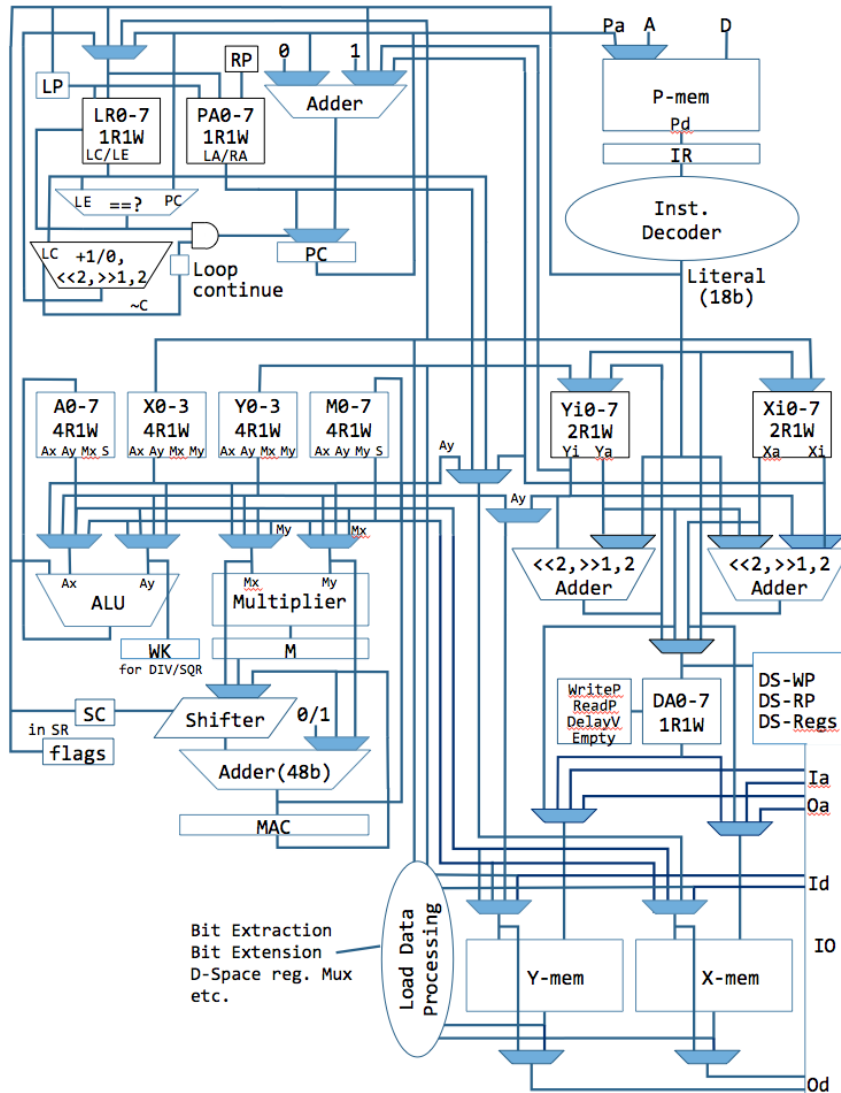


Even Further : Automated ISA -> RTL Sim->FPGA->ASIC->SW Tools Flow



18/36b S-Core DSP

- Development in Progress (Target Completion: August 2015)



DSP ISA and Code Map

ALU & MAC	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	primary set		
Xa	0	0																																					Xa0(x14),xc0(x16),x10,x12
Xd																																							Yd0(y14),yc0(y16),y10,y12
X																																							secondary set
Y																																							Xa1(x15),xc1(x17),x11
Y																																							Yd1(y15),yc1(x17),y11
Y																																							x13,y13: always "0"
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y																																							
Y															</																								

What can you do with this now?

- Download bitstream + vmlinux/initramfs, install on fpga board, boot kernel to shell prompt
- HOWTOs with background info
 - Where to order FPGA board(s)
 - Download and install xilinx/digilent tools
 - Free download for linux/mac, but alas no open source bitstream compiler yet. OpenOCD installer is todo item.
 - Build new bitstream from source
 - Program nommu Linux (gcc/binutils/musl toolchain)

Web Resources

- <http://0pf.org> - Open Processor Foundation
 - Yes, that's a zero
- <http://lists.nommu.org> - mailing lists
- <http://nommu.org> - nommu Linux docs
- <http://nommu.org/jcore> - board bringup
 - HOWTO buy fpga board, install hardware and software build environments, build and install, boot linux to shell prompt on serial console