

Reverse Engineering DSP Code

Pierre Bourdon

delroth@lse.epita.fr
<http://lse.epita.fr>

February 12, 2013

- Core developer of the Dolphin Emulator (GameCube/Wii)
- Recently working mainly on sound processing emulation
- Had to understand how it worked and reverse engineer the code running on it to reimplement it

- Digital Signal Processor
- Highly specialized CPUs with several ways to make signal processing fast
- Applications: sound mixing, sound effects processing, signal demodulation, etc.

- Mixing sounds together: $s = a + b$
- Setting a volume: $s = v \times i$ ($0 \leq v \leq 1$)
- You can only mix together sounds at the same sample rate, so resampling might be needed (linear, cubic, FIR)
- Sound delaying to simulate precise 3D positioning
- Filters: LPF, FIR, etc.

- Unless you need to cover a large range of values, floating point numbers are bad compared to fixed point numbers
- Sound samples are in $[-1.0, 1.0]$
- Volume is in $[0.0, 1.0]$
- Each sound sample can be represented as a 16 bit number in $[-32768, 32767]$
- Volume can be represented as a value in $[0, 32767]$
- Big optimization: ALU computations are a lot faster than FPU
- Need to be careful with overflows in intermediate computations

- The DSP also needs to communicate with several external components: CPU, RAM, hardware decoder, ...
- Often has interrupts and in/out ports support to get events from the CPU
- Data from RAM is fetched and written using DMA to an internal, smaller RAM

- Custom Macronix DSP design
- Runs at 81MHz (fast!)
- Hardware 32 bit multiplier with overflow handling
- 4K IRAM, 4K DRAM
- 4K IROM, 8K DROM
- DMA access to the GameCube RAM and ARAM
- Hardware PCM8, PCM16 and ADPCM decoding from ARAM

- 4 Address Registers: \$AR0, \$AR1, \$AR2, \$AR3
- 4 Index Registers: \$IX0, \$IX1, \$IX2, \$IX3
- 4 Wrapping Registers: \$WR0, \$WR1, \$WR2, \$WR3
- 2 32 bit "general" registers: \$AX0, \$AX1
- 2 40 bit accumulators: \$ACC0, \$ACC1
- 1 40 bit multiplication result register: \$PROD

- \$AX0.H, \$AX0.L (16 bit)
- \$ACC0.H (8 bit), \$ACC0.M, \$ACC0.L (16 bit)
- Access to \$ACC0.H can be either zero-extended or sign-extended

- 16 bit bytes: addresses index 16 bit values
- Instructions can be either 16 or 32 bits long (usually with a 16 bit immediate)
- Some instructions can be merged with an "extended operation" to perform 2 operations at once
- Strange control flow instructions using an internal loop register stack: LOOP, BLOOP, IFC, ...

```
CLR $ACC0           // ACC0 = 0;
LOOP $ACC1.M        // while (ACC1.M--)
SRRI @$AR0, $ACC0.M // *AR0++ = ACC0.M;
```

- Explicit parallelization of some operations that can be performed at the same time
- For example, "load from memory" and "multiply two numbers"
- Used a lot to make loops faster: load and store data at the same time you perform operations
- More than memory access: moving data from a register to another, adding an index register to an address register, etc.
- Uses parts of the CPU not used by the main instruction

opcode: bcf0

disasm:

```
MULAX'LD $AX0.H, $AX1.H, $ACC0 : $AX0.H, $AX1.H, @$AR0
```

pseudocode:

```
ACC0 += PROD;
```

```
PROD = AX0.H * AX1.H;
```

```
AX0.H = *AR0++;
```

```
AX1.H = *AR3++;
```

Example 2

opcode: f2e7

disasm:

```
MADD'LDN $AX0.L, $AX0.H : $AX0.H, $AX1.L, @$AR3
```

pseudocode:

```
$PROD += AX0.L * AX0.H;
```

```
AX0.H = *AR0++;
```

```
AX1.H = *AR3;
```

```
AR3 += IX3;
```

- Only one disassembler available, no real static analysis tool
- Wrote an IDA plugin for the GCN DSP in November 2011
- IDA handles surprisingly well most of the strange features of this DSP (including 16 bit bytes)
- Made it a lot easier to do cross-references, renaming symbols, etc.
- Writing IDA plugins will make you hate it, but it's worth the trouble in the end

Reverse Engineering DSP Code

Pierre Bourdon

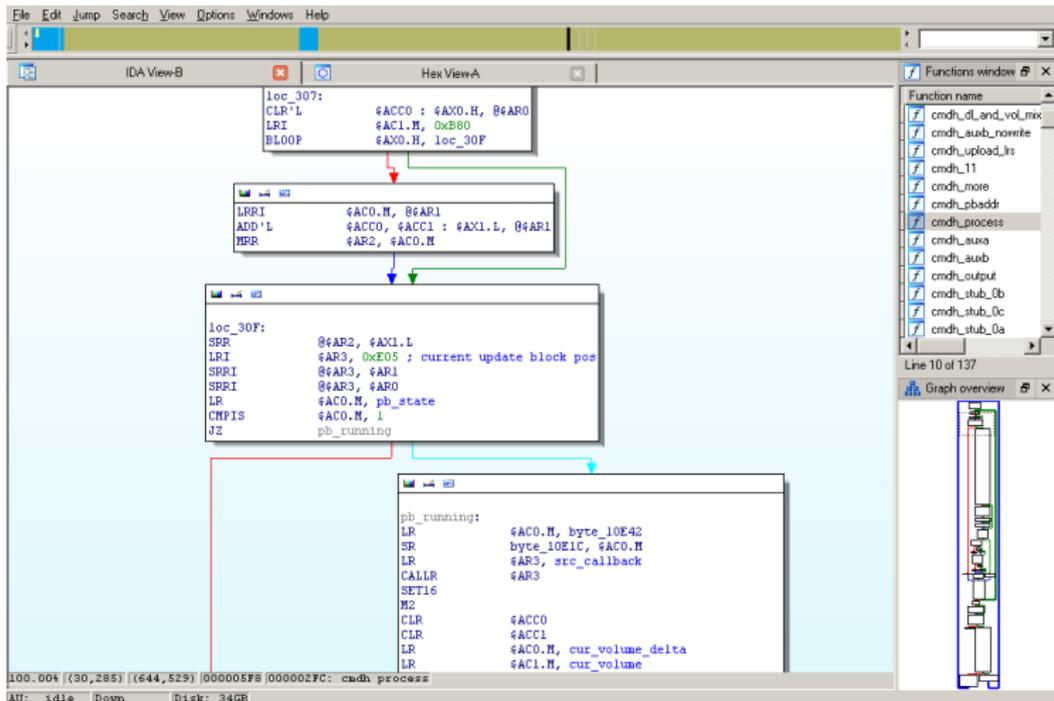
Introduction

DSP

GameCube DSP

Analyzing GCN DSP code

Conclusion



File Edit Jump Search View Options Windows Help

IDA View-B Hex View-A Functions window

```

loc_307:
CLR'L    %ACCO : %AX0.H, @%AR0
LRI      %ACL1.H, 0xB80
BLOOP   %AX0.H, loc_30F

loc_30F:
SRR      @%AR2, %AX1.L
LRI      %AR3, 0xE05 ; current update block pos
SRR1     @%AR3, %AR1
SRR1     @%AR3, %AR0
LR       %ACO.M, pb_state
CHPIS   %ACO.M, 1
JZ      pb_running

pb_running:
LR       %ACO.M, byte_10E42
SR       byte_10E1C, %ACO.M
LR       %AR3, src_callback
CALLER  %AR3
SET16
M2
CLR      %ACCO
CLR      %ACC1
LR       %ACO.M, cur_volume_delta
LR       %ACL1.M, cur_volume
  
```

Function name

- cmdh_d_end_vol_mix
- cmdh_sub_nwrite
- cmdh_upload_irq
- cmdh_11
- cmdh_more
- cmdh_sbaddr
- cmdh_process
- cmdh_sava
- cmdh_sub
- cmdh_output
- cmdh_stub_0b
- cmdh_stub_0c
- cmdh_stub_0a

Line 10 of 137

Graph overview

100.004 (30,285) (644,529) 000005F8 000002FC: cmdh_process

AU: idle Down Disk: 34GB

- All of the code is written directly in assembly, without respect for any kind of calling convention
- Branching has an impact on speed, so loops are sometimes manually unrolled
- Wrapping registers used to implement circular buffers
- Automatic multiply by 2 for volume mixing

Used to increase the throughput of a loop by taking advantage of the explicit parallelization.

```
LRRI          $AX0.H, @$AR3
LRRI          $AX0.L, @$AR3
MULX         $AX0.L, $AX1.L
MULXMV       $AX0.H, $AX1.L, $ACCO
BLOOPI       0x30, 0x0655
    ASR16'L   $ACCO : $AC1.M, @$AR1
    ADDP'LN   $ACCO : $AC1.L, @$AR1
    LRRI      $AX0.H, @$AR3
    ADD'L     $ACC1, $ACCO : $AX0.L, @$AR3
    MULX'S    $AX0.L, $AX1.L : @$AR1, $AC1.M
    MULXMV'S  $AX0.H, $AX1.L, $ACCO : @$AR1, $AC1.L
```

- @delroth_
- <http://dolphin-emu.org/>